

DOCUMENT RESUME

ED 289 738

SE 048 823

AUTHOR Sheridan, Thomas B., Ed.; And Others
TITLE Human Factors in Automated and Robotic Space Systems: Proceedings of a Symposium (Washington, D.C., January 29-30, 1987).

INSTITUTION National Academy of Sciences - National Research Council, Washington, DC. Commission on Behavioral and Social Sciences and Education.

SPONS AGENCY National Aeronautics and Space Administration, Washington, D.C.

PUB DATE 87
CONTRACT NASW-4071
NOTE 465p.

AVAILABLE FROM Committee on Human Factors, Commission on Behavioral and Social Sciences and Education, National Research Council, 2101 Constitution Ave., NW, Washington, DC 20418.


PUB TYPE Collected Works - Conference Proceedings (021)

EDRS PRICE MF01/PC19 Plus Postage.
DESCRIPTORS *Artificial Intelligence; Cognitive Processes; Computers; Computer Software; *Cybernetics; Decision Making; Expert Systems; Information Technology; Intelligence; *Man Machine Systems; Research; *Robotics; *Space Exploration; *Space Sciences; Technological Advancement

IDENTIFIERS *National Aeronautics and Space Administration; Space Colonization; Space Travel

ABSTRACT

This document attempts to identify and promote human factors research that would likely produce results applicable to the evolutionary design of a National Aeronautics and Space Administration (NASA) national space station to be launched in the 1990s. It reports on a symposium designed to yield information applicable to future space systems. The opening session includes a keynote address by Allen Newell on "Human Factors Research for the NASA Space Station." Session one, "System Productivity: People and Machines," focuses on productivity in the space station. Session two, "Expert Systems and Their Use," discusses artificial intelligence systems in the space station and their uses. Session three, "Language and Displays for Human-Computer Communication," addresses some changes needed in human-computer interfaces and cognitive factors in software design and development. Session four, "Computer Aided Monitoring and Decision Making," discusses robustness and transparency of intelligence systems. Session five, "Telepresence and Supervisory Control," deals with research needs for teleoperation, telepresence and telerobotics for the evolving space station. Session six, "Social Factors in Productivity and Performance," includes discussions about social issues involved with life in a space station. Session seven, "The Human Role in Space Systems," discusses human roles and sharing tasks between people and computers. (CW)



U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

This document has been reproduced as received from the person or organization originating it.

Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

ED289738

Human Factors in Automated and Robotic Space Systems

Proceedings of a Symposium

NASA

HUMAN FACTORS IN
AUTOMATED AND ROBOTIC SPACE SYSTEMS:
PROCEEDINGS OF A SYMPOSIUM

Thomas B. Sheridan, Dana S. Kruser, and
Stanley Deutsch, editors

Committee on Human Factors
Commission on Behavioral and Social Sciences and Education
National Research Council

Washington, D.C. 1987

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

This report has been reviewed by a group other than the authors according to procedures approved by a Report Review Committee consisting of members of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine.

The National Academy of Sciences is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Frank Press is president of the National Academy of Sciences.

The National Academy of Engineering was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. Robert M. White is president of the National Academy of Engineering.

The Institute of Medicine was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Samuel O. Thier is president of the Institute of Medicine.

The National Research Council was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Frank Press and Dr. Robert M. White are chairman and vice chairman, respectively, of the National Research Council.

This work relates to the National Aeronautics and Space Administration Contract Number NASW-4071. However, the content does not necessarily reflect the position or the policy of the government, and no official endorsement should be inferred.

The United States government has at least a royalty-free, nonexclusive and irrevocable license throughout the world for government purposes to publish, translate, reproduce, deliver, perform, dispose of, and to authorize others so as to do, all or any portion of this work.

Available from:
Committee on Human Factors
Commission on Behavioral and Social Sciences and Education
National Research Council
2101 Constitution Avenue, N.W.

SYMPOSIUM STEERING GROUP

THOMAS B. SHERIDAN (Chair), Department of Engineering and Applied
Psychology and Department of Mechanical Engineering, Massachusetts
Institute of Technology

OSCAR GRUSKY, Department of Sociology, University of California, Los
Angeles

ALLEN NEWELL, Computer Science Department, Carnegie-Mellon University

JUDITH REITMAN OLSON, Graduate School of Business Administration,
University of Michigan

ROBERT C. WILLIGES, Department of Industrial Engineering and
Operations Research, Virginia Polytechnic Institute and State
University

HARRY L. WOLBERS, Space Station Program, McDonnell Douglas
Aeronautics Company, Huntington Beach, California

STANLEY DEUTSCH, Study Director

DANA S. KRUSER, Project Coordinator/Consultant

BEVERLY M. HUEY, Research Associate/Consultant

ELIZABETH F. NEILSEN, Research Assistant

MARGARET A. CHENG, Administrative Secretary

COMMITTEE ON HUMAN FACTORS

1986-1987

THOMAS B. SHERIDAN (Chair), Department of Engineering and Applied Psychology and Department of Mechanical Engineering, Massachusetts Institute of Technology

NANCY S. ANDERSON, Department of Psychology, University of Maryland

CLYDE H. COOMBS, Department of Psychology, University of Michigan

JEROME I. ELKIND, Information Systems, Xerox Corporation, Palo Alto, California

OSCAR GRUSKY, Department of Sociology, University of California, Los Angeles

ROBERT M. GUION, Department of Psychology, Bowling Green State University

DOUGLAS H. HARRIS, Anacapa Sciences, Santa Barbara, California

JULIAN HOCHBERG, Department of Psychology, Columbia University,

THOMAS K. LANDAUER, Information Sciences Division, Bell Communication Research, Morristown, New Jersey

JUDITH REITMAN OLSON, Graduate School of Business Administration, University of Michigan

RICHARD W. PEW, Computer and Information Sciences Division, Bolt Beranek and Newman Laboratories, Inc., Cambridge, Massachusetts

STOVER H. SNOOK, Liberty Mutual Insurance Company, Hopkinton Research Center, Hopkinton, Massachusetts

CHRISTOPHER I. WICKENS, Aviation Research Institute, University of Illinois

ROBERT C. WILLIGES, Department of Industrial Engineering and Operations Research, Virginia Polytechnic Institute and State University

HAROLD P. VAN COTT, Study Director

STANLEY DEUTSCH, Study Director (1984-1987)

SYMPOSIUM PARTICIPANTS

- THOMAS B. SHERIDAN (Chair), Department of Engineering and Applied Psychology and Department of Mechanical Engineering, Massachusetts Institute of Technology
- DAVID L. AKIN, Assistant Professor of Aeronautics and Astronautics, Massachusetts Institute of Technology
- ANTAL K. BEJCZY, Telerobot Research Program, Jet Propulsion Laboratory, Pasadena, California
- BRUCE G. BUCHANAN, Knowledge Systems Laboratory, Stanford University
- RAYMOND S. COLLADAY, Office of Aeronautics and Space Technology, National Aeronautics and Space Administration
- KAREN S. COOK, Department of Sociology, University of Washington
- RANDALL DAVIS, Associate Professor, Sloan School of Management, Massachusetts Institute of Technology
- BARUCH FISCHHOFF, Department of Engineering and Public Policy, Carnegie-Mellon University
- OSCAR GRUSKY, Department of Sociology, University of California, Los Angeles
- PHILIP J. HAYES, Natural Language Systems, Carnegie Group, Inc., and Carnegie-Mellon University
- WILLIAM C. HOWELL, Department of Psychology, Rice University
- H. ANDREW MICHENER, Department of Sociology, University of Wisconsin
- THOMAS M. MITCHELL, Computer Science Department, Carnegie-Mellon University
- ALLEN NEWELL, Computer Science Department, Carnegie-Mellon University
- RAYMOND S. NICKERSON, Bolt Beranek and Newman Laboratories, Inc., Cambridge, Massachusetts
- JUDITH REITMAN OLSON, Graduate School of Business Administration, University of Michigan
- PETER G. POLSON, Computer Laboratory for Instruction in Psychological Research, University of Colorado
- WILLIAM H. STAREBUCK, Department of Creative Management, New York University
- LAWRENCE STARK, Telerobotics Unit, University of California, Berkeley
- ROBERT C. WILLIGES, Department of Industrial Engineering and Operations Research, Virginia Polytechnic Institute and State University
- HARRY L. WOLBERS, Manned & EVA Systems, Space Station Program, McDonnell Douglas Astronautics Company, Huntington Beach, California

FOREWORD

The Committee on Human Factors was established in October 1980 by the Commission on Behavioral and Social Sciences and Education of the National Research Council. It is sponsored by the Office of Naval Research, the Air Force Office of Scientific Research, the Army Research Institute for the Behavioral and Social Sciences, the National Aeronautics and Space Administration, and the National Science Foundation.

The principal objectives of the committee are to provide new perspectives on theoretical and methodological issues, identify basic research needed to expand and strengthen the scientific basis of human factors, and to attract scientists both inside and outside the field to perform needed research. The goal of the committee is to provide a solid foundation of research on which effective human factors practices can build.

In order for the committee to perform its role effectively, it draws on experts from a wide range of scientific and engineering disciplines. The committee includes specialists in the fields of psychology, engineering, biomechanics, cognitive sciences, machine intelligence, computer sciences, sociology, and human factors engineering. Participants in the working groups, workshops, and symposia organized by the committee represent additional disciplines. All of these disciplines contribute to the basic data, theory, and methods required to improve the scientific basis of human factors.

CONTENTS

PREFACE	xiii
SYMPOSIUM SUMMARY	1

OPENING SESSION

WELCOME	13
Thomas B. Sheridan	
INTRODUCTION	15
Raymond S. Colladay	
KEYNOTE ADDRESS: HUMAN FACTORS RESEARCH FOR THE NASA SPACE STATION	17
Allen Newell	

SESSION I: SYSTEM PRODUCTIVITY: PEOPLE AND MACHINES

PRODUCTIVITY IN THE SPACE STATION	31
Raymond S. Nickerson	
DISCUSSION: COMMENTS ON SYSTEM PRODUCTIVITY: PEOPLE AND MACHINES	82
Robert C. Williges	
SYNOPSIS OF GENERAL AUDIENCE DISCUSSION	87

SESSION II: EXPERT SYSTEMS AND THEIR USE

AI SYSTEMS IN THE SPACE STATION	91
Thomas M. Mitchell	
EXPERT SYSTEMS: APPLICATIONS IN SPACE	113
Bruce G. Buchanan	
DISCUSSION: COMMENTS ON EXPERT SYSTEMS AND THEIR USE	142
Allen Newell	
SYNOPSIS OF GENERAL AUDIENCE DISCUSSION	147

SESSION III: LANGUAGE AND DISPLAYS FOR
HUMAN-COMPUTER COMMUNICATION

CHANGE IN HUMAN-COMPUTER INTERFACES ON THE SPACE STATION: WHY IT NEEDS TO HAPPEN AND HOW TO PLAN FOR IT Philip J. Hayes	151
COGNITIVE FACTORS IN THE DESIGN AND DEVELOPMENT OF SOFTWARE IN THE SPACE STATION Peter G. Polson	176
DISCUSSION: DESIGNING FOR THE FACE OF THE FUTURE: RESEARCH ISSUES IN HUMAN-COMPUTER INTERACTION Judith Reitman Olson	201
SYNOPSIS OF GENERAL AUDIENCE DISCUSSION	208
SESSION IV: COMPUTER-AIDED MONITORING AND DECISION MAKING	
ROBUSTNESS AND TRANSPARENCY IN INTELLIGENT SYSTEMS Randall Davis	211
DECISION MAKING--AIDED AND UNAIDED Baruch Fischhoff	234
DISCUSSION: ISSUES IN DESIGN FOR UNCERTAINTY William C. Howell	263
SYNOPSIS OF GENERAL AUDIENCE DISCUSSION	275
SESSION V: TELEPRESENCE AND SUPERVISORY CONTROL	
TELEOPERATION, TELEPRESENCE, AND TELEROBOTICS: RESEARCH NEEDS FOR SPACE Thomas B. Sheridan	279
TELEROBOTICS FOR THE EVOLVING SPACE STATION: RESEARCH NEEDS AND OUTSTANDING PROBLEMS Lawrence Stark	292
DISCUSSION: COMMENTS ON TELEPRESENCE AND SUPERVISORY CONTROL Antal K. Bejczy	320
SYNOPSIS OF GENERAL AUDIENCE DISCUSSION	323

SESSION VI: SOCIAL FACTORS IN PRODUCTIVITY AND PERFORMANCE

SOCIAL STRESS, COMPUTER-MEDIATED COMMUNICATION SYSTEMS, AND HUMAN PRODUCTIVITY IN SPACE STATIONS: A RESEARCH AGENDA Karen S. Cook	329
CONTROL, CONFLICT, AND CRISIS MANAGEMENT IN THE SPACE STATION'S SOCIAL SYSTEM (YEAR 2000) H. Andrew Michener	356
DISCUSSION: CONFLICT AND STRESS IN THE SPACE STATION Oscar Grusky	390
SYNOPSIS OF GENERAL AUDIENCE DISCUSSION	402

SESSION VII: THE HUMAN ROLE IN SPACE SYSTEMS

THE ROLES OF HUMANS AND MACHINES IN SPACE David L. Akin	405
SHARING COGNITIVE TASKS BETWEEN PEOPLE AND COMPUTERS IN SPACE SYSTEMS William H. Starbuck	418
DISCUSSION: COMMENTS ON THE HUMAN ROLE IN SPACE SYSTEMS Harry L. Wolbers	444
SYNOPSIS OF GENERAL AUDIENCE DISCUSSION	451

CONCLUSION

CONCLUDING REMARKS Allen Newell	455
CONCLUDING REMARKS Thomas B. Sheridan	457
APPENDIX: SYMPOSIUM PROGRAM	463

PREFACE

A steering group formed by the Committee on Human Factors was charged to identify the types of human factors research that, if funded and begun immediately, would be likely to produce results applicable to the evolutionary design of a National Aeronautics and Space Administration national space station to be launched in the 1990s. The steering group was instructed to consider human factors research relevant to such future space systems as the space station, lunar bases, and possibly interplanetary travel. The symposium, which was planned by the steering group and is reported in these proceedings, did indeed yield information applicable to future space systems. In addition, it provided information and offered insights of potential interest to many other civilian and military endeavors. It was our hope that this potential for transfer would occur.

I would like to thank the participants in this project for their time, effort, and contributions to the symposium. Individual authors accept primary responsibility for each paper and this authorship is acknowledged at the beginning of each paper. Steering group members deliberated, reviewed, and contributed to improvements in the content of each paper. I am especially grateful to them for their generous contribution of time both before and after the symposium.

The steering group, and the other principals in the production of this symposium, received a great deal of guidance and assistance from NASA personnel. On behalf of us all, I would like to thank Melvin Montemerlo and Michael McGreevy of the Office of Aeronautics and Space Technology, Richard Carlisle and Bryant Cramer of the Space Station Office, and Owen Garriott, astronaut, for their extensive summaries of the space station planning activities during the initial October 1985 steering group briefings. Special thanks are also due to Jesse Moore, the director of the Johnson Space Center, Joseph Loftus, assistant director, David Nagel from the Ames Research Center, and the many NASA personnel who participated in the briefings of the steering committee held at the Johnson Space Center.

Finally, thanks are due to the people who have worked behind the scenes to ensure that the symposium was conducted, and the proceedings prepared, in an organized and timely manner. Appreciation is extended to Stanley Deutsch, study director at the time of the symposium, for his contributions to its planning; to Dana Kruser, project coordinator,

for her efforts in the organization and execution of the symposium and assistance in the editing of this report; to Elizabeth Neilsen, research assistant, for her managerial and logistic support; to Beverly Huey, who also provided logistic support; to Christine McShane, of the Commission staff, for editorial support; to Margaret Cheng, who provided secretarial assistance in preparation for the symposium; to Marian Holtzthum, for secretarial assistance in preparing this document for review; and to Martha Seijas, for preparing the document for publication. I express my sincere thanks to each of these individuals for their significant contributions.

Thomas B. Sheridan, Chair
Committee on Human Factors

SYMPOSIUM SUMMARY

We can follow our dreams to distant stars, living and working in space for peaceful economic and scientific gain. Tonight, I am directing NASA to develop a permanently manned Space Station and to do it within a decade.

President Ronald Reagan, State of the Union Message,
January 5, 1984.

In response to this presidential mandate, the National Aeronautics and Space Administration (NASA) is planning to launch a national space station in the early 1990s. To implement this commitment, and in concurrence with a congressional mandate, NASA is focusing serious attention on the use of automation and robotics in future space systems.

There is a tendency, particularly in the public sector, to view the emergence of new computer capabilities and automation and robotic technologies as a basis for replacing humans in space and thereby avoiding tragedies such as those of the Apollo 7 and the Challenger. However, it is unlikely that artificial intelligence comparable to human intelligence will be available to replace humans during the last part of the twentieth Century and the early part of the twenty-first. Therefore, people and automated systems will work together in space for the foreseeable future.

NASA is planning new research programs aimed at acquiring a better understanding of how computers, automation, and robotics can be made to work in partnership with people in complex, long-duration space system missions. These programs will address important questions concerning the relationship between what are called intelligent systems and the people who will use them as astronauts inside a space vehicle and in extravehicular activities, as scientists and technicians in space and on the ground, and as controllers on the ground.

Space offers significant challenges for the exploration and demonstration of human-computer-robot cooperation. Recognizing the size, complexity, and importance of this challenge, the Aeronautics and Space Technology Office approached the Committee on Human Factors for assistance. The specific question posed was "What research is, or should be, going on now that might produce new technologies that could, or should be, integrated into the space station after its initial operating capacity has been established?" The committee responded to NASA's question by proposing to assemble a group of eminent scientists

to address this issue and to present its views to the research community by means of a symposium on human factors research needs in advanced space station design.

DEVELOPMENT OF THE SYMPOSIUM

The Committee on Human Factors initially formed a small Steering Group composed of six researchers representing a broad range of relevant disciplines (i.e., human factors, artificial intelligence, expert systems, decision science, robotics and telepresence, and social science and space system design). The steering group was introduced to the task at hand through briefings from various NASA headquarters offices, including the Office of Aeronautics and Space Technology and the Space Station Office. Based on the information gathered during these briefings, the steering group then developed the following list of symposium topics and questions for consideration by prospective speakers.

- System Productivity/People and Machines
 - How can human performance and productivity be defined?
 - How can system productivity be measured and evaluated?
- Expert Systems and Robotics and Their Use
 - What are the requirements for reliability?
 - How can people, expert systems, and robots form an effective partnership?
- Language and Displays for Human-Computer Communication
 - How much structure does a computer language need?
 - What types of displays are most effective?
- Telepresence and Supervisory Control
 - What are the relative merits of various telepresence displays? (e.g., touch or stereopsis)
 - What can be done to increase the precision of control for remote manipulators?
- Computer-Aided Monitoring and Decision Making
 - What types of routine operations could be automated?
 - How will people use these types of aids?
- Social Factors
 - What factors affect group productivity and performance?
 - What are the potential effects of increased crew diversity with respect to such variables as gender, professional training, and interest differences?
- Human Role in Space
 - How should system functions be allocated in manned space systems?
 - Who or what instrumentality should take ultimate responsibility for system performance and safety, a human or a computer?

The general framework for the symposium was planned as follows. Each topic area would constitute a different session. Each session would consist of two formal presentations of papers prepared especially for the symposium and would be followed by a formal commentary on the papers by a preassigned discussant and would conclude with an open discussion. Members of the audience would be active participants and would be selected with this in mind.

The steering group identified and recruited three experts in each topic area: two authors and a designated discussant. The session on system productivity was an exception, having one author and one discussant. Before the symposium, all the prospective authors and discussants were invited to visit the Lyndon Johnson Space Center for briefings and discussions with key personnel involved in manned space flight research and development. Speakers and advisors were present from NASA headquarters, the Johnson Space Center, the Ames Research Center, and the Jet Propulsion Laboratory.

Following the extensive overview of NASA research efforts aimed at the space station effort provided by NASA personnel, symposium authors and discussants began preparing materials for the symposium. Individuals involved in each session worked together using an iterative peer review and revision approach in writing the papers and the formal commentary on them that was to be included in the symposium proceedings. Each group took responsibility for the completeness and technical accuracy of the material representing its area of expertise. Prior to the symposium, authors and discussants received a complete set of papers and commentary for each of the sessions.

The symposium was held at the National Academy of Sciences on January 29-30, 1987. Following the symposium, authors were asked to revise their papers and to suggest revisions to papers written by others based on the information and insights gained during the symposium.

The steering group did not consider its mandate to encompass the task of developing specific research recommendations for NASA. The symposium presentations and commentary serve that purpose. However, the closing remarks of the keynote speaker and the chair, which appear at the end of these proceedings, stand as their personal interpretation of what was said that was the most important.

SYMPOSIUM ABSTRACTS

This section summarizes the contents of each of the symposium papers and provides the interested reader with an overview of the symposium program.

System Productivity: People and Machines

Productivity in the Space Station (Raymond S. Nickerson) The concept of productivity, while elusive, has been an important one in economics and engineering psychology and is frequently encountered in discussions

of the space program and of the space station in particular. Nickerson begins with a discussion of what productivity means and how it has been assessed in earth environments. Several variables that have been shown to affect it are identified. Factors that are likely to have an impact on productivity in space are discussed, with emphasis on a variety of stressors that may be expected to characterize the space station environment. The paper ends with a set of recommendations for research.

Expert Systems and Their Use

AI Systems in the Space Station (Thomas M. Mitchell) Among the technologies that will help shape life in the space station, artificial intelligence (AI) seems certain to play a major role. The striking complexity of the station, its life support systems, and the manufacturing and scientific apparatus it will house require that a good share of its supervision, maintenance, and control be done by computer. At the same time, the need for intelligent communication and shared responsibility between such computer programs and space station residents poses a serious challenge to present interfaces between people and machines. Hence, the potential and need for contributions from AI to the space station effort are great.

This paper suggests areas in which support for new AI research might be expected to produce a significant impact on future space station technology. The paper focuses on two areas of particular significance to the space effort: (1) the use of knowledge-based systems for monitoring and controlling the space station and (2) issues related to sharing and transferring responsibility between computers and space station residents.

Expert Systems: Applications in Space (Bruce C. Buchanan) The technology of artificial intelligence (AI), specifically expert systems, is reviewed to examine what capabilities exist and what research needs to be conducted to facilitate the integration of humans and AI technology in future space systems. An expert system is defined as a flexible, symbolic reasoning program that uses heuristics to manipulate symbolic data in order to generate plausible answers to questions. Four goals are identified for expert systems: (1) performance (at a standard comparable to the best specialists); (2) reasoning (as opposed to straight "number crunching"); (3) understandability (the ability to explain why an answer is plausible and how it was generated); and (4) flexibility (the ability to deal with novel situations). Methodological techniques for achieving these goals are discussed, including modularity (keeping domain knowledge separate from decision rules, and independent clusters of domain knowledge separate from one another) and uniformity of language and constructs (both internally between segments of the program, and externally between the program and the intended users). The problems of collecting, representing, storing, maintaining, and manipulating

domain knowledge are reviewed. Buchanan concludes that existing expert system technology is adequate for some problems but can be improved to use the very large knowledge bases required by a system as complex as the space station.

Language and Displays for Human-Computer Communication

Change in Human-Computer Interfaces on the Space Station (Philip J. Hayes) The planned longevity of the space station will require modularity in its design to allow components to be changed and updated as independently of one another as possible. This paper explores the issue of modularity in the design of human-computer interfaces for the space station. The need for modularity centers on the rapid rate of expansion in the kinds and combinations of modalities (typing, graphics, pointing, speech, etc.) available for human-computer interaction, and on the techniques available to effect their implementation and interaction. The paper assesses the appropriateness of current and forthcoming modalities according to task, user, and space station environment. A secondary factor that makes change in human-computer interfaces inevitable for the space station is the development of intelligent interfaces. The paper discusses methods of achieving intelligence in interfaces and in what circumstances it is desirable. The question of how to achieve the necessary changes in human-computer interfaces is considered, focusing on methods of obtaining a clean separation between the interface and the underlying space station system application. User interface management systems and interaction interface development environments are also addressed. The paper concludes with a set of research recommendations covering both research into new interface technology and methods for dealing with the consequent need for change in interfaces.

Cognitive Factors in the Design and Development of Software in the Space Station (Peter G. Polson) The paper describes major problems in the design of human-computer interfaces for systems on the space station and shows how systematic application of empirical and theoretical results and methodologies from cognitive psychology and cognitive science can lead to the development of interfaces that reduce training cost and enhance space station crew productivity. The paper focuses on four issues: (1) transfer of user skills; (2) comprehension of complex visual displays; (3) human-computer problem solving; and (4) management of the development of usable systems. Four solutions to the problems are proposed: (1) use of information processing models of tasks in the design process; (2) allocation of adequate resources to user-interface development; (3) use of user interface management systems; and (4) use of existing expertise in NASA.

Computer-Aided Monitoring and Decision Making

Robustness and Transparency in Intelligent Systems (Randall Davis)

Building and operating a manned space station will give rise to problems of enormous complexity in an environment that is both hostile and unfamiliar. The complexity of the station and the novelty of the environment preclude the creation of an exhaustive list of contingency procedures. Unforeseen events will inevitably occur, requiring real-time interpretation, diagnosis, and response.

The paper reviews the failure of a fuel cell during the second space shuttle mission in order to give an example of the kind of unanticipated event that can occur and examines the varieties of knowledge and engineering reasoning required to deal with such an event. Davis considers what might be required to have a computer assist in this task by giving it an understanding of "how something works". Some nonsolutions to the problem are discussed to demonstrate why existing technology is insufficient, and several research themes are then explored. The nature and character of engineering models are considered and it is suggested that their creation, selection, and simplification are key issues in the sort of understanding that should be created. Recalling the difficulties involved in the capture of Solar Max, the paper argues for the necessity of complete design capture and speculates about what it would take to create a design capture system so effective that it would be almost unthinkable to create or modify a design without it. The paper also considers what can be done at the design stage to create models that are easier to use and more effective; that is, how to design in such a fashion that interpretation, diagnosis, and response are made less complex processes.

Decision Making--Aided and Unaided (Baruch Fischhoff) There are few aspects of space station design and operation that do not involve some decision making, whether it be choosing critical pieces of equipment, choosing to trust automated systems, choosing where to look first for the source of an apparent anomaly, or choosing the range of conditions for pre-mission testing. Knowing how people intuitively make such decisions provides a basis for determining where they need help, in the form of automated decision aids, specialized training, or designs that are robust in the face of fallible decision making. Although it has much in common with decision making in other contexts, space station decision making presents some special demands. These include: (1) the need to create a shared model of the space station and its support systems, which will coordinate the widely distributed decision makers capable of affecting its performance; (2) the need to make decisions with imperfect systems, whose current status and future behavior are incompletely understood; (3) the need to make novel decisions, responding to nonroutine situations. The human factors research needs in each of these areas are identified, using as a point of departure the literature of behavioral decision theory. Meeting these demands

will require the sort of programmatic research effort that has distinguished NASA in the past.

Telepresence and Supervisory Control

Teleoperation, Telepresence, and Telerobotics (Thomas B. Sheridan) The problems of integrating humans and automated or robotic systems in space environments are discussed, beginning with brief definitions of key terms like teleoperation, telepresence, telerobotics, and supervisory control. The early development of teleoperators is summarized, from the crude mechanical earth-moving and construction equipment available prior to 1945, to the industrial robots, equipped with primitive computer vision, wrist force sensing, and "teach pendant" control boxes that were in use by the early 1980s. The current status of teleoperator development is evaluated, and multifingered manipulators, touch sensing, and depth perception are cited as areas in which promising research is occurring. A need is identified for a formal theory of manipulation to guide the development of human-machine integrated sensory-motor control systems. Research needs are identified in the following areas: (1) telesensing (including resolved force, touch, kinesthesia, proprioception, and proximity); (2) teleactuating (including multi-degree-of-freedom end effectors, two-arm interaction, and multiperson cooperative control of teleoperators); (3) human-computer interaction in a computer-aided environment (including simulation, planning/decision-aiding, and command/communication/control). It is concluded that research in the areas discussed is critical for the development of teleoperator/telerobotic capabilities, which will permit the best relative use of both human and machine resources in future space systems.

Telerobotics for the Evolving Space Station (Lawrence Stark) In this paper, telerobotics is used to mean remote control of robots by a human operator using supervisory and some direct control. By robot is meant a manipulator/mobility device with visual or other senses. This is an important area for the evolving NASA space station. The paper suggests that triplicate or three way planning should be employed. It is important to carry out research to accomplish tasks: (1) with people alone, if possible, such as in extra-vehicular activities; (2) with autonomous robots (AR); and (3) with telerobotics. By comparing and contrasting the research necessary to carry out these three approaches, present problems may be clarified.

The paper describes an experimental telerobotics simulation suitable for studying human operator performance. Simple manipulator pick-and-place and cranking tasks allowed quantitative comparison of a number of calligraphic display viewing conditions. The Ames-Berkeley enhanced perspective display was utilized in conjunction with an experimental helmet mounted display system. A number of control modes could be compared in this telerobotics simulation, including

displacement, rate, and acceleratory control using position and force joysticks. Communication delay was introduced to study its effect on performance.

The paper suggests that the impetus and support for telerobotics research technology should come from NASA and from private industry and that such research could also be conducted, with support from NASA, in university laboratories.

Social Factors in Productivity and Performance

Social Stress, Computer-Mediated Communication Systems, and Human Productivity in Space Stations (Karen S. Cook) The paper has two distinct but related foci. First, it considers the issue of stress and reviews the social psychological literature relating stress to individual and group functioning. Primary attention is focused on the link between stress and group productivity. The paper identifies promising lines of research in the social sciences and poses issues that might be of particular interest to NASA for future research. Second, the paper considers a broad class of problems that arise from the fact that life aloft requires, almost exclusively, mediated communication systems. This section of the paper addresses the psychological and social aspects of mediated communication (primarily, computer-mediated communication systems) and its impact on individual and group performance or productivity. The concluding section of the paper proposes a critical set of research needs that NASA might take as recommendations for programmatic research. These complement research currently being supported by NASA's Human Factors Division. Emphasis is placed on what are termed critical social contingencies, namely, those psychological and sociological aspects of life as envisioned on space stations that, if not managed well organizationally, could create major problems for crew productivity and viability in space.

Control, Conflict, and Crisis Management in the Space Station's Social System (H. Andrew Michener) The paper discusses two social systems: (1) the space station social system in the year 1993 and (2) the space station social system as it may have evolved by the year 2000. Because neither of these social systems exists today, they cannot be investigated by empirical techniques; thus, the discussion in this paper is necessarily theoretical and conjectural. It is proposed that the year 2000 social system, in contrast with the 1993 system, will be larger in size and more differentiated in composition, will make greater use of on-board computerization (artificial intelligence), and will pursue different goals and subgoals. These changes will, in turn, create a year 2000 social system that is more complex, more differentiated into subgroups, and more decentralized with regard to decision making than the year 1993 system. It is suggested that several consequences will follow from increases in complexity, differentiation, and decentralization. Specifically, it is likely that: (1) the supervisory-control system on board the space station

will shift from a hierarchical form to a heterarchical form; (2) the potential for, and severity of, interpersonal conflict will be greater; and (3) the logistics of responding to crises will be different. Each of these points is discussed in detail. The paper closes with suggestions regarding research that might usefully be conducted today in anticipation of these changes.

The Human Role in Space Systems

The Roles of Humans and Machines in Space (David L. Akin) The fundamental requirements for any self-contained device performing a useful function in space are identified as follows: (1) sensation (the ability to detect objects); (2) computation (the ability to formulate a plan of action); (3) manipulation (the ability to interact with, and to alter, the environment); (4) locomotion (the ability to maneuver within the environment); (5) support (power, cooling, etc.). The past and present roles of human and mechanical systems in fulfilling these functions in space activities are reviewed, with emphasis on the special contributions of people to the performance of space systems. The need to take an earthlike environment into space in order to accommodate humans is also discussed, including the constraints of atmosphere, consumables, volume, work cycles, and gravity. It is concluded that there will continue to be necessary and sufficient roles for both humans and machines in space systems for the foreseeable future. Research needs are identified in the following areas: (1) development of a meaningful data base on human and machine capabilities and limitations in space environments; (2) identification of appropriate roles for humans and machines in space systems; (3) development of appropriate metrics of human and machine performance; and (4) an assessment of anthropocentrism (the tendency to design autonomous machines based on a human model).

Sharing Cognitive Tasks Between People and Computers in Space Systems (William H. Starbuck) The differences between people and computers are persistent and profound. Although computers' capabilities have been developing rapidly, computer simulation of human thought has had little success. However, the differences between people and computers suggest that combinations of the two can achieve results beyond the capabilities of each alone. For that reason, NASA should devote research to improving the interactions and synergies between people and computers.

Nearly all the research on human-computer interaction has focused on people who lacked thorough training and who had little experience with computers. Since most of these findings may not extrapolate to the well-trained and experienced operators of space systems, there is need for studies of such users. Five research topics seem especially interesting and important: (1) fostering trust between people and expert systems; (2) creating useful workloads; (3) anticipating human errors; (4) developing effective interface languages; and (5) using

meaningful interface metaphors. Inherent in these topics is an implication that NASA should develop a user interface management system that will recognize the needs of different users, allow different users to express their personal preferences, and protect users' individuality. The paper concludes that to improve the quality of designs and to improve users' acceptance of designs, experienced astronauts and controllers should participate in the designing of interfaces and systems.

OPENING SESSION

WELCOME

Thomas B. Sheridan

Welcome to the Symposium on Human Factors in Automated and Robotic Space Systems. I will start by saying a few words about why we're all here. A bit over a year ago, actually before the Challenger accident, Melvin Montemerlo, the Manager of the Human Factors Research Program and Co-Manager of the Automation and Robotics Program for the Office of Aeronautics and Space Technology in NASA Headquarters, requested the Committee on Human Factors of the National Research Council to consider the needs for human factors research in evolutionary manned space stations. Mel asked the committee to look at future manned space systems beyond the Initial Operating Configuration (IOC); looking ahead into the late 1990s and beyond. (I might mention that Mel is on sabbatical leave and Mike McGreevy is currently managing the programs.)

It was clear to us that any new research started now could not have much effect on the design of the IOC, so we knew we had to speculate for a period beyond this first space station. It was also clear to us, as we thought about it, that if a single issue could be considered to have the most effect on human factors in the space program, it would be the computer. And while much of the public, even the Congress, and even some in NASA management, have come to think in terms of the astronaut versus the computer and automation and robotics, I believe the science and technology community and many in NASA know better. It's really the astronaut, or human beings, working together with the computer and automation and robotics in close cooperation, that will result in the greatest mission success. But simply to say that and to have it really happen, are, of course, not the same. We have a long way to go to piece it all together.

So we were asked to think about this major issue and to organize a symposium, composed of experts who, in our judgement, represented the most critical areas of human-machine interaction, even though we could not cover all of the major aspects of human factors. The committee decided that it would be most effective if it concentrated on human factors issues in relation to computers, automation, robotics, and the roles of people in the space stations of the future. A reason for selecting the symposium format was the opportunity that it would afford an exchange with other people in the scientific community (including NASA) and other organizations who might make cogent contributions to the discourse.

Let me identify the people who worked hard with the Symposium Steering Group to put this symposium together, the staff of the Committee on Human Factors: Dr. Stanley Deutsch, the Study Director for the committee; Dana Kruser, a consultant to the committee, who is largely responsible for having all of the symposium papers ready on time; Elizabeth Neilsen, the committee's staff assistant, whose support on the logistics was invaluable; and Beverly Huey, also a consultant, who helped us to meet our schedule in myriad ways. They will all be available during the meeting if you have any needs.

We ask you to listen to our thoughts, and possibly some irreverent comments about the space program and the research that's been done or should be done, and to participate in the discussion. One reason that the proceedings are available out at this time is so that we can capture your ideas and include them in the proceedings of the meeting.

I thank you for participating and I hope we can make this an interactive meeting.

Now, I want to introduce Dr. Raymond S. Colladay, the Associate NASA Administrator for the Office of Aeronautics and Space Technology, to say a few words about the NASA organization. I will then ask Dr. David A. Goslin, the Executive Director of the Commission on Behavioral and Social Sciences and Education (CBASSE), to say a few words about the National Academy of Sciences and the National Research Council. The Committee on Human Factors is located organizationally within CBASSE.

So, first, Ray Colladay.

INTRODUCTION

Raymond S. Colladay

I'm delighted to see that in spite of the snow here in Washington, there is such a good turnout. I was talking to Stan Deutsch before the meeting and he told me that attendance had to be restricted so that the group would be small and intimate to encourage good interchange and dialog. I'm pleased with that because it provides a focus on a subject that is extremely important to NASA. And I'm further pleased by the fact that Human Factors is being considered at this symposium in the context of Automated and Robotic systems, because that's precisely that way we should look at that subject. This reflects what NASA is trying to do to bring those disciplines together.

I think that when you look into the subjects which you are addressing in this symposium, you're going to see a discrepancy between our goals and our current capability, specifically in the NASA program. Your feedback in the discussions and in the proceedings of this meeting will be very important to us in planning the program and in trying to get our capability on track with our expectations and our vision. We have great plans for extending human presence in space. The space station is only the first step in that vision, which is taking shape right now as we contemplate lunar bases, expeditions to Mars, and other missions beyond the space station.

It is my pleasure to welcome you to this symposium on Human Factors in Automated and Robotic Space Systems, and I'd like to thank the National Research Council's Committee on Human Factors for their efforts in conducting this symposium, and for their valuable contributions over the years to NASA's Aeronautics and Space Human Factors research programs. The committee has helped us to formulate and develop the kinds of programs we need in this area.

The subject of this symposium is timely indeed. Yesterday was the first anniversary of the Challenger accident, a day of rededication to excellence in memory of the Mission 51-L Challenger astronauts. It was a day when, as a nation, we rededicated ourselves to the excellence that characterizes America. For our part, we at NASA are developing a clear vision of the future in space and are currently refining our research and technology development plans to ensure the health, safety, and productivity of humans in space throughout the coming decades. Although it was only formalized as a research discipline about five years ago, our Space Human Factors Program is built upon a long history

of aeronautical human factors research, and extensive agency experience in life sciences research and manned space flight.

Something else is happening in the NASA program which pleases me, and that is the start of a new building for human performance research for the space program at our Ames Research Center in California. I intend this building to be the first leg of a major facility that combines human performance and automation research. We are, in fact, putting a building in place to reflect exactly the kind of merger of those disciplines that this symposium is addressing. We'll call it the Human Performance and Automation Laboratory. This will pull those disciplines together in a very realistic way, and will get researchers working in the laboratory in computer science, artificial intelligence, automation, and human factors.

I look forward to the results of this symposium. I think it will be extremely helpful to us. We welcome this opportunity to interact with you and I wish you luck in the proceedings and the discussions that follow. Thank you very much.

KEYNOTE ADDRESS:
HUMAN FACTORS RESEARCH FOR THE NASA SPACE STATION

Allen Newell

Symposia are held for many reasons. This one is to do a task. I hope you are intellectually stimulated by what you hear and will take away some new knowledge that you do not already possess. I also hope the symposiasts have enhanced their own knowledge by their studies and are gratified by the chance to advance their views from this earth-bound, but otherwise splendid, platform. But neither of these has much to do with the actual reason for this symposium. We are here to help shape the research that NASA will perform on how humans interact with the technology of the space station.

In a nutshell, and to state what I hope is already shared knowledge among us, operating in space poses extreme challenges. It is a hostile, aversive, constraining and unforgiving environment. Our intent as a species to make such operations successful and to continuously extend their scope in complexity, duration, and usefulness is epitomized in NASA, and other space agencies around the world. It takes its concrete form by the posing of specific projects, each more daunting than the last, but (skillfully we hope) set just within the bounds of the reachable. For us today that project is the space station, a project with an initial development phase prior to launch of a decade and a total lifetime of several more decades. Such projects force us to not only use the best available technology and science, but to extend them substantially. For us today, the question is what research is most needed that could have important payoff for the space station. It is not possible, of course, to encompass in a single effort all the technologies and sciences that support the space station. Thus, we focus on how humans and technology can cooperate to carry out the operations of the space station.

We will speak today almost entirely of the space station. That is proper, because we need projects to give as much form as possible to a future which is almost agonizably open. But, such far-future projects are emblematic of entire technological futures. Thus, behind the space station is to be seen an entire spectrum of future space systems, replete with automated and robotic devices, while also being a habitat for humans. Indeed, the space station is itself not a single envisioned system, but a projected series that stretches out in time and evolves in significant ways. Only occasionally will we have need to distinguish even between such relatively concrete visions as the IOC

and SSOC. The research talked about here in the context of the space station is what we see as necessary to this entire technological future.

This is our task. Its success can be measured by the influence of this symposium on how humans and technology actually work together. Do the crews that run the space station, both on the ground and in space, have an easier, safer, more productive time than would otherwise have happened? We are only one player in the hundreds of individuals, groups and organizations that affect what goes into the space station, and a highly transient one at that. Our only leverage is the cogency of the ideas we put forth. Still, we fail if nothing down stream is different because of what we say here today. It is not enough to have an effect, it must be the right kind of effect at the right place and time. Furthermore, the effect depends not only on NASA decisions about its research program, but also on the quality of the research that is thereby enabled, and whether its results transfer into the operational space station--a notoriously tenuous conjunctive chain. Still, though we talk here today of possible research, we hope for operational results.

My task, right now, is to get us launched--to set the stage and provide the context for the papers you will hear over the next two days. I will only take a few minutes to do this. But some overview will help us to keep on track throughout the meeting.




FIGURE 1 Artist's construction of the Space Station, due to launch 1996

THE SPACE STATION

Let's start with the space station itself--although I am hardly the one to do so, with an audience that contains many with active responsibilities for it. Still, even I know enough to start with the obligatory picture (Figure 1). This is of course a fantasy, composed from the minds of many persons and living only there and derivatively in the minds of receptive audiences. In accordance with its fantasy character, it changes continually--if not daily, at least monthly. The planners harden the fantasy with physical mockups that can be walked in and gawked at. That helps, but the time constants to realization are still of the order of half-decades.

This way of talking about the space station may induce a sense of fragility. That could be a good thing, if it brings with it an increased sense of commitment to making it happen. However, my actual objective is to induce a sense that much can change in the space station before it takes its place in the sky and, indeed, after it does. If we are to consider launching research in 1987 and expect it to have operational impact, then the time scale of that operational world must be sufficiently long and its character sufficiently malleable.

Planning--even research planning--must have some grip on reality. Thus, we need to focus on the hard constraints on the space station--the ones that appear to hold no matter what, and on which we can build securely. Table 1 presents three handfuls--already more than can be assimilated in an introduction. These constraints are what strike a technically observant human-factors specialist immediately upon hearing a briefing on the station. They are the constraints that shape the roles that humans must play and the tasks they must perform to make the space station function. What makes them unyielding is the

TABLE 1 The Hard Constraints that Apply to the Space Station

-
1. Long lifetime of the station (decades).
 2. Medium term crew residence on board (months).
 3. Small group of residents aloft (less than ten, to begin with).
 4. Large group of operators (non-residents) aground (hundreds).
 5. Very small amounts of resources available per resident.
 6. Very small amounts of space available per resident.
 7. Infrequent physical communication (months).
 8. Continuous, but limited-bandwidth communication.
 9. Time delay of station communication of .5 to 2 seconds.
 10. Modest time constants of action (minutes to hours).
 11. Weightlessness.
 12. Continuous, high task load.
 13. Continuous high threat-level of many potential errors.
 14. Continuous public exposure.
 15. Completely artificial environment.
-

limited state of our space technology, the primary goals set for the station, and the necessity of acquiring certain experiences as stepping stones to future technological frontiers. No matter how technology changes we must pass this way to move forward--not, of course, with the exact particularities of the space station we will build, but through something with the general characteristics listed in Table 1.

Many familiar things follow from this: the general strangeness of the weightless world and its frustrations; the isolation of the station group, coupled with the lack of privacy and the extent to which members are locked in; the public work-oriented, regimented world; the complete dependence on the efforts of others; the stress of continually living close to fatal errors. By and large, humans respond adaptively to all these conditions. Still, humans in space must spend their psychic resources to cope with these conditions, rather than spend it in other more productive ways.

One striking thing is how saturated with technology the life of the station will be. This is completely true of those stationed aboard, but is almost as true of those aground for their workaday world, although they get to go home to the grass each evening.

Another striking thing is that the residence time-scale is long enough so that many functions have to be accommodated that can be avoided in shorter flights. The station appears to be a microcosm of life--so many activities must occur that one can find any problem or task one looks for, or at least a close analog. Now, in fact, this is not quite so. Many functions, such as raising a family, becoming educated, moving to a new home, and planning retirement, do not show up at time scales even as long as months. And to those concerned with the man-machine system in the modern fighter plane, where the focus is on actions in the subsecond range, the station will appear downright leisurely. That the space station occupies a middle range in the total timescale of human action is a significant simplification--as we will discover when we have to plan permanent space or lunar stations. But even so, from the perspective of a human factors analyst, the space station has moved a long ways toward total living and not just temporarily occupied workspace. Along with that has come an almost un-enumerable collection of tasks that humans must perform, and the need for designing the artificial environment in which to perform them.

Still the tasks must be enumerated. One of the great liabilities of technological environments is that they don't take care of themselves--not yet and for some time to come. The tasks to be performed in the station and between ground and station must be enumerated and explicitly planned for. What we fail to enumerate here below is in parlous state up above. There will always be true stories about the novel activities of intelligent astronauts, solving life-critical problems or having fun in ways we could not predict. Bless them for that. But let no one argue back from that blessed fact to the need for less preparation. And preparation implies explicit task description and enumeration. NASA, of course, has gone to great lengths to do this. Table 2 provides some reminders of what those enumerations cover.

The left hand column simply lists the various subsystems involved, so one gets some notion of diversity. With respect to each of these there are many actual tasks to be performed. To enumerate them is to descend into the technological gritty of each type of system. But various types of activities that go into these tasks can be identified, which is what the right hand column shows. These generic activities come in indefinite variety as well, in terms of what must actually be accomplished, with what initial knowledge, and against what constraints. Finally, I have put across the bottom what is perhaps the most important factor, namely, that the time scale over which these tasks endure stretches from less than a second to about four months—seven powers of ten. Each task in its individuality fits into this time-stretch at some point. But every duration contains tasks of every type.

There are two points and one conclusion to be made from all this. First, I would impress upon you that there are an almost unimaginable variety of tasks, which contain almost any combination of task demands one cares to contemplate. Second, the vast majority of these tasks are to be accomplished by some combination of humans and technology. To be sure, at the top ultimately there is a pure human, if only a congressman; and at the bottom there is a pure machine, if only a pushbutton making an electrical contact. It follows that we can consider today only a selection of all the problems. We will of course seek for research that is generic in its character and that will impact large classes of these tasks. But much that is important will not even be mentioned.

TABLE 2 Types of On-Station Tasks for the Space Station

Subsystems	Generic functions
Power	Power handling
Guidance & navigation	Checkout
Communication & tracking	Mechanical actuation
Data handling	Data handling and communication
Propulsion	Monitoring/control
Environmental control and life support	Computation, decision and planning
Thermal	Fault diagnosis and handling
Structures/mechanisms	Sensing
Crew systems	
Payloads (experiments, manufacturing, observations)	

.1 sec	1 sec	10 sec	100 sec	10 ³ sec	10 ⁴ sec	10 ⁵ sec	10 ⁶ sec	10 ⁷ se
	1 sec		1 min		1 hour	1 day		1 month

THE TECHNOLOGY OF INTERACTION

The classical situation of human factors has been that an industrial or military organization develops some machine to do some task. The human-operator aspects of controlling this machine and of being trained to do so are dealt with in due course. In the best of cases, this occurs early enough to permit modest alteration of the engineering of the interface. But in the main, the technology of the machine is autonomous and fixed.

With the advances in artificial intelligence and computer science in general, and in computer interfaces in particular, the situation is changing--and changing in several ways simultaneously. First, the machines are becoming more complex, yet capable of more autonomy and intelligence at the same time. Second, the interfaces themselves are becoming more intelligent so that they can aid the user and operate cooperatively with him. Third, all interfaces are becoming alike in their utilization of a common hardware and software technology. Finally--and of a different order entirely--the technology on which all this is based in itself undergoing rapid evolution, so that all the features just mentioned are not new fixities that can be depended upon, but are themselves on the move. All of these current truths have double force for the space station, which is located a long ways in the future. Let us focus on each of them in turn.

Machines are controllable arrangements of matter and energy that do things to the physical world. (Thus, tools are machines.) The ability to be controlled is of their essence, for it is what changes them, as micro parts of the world, from a thing that can be taken advantage of (as to drink from a brooklet happened upon) to a thing that can be used at will (as to turn on a faucet whenever thirsty). So machines bring with them the problem of the human-machine interface, and necessarily those interfaces are dynamic and continue throughout the duration of use.

As machines become more capable, through the rational foresight of their designers and the skill of their builders, the tasks that machines can do without human intervention increase. Although the real measure is in the total range of useful tasks they can accomplish with acceptable reliability, an appropriate indicator is the length of time machines can go without interaction with humans. With this increased scope comes inevitably the problem of who should do a task, the human or the machine. Formally, this is exactly the same as the problem of whether this human or that should do a task, or whether this machine or that. However, because of the category difference, the human-machine question is taken to have a more profound character and it becomes the focus of scientific attention. It is a surrogate, of course, for our need to understand the advancing capabilities of machines.

That question is finally about to change its form radically. The advances in computers and computation have now been driving exponentially for forty years. All parts of that advance are significant for us today, in part because they all interrelate. The driver of it all, we always say, is the cost/performance of the computing devices and the level of their integration. But by this time

that itself depends on software design systems with quality graphics. So it is all one ball of wax. Nevertheless, the parts where the advances touch us the most here today is in robotics, artificial intelligence and the technology of the human-computer interface. Through these, the amount of intelligence that can be incorporated into machines is now reaching the place where the problem of assignment of functions to men or to machines no longer holds any charm. The question must be phrased--How can humans and technologies cooperate to attain a set of system-level goals.

The situation at the interface between the human and some machines provides a good example of the increase in the capabilities that are available, with a concomitant increase in the complexity for those of us who design and understand these systems. As machines increase in capability, interfacing to them becomes a complex task in its own right and requires substantial knowledge about what is required to communicate knowledge back and forth--languages, protocols, communication over intermediate links, the status and location of the communicants, and on into the night. The solution is to have special agents that have this knowledge or know how to acquire it, in short, intelligent interface agents. But such agents imply that knowledge about how things work will be distributed--of what good are such agents unless they relieve other parts of the system of the responsibility for having certain knowledge and skills? But this reinforces the point made earlier that it no longer makes any sense to cast the problem of how humans work with technology in exclusive terms of who controls whom. Rather, it must be in how agents embodying distributed sources of knowledge cooperate.

One more point about the technology and I am done with it. If NASA had to settle for the level of intelligence in current robotic and expert systems, this symposium would have a very different character. We have, of course, come a long way in computer science in the last forty years and this is plainly evident in existing robotic and intelligent systems. But the changes are proceeding very rapidly and substantially more capabilities can be expected to be available in another five years or in five years more again. This introduces uncertainty into our proceedings, for we must not only talk of what new research might bring, but must place this against a background that will increase in possibilities no matter what NASA does. But this same motion also adds to the sense of excitement of the new powers that are possible in the space station. The space station, by being a project measured in terms of decades, both suffers and benefits in the extreme from this motion of technology.

RESEARCH OBJECTIVES

Given the picture just sketched of tasks and technology, the question of the day is what research should be done. The substantive answers to that question are the responsibility of the speakers of this symposium. I would only ask you to keep three general considerations in mind.

First, the research topics raised here range widely--from artificial intelligence, to the human-computer interface, to telerobotics, to issues of social organization. These are not just a congeries, brought together to obtain coverage. They are all facets of how humans are to interact with the primary technology of the space station, and what technologies are involved in that interaction. A research program needs to address all these aspects in some coherent way, and not treat them as separate questions.

Second, we have had to sample--to focus on some issues and to neglect others. But the research program needs to consider the full range of phenomena. It is in research plans, and the study efforts that support them, that one engages in the compulsive attempts to taxonomize the domains and worry seriously about coverage and missing elements. A symposium is to make clear the fruitfulness of research areas and to show that there are exciting research questions. Attempts at completeness and evenhandedness would only dull the senses. Third, with more glibness than honesty, I have just shifted an immense burden from the symposium speakers to the symposium participants--or at least some of them. For, of course, the domain of research is so broad that coverage is a chimera. That is especially true if one thinks of research as devoted to getting answers to specific questions about a specifically configured space station. Such answers must be obtained--that is what engineering requires. And in the present context it is human engineering and even organizational engineering. A research program that is in effect a systematic and planned program of human and organizational engineering, with the resources to do some background studies, cannot possibly provide the coverage that is necessary. Thus, the research program must be aimed at discovering conceptual, theoretical and technical tools that will permit the human and organizational engineering of the space station to proceed with greater efficiency and accuracy. Only if a research program advances the theoretical state of the art, including therein systematic organizations of data that permit answering a multitude of questions, will it serve NASA in the decades it takes to achieve the space station.

THE INSTITUTIONAL CONTEXT

Thus far, like a good cobbler, I have stuck to my last, discussing the substantive issues. But it is important to say something about the institutional context in which the symposium occurs. Blessedly, I need not take my point of departure from the current spasm of reflection, critique and renewal that has been the fallout of the Challenger disaster. For our timescale is too long for that to count as more than a transient. At least that is true if NASA can continue in its planful ways, which it shows every sign of doing. Thus, in setting out the institutional context I will not talk about the microstructure of command and timing that will, in fact, have the lion's share of responsibility for whether any trace of this symposium's efforts survives these two days. Instead, I will point to larger entities.

Let us start with NASA. It is, of course, the primary player. It is its space station, after all. Its primary view of human factors considerations has got to be simply as an ingredient to make the space station better--as a factor of production, in the economist's sense. That view leads inevitably to working backwards from specific questions about the space station to specific studies to answer them. After all, in the logic of planful organizations: To get X, set up a plan for getting X. Furthermore, the cogency of a plan can only be apparent if it explicitly and recognizably puts down each step, from what is available initially to the obtaining of X. This leads to a thoroughly applied effort and one characterized by short-range goals with tight loops of justification. Such a logic is certainly appropriate in part--after all, if NASA doesn't do the studies to deliver the answers it needs on the nitty-gritty of the space station, who else will? But the timescale of the space station is long enough so that other attitudes are appropriate as well. NASA can change the available science enough to make a difference to the space station itself. And to do that the research must be launched on a broader and freer path, letting it pick its way among the interesting questions of today to the different questions of tomorrow. The issue for NASA then is whether it will rise above the immediate applied questions of human factors--to which the safety and productivity of the astronauts will force attendance in any event--to the faith that major gains for the space station can be attained from supporting basic long-term research.

Each of us has our own stories of where such long range research by an institution has made immense differences to the downstream operation of that institution. Not being a NASA insider, my stories of that ilk do not come from NASA. But even to an outsider it is apparent that there must be a whole book full of such stories. After all, space science is an almost new science, even though, as always with science, it has a whole tangle of historical roots in early rocketry, astronomy, and more. And space science is practically a creature of NASA, so NASA must know all about the gains from bringing a new science along.

Nevertheless, it may be worth recounting briefly one of my own stories. This is DARPA's creation of the field of artificial intelligence and expert-systems technology. DARPA did not start artificial intelligence, that occurred in the mid 1950s. But only a few years afterwards, in the early 1960s, DARPA began its open support of that part of computer science. It did so in an essentially free spirit and mixed with the many other things it was also supporting, such as time sharing, graphics, multiprocessors (Illiac IV) and networking. The support was substantial, but was far from being the dominating item in the mix of DARPA programs. The important aspect, from the present point of view, is that DARPA started its support in 1962. By 1972, a mere decade later, the first expert systems had begun to emerge--Dendral and Mycin. By 1982, only one more decade, the commercialization of expert systems had begun. Today, five years later, though still a green and nascent technology, it has become the property of us all. It has become integral to much of DOD's own future and is now integral to our discussions here. But for almost all the first twenty years, DARPA was essentially the only support for

artificial intelligence. (My friends from ONR and from NIH-supported AIM (AI in medicine) may be a little annoyed at so sweeping a claim; yet it remains close to true.) Thus, we see that an agency can bring into existence wholly new techniques and ideas for its own use downstream. It cannot usually be done in less than a decade. But in timescales that are commensurate with the space station, such things are possible. And their payoff is incalculable.

The second major player is the collection of scientists and engineers who will conduct the research. This is not a homogenous group. Most immediately, the scientific cadres within NASA concerned with human factors and artificial intelligence are to be distinguished from the scientists in the universities and research organizations across the country. Each clearly plays a different role, although, in the style of the times, strong attempts exist to weld these into a more continuous community, with the establishment of places such as the NASA Research Institute at Stanford University.

The more important inhomogeneity is among the social institutions we call professions and disciplines. Focus narrowly on the human-science issues concerning the space station, and ignore totally the half-hundred natural-science and engineering disciplines concerned with the physical structures in the space effort. However, a gaggle of disciplines are still gathered around this seemingly narrow focus. Alphabetically, they are: artificial intelligence, cognitive psychology, computer science, human factors, industrial engineering, organization theory, robotics, social psychology, sociology. I have no doubt overlooked some, but all these, at least, are represented among the speakers of this symposium. The inhomogeneity here arises from two sources. First, the issues of the space station involve multiple technologies, and the relevant human phenomena are so diverse that they necessarily make contact with different human sciences. But second, multiple human-science disciplines focus on the same phenomena, but do so from different perspectives. In particular, the emergence of the computer as a mass phenomena has raised the problem of human-computer interaction to prominence. At least four disciplines--artificial intelligence, cognitive psychology, computer science (mostly graphics and interface programming) and human factors--are currently engaged in forming an interdiscipline called human-computer interaction (HCI). The effort is currently focused on the individual in interaction with the computer via a system of interaction mechanisms (displays, keyboards, pointers, etc.). It is acknowledging, though only gradually, social and communicative dimensions. The conceptual and disciplinary turbulence involved in all this is both part of the inhomogeneity of the current scene and revelatory of it. HCI is only one part of the human-related issue of the space station, though a significant one.

The NASA situation that we discuss at this symposium provides an opportunity for these disciplines. They can, of course, treat the NASA problems as if they were just another collection of interesting situations in which to ply their investigatory trade. Our nation--blessedly, once again--is extraordinarily pluralistic. Thus, NASA research contracts and grants can be taken as providing additional

micro-research opportunities in a larger mix. This is one view and an important one.

But the NASA situation provides a larger opportunity, or at least it does if NASA chooses to make that opportunity available. The space station provides a unique focus for the development of the science of how humans interact in a technology-saturated environment. By reaching so far ahead of the degree of saturation in the rest of current society, it offers a chance to study a world well ahead of its time. It is a unique opportunity in this historical moment, although it will become less so as the saturation of the rest of the world proceeds.

It is important to realize that in applied sciences technological foci have an immense influence on the character of the science. One has only to think of the influence on human factors of its being nurtured by the aircraft industry, while being relatively ignored by other industries. Thus, NASA has a fleeting opportunity to bend the twig of HCI to a long-term concentration on aspects especially relevant to NASA's interests.

ENVOI

The ingredients of the symposium have now been assembled before your very eyes—the space station; the tasks of human-technology interaction; the technologies that are both the object of that interaction and the means to make it work; the orientation towards the research that needs to be done; and the institutional setting within which this symposium must make its contribution. Let us now move to the substantive papers.

SESSION I

SYSTEM PRODUCTIVITY: PEOPLE AND MACHINES

PRODUCTIVITY IN THE SPACE STATION

Raymond S. Nickerson

INTRODUCTION

What is productivity? How do we measure it, predict it and control it on earth? To what extent can that knowledge be extrapolated to a space context? What do we not know about productivity on earth that might be found out--and is worth finding out--through research? How might the expected findings be applied to space? How should the research be directed to ensure its applicability to space? Are there important questions about productivity in space that earth-based research is not likely to help answer?

I wish I could promise to answer these questions here. Unhappily, I cannot. These are the kinds of questions that I have had in mind, however, in preparing this paper. In what follows I will focus first on the notion of productivity and on how it has been measured and manipulated in earth environments, and then turn to the question of productivity in space, or more specifically, the Space Station. The paper ends with a set of recommendations for research.

WHAT IS PRODUCTIVITY?

Productivity is an elusive concept. It seems straightforward enough when one begins to consider it. It is easy to think about the productivity of chickens or dairy cows in terms of eggs laid or milk produced per unit time; here we are dealing with output in a very literal sense. And it does not tax one's imagination to think about comparing the output of the one producer with that of the other. To do this we need a way to describe eggs and milk quantitatively in the same terms, which is not difficult. Since eggs and milk are valued as foodstuffs, we could describe them both with respect to their nutritional ingredients. But quantifying productivity only in terms of output is not very useful from an economic point of view, and as it relates to chickens and cows as producers it would be grossly unfair to the chickens; we must also take into account how much chickens and cows consume in order to produce a given amount of nutritive capital by means of eggs and milk respectively. And to round out the picture we must factor into the equation not only what the producers eat, but

other resources upon which their continuing production depends. To do all this we may find it convenient, since not all the factors that must be considered are nutritional, to quantify everything in monetary terms. But this gives us no serious problem. The situation is still fairly simple conceptually: chickens and cows produce foodstuffs that can be given a monetary value, and to do so they consume resources that have a monetary cost; productivity can be thought of in terms of the value of what is produced and the cost of producing it. This all makes intuitive sense.

When one tries to apply the same type of thinking to human productivity, one has no trouble as long as the human activity involved is analogous to laying eggs and giving milk, in the sense of producing tangible goods that can be used to satisfy basic human needs, and consuming resources in the process of doing so. The picture gets less clear quickly, however, when what is produced is not so tangible--perhaps not even readily identifiable--and not easily quantified in monetary terms. How does one measure the productivity, for example, of the teacher, the scientist, the poet, the philosopher, the salesperson, the physician, the corporate executive, the athlete, the entertainer--or the astronaut?

Lack of definitional precision has seldom been a great deterrent to the use of words, and "productivity" is no exception in this regard. It is a popular word in economics, and like "truth" and "beauty," connotes something much to be desired, whatever it means. Within the literature pertaining to space exploration, one finds references to increases in the productivity of spacecraft crews resulting from changes in displays, control procedures or other variables, but seldom is it clear exactly what this means. The word is also seen throughout the human factors literature more generally; although Muckler (1982) has commented that the unconstrained way in which it is used here makes its meaning difficult to discern in this context. In practice, productivity is often used more or less as a synonym for performance; if performance improves, by nearly any criterion, productivity is said to go up; if performance degrades, productivity is said to go down.

Sometimes the word is given a precise quantitative meaning by virtue of the variables that are involved in its measurement. Indices of productivity are typically expressed as a ratio where the numerator is some measure of output (what is produced or the value of same), and the denominator is some measure of input (what is used up in the production process or the cost of same). What constitutes input and output, and how they are quantified, differs considerably from case to case, however; and changes in productivity indices over time can sometimes be difficult to interpret (Baily, 1986). Moreover, often the word is used as though it were intended to connote a quantitative entity, but there is no clue as to what the input and output variables are or how they could be measured.

Two concepts that are closely related to productivity are those of production and efficiency. Productivity implies production, or more specifically, product and producer. Productivity is an attribute of a producer; and a producer, by definition, is one who produces something. What is produced may be tangible (paper clips, a household

appliance, an airplane) or intangible (an educational service, entertainment). A producer may be a person, a person-machine system, a team, a factory, an industry, an economic sector (agriculture), a nation, the world.

But although productivity and production are closely related concepts they are not the same. As we have noted, productivity is usually expressed as a ratio of some measure of output or product value to some measure of input or production cost, and the goal, in most cases, is to make this ratio as high as possible. Production usually refers only to output quantity. Given these connotations, it is easy to imagine production increasing or decreasing independently of changes in productivity. If, for example, a manufacturer produced 10 percent more items in a given year than in the preceding year, but doing so required a 15 percent increase in the number of employees, we might say that production increased while the productivity of the employees declined.

The concept of efficiency, like that of productivity, relates output to the resources consumed in obtaining it. Efficiency has to do with getting the most out of given resources; the challenge is to organize a production process so as to minimize wasted effort. A process is said to be made more efficient when the unit costs of output are decreased or when the consumption of a fixed amount of resources yields a greater output than before.

Techniques for measuring the efficiency of assembly line workers were among the earliest contributions of engineering psychology to the manufacturing process and have been used extensively in the work place. These have typically involved analyzing production tasks into observable components. The development of task-analysis techniques has received considerable attention from human factors engineers (Van Cott and Kincaid, 1972; Woodson, 1981). Such techniques have been more readily applied to psychomotor tasks than to tasks that are primarily cognitive in nature or even those that have major cognitive components. Attention has been focused increasingly, however, on the problem of analyzing cognitively-demanding tasks, as an increasing percentage of the tasks performed by people in the work force are defined more by cognitive than by psychomotor demands.

We cannot hope to settle terminological issues here. Moreover, definitions are of limited utility when dealing with terms that are widely used, with a variety of connotations, within a field. For present purposes, productivity will be taken to be very close, but not quite identical, in meaning to efficiency. An entity (person, group, system) will be considered highly productive when it uses its resources to maximum advantage in accomplishing its goals. One can be efficient in the sense of not wasting resources simply by using those resources very sparingly, but that type of efficiency could be counterproductive if resources are husbanded to the point of precluding getting the task done. To be productive one has to use one's resources and use them well.

As a working definition of productivity I will use: effective and efficient use of resources in accomplishing a goal. The emphasis is on both effectiveness and efficiency. A productive system is one that

gets the intended job done and does so with a minimum of wasted effort and resources. I do not mean to split hairs here in making a distinction between efficiency and productivity; if one's idea of efficiency incorporates effectiveness, then I see no objection to thinking of efficiency and productivity as more or less synonymous. Effort and resources can be wasted as a consequence of many factors, such as poor training, lack of motivation, mismanagement, faulty organization, misscheduling, and a host of others. Productivity will be said to increase when either more is accomplished with no increase in consumed resources or the same objectives are attained with a smaller expenditure of resources.

These are still somewhat imprecise notions, but not so imprecise as to be useless. In the Space Station context, as elsewhere, when modifications in design or operating procedures have big effects on productivity, there probably will be no difficulty in getting a consensus that productivity has really been improved. When tasks are performed more easily, more reliably, and with fewer costly errors, most interested observers will probably be willing to describe what has happened as an increase in productivity, and even if not, they are likely to agree that changes for the better have occurred. It seems to be generally assumed, if only tacitly, that anything that improves human performance (increases speed, accuracy, reliability) probably increases human productivity. This appears to me to be a reasonable assumption, and a very useful one. Frequently in this paper, the discussion focuses on variables that influence performance, the justification being the assumption that what affects performance for better or worse will affect productivity in a comparable way.

ASSESSING PRODUCTIVITY

It is helpful in the present context to distinguish between the problem of determining what the level of productivity is at any given time and that of determining whether productivity is changing, or has changed. One might assume that the second problem is more difficult than the first, inasmuch as a measure of change, or difference, is derived from the more fundamental measure of absolute value: to determine whether productivity is more or less this week than it was last, one simply takes the difference between this week's measure and last week's. But this is so only if one wishes to know the magnitude of the difference. If one is content to know only the direction of the difference, it may not be necessary to know the individual magnitudes, at least if the magnitude of the difference is relatively large. One does not have to know the precise weight of each of two objects to know which one weighs more, especially if the difference is sizeable.

Productivity as a Percentage of Capacity

Productivity is sometimes quantified in terms of performance relative to a maximum. When this is done, maximum output or performance is used

as the standard against which to evaluate the actual output or performance, whether the performer is an individual, a system (say a factory), or an economy. Thus one might encounter the claim that the productivity of a given industry in a particular region is currently at about 70 percent, which would mean that that industry is operating at 70 percent of what, under certain assumptions, is the maximum possible. Economists often refer to how close to capacity factories and other manufacturing facilities are operating. The ability to specify how close to capacity some entity is operating presupposes a metric in terms of which to quantify the operation. Determining what constitutes maximum capacity can sometimes be a complicated and controversial process. Further, maximum must be understood as maximum within a particular context. The maximum output of a given factory, for example, could mean maximum obtainable with the present tooling, layout, manpower and stock; alternatively it could refer to what would be obtainable if one or more of these constraints on output were relieved.

As applied to individual human beings, capacity connotes the best (which often, but not always, equates to most) one can do in a given situation, the limit of human performance--or, more accurately, the limit of the individual performer. Conceptually, there are two ways to determine capacity in any given instance: one is to derive it from theoretical considerations; the other is to measure performance under ideal conditions. Neither works very well. While information theory once provided a basis for the hope of defining capacity theoretically, it proved to be a false hope, and psychologists have not yet found or developed an alternative that can do the job. Ideal conditions for performing a given task--which would have to include an optimally motivated performer--have proved also to be easy to conceptualize but difficult if not impossible to actualize.

Differential Productivity

Differential productivity in a business context is sometimes measured in terms of changes in the number of employees or amount of employee time required to get a fixed amount of work done, or conversely by changes in the amount of work accomplished by a fixed staff. Thus a retail company is said to have doubled the productivity of its bill collection departments when it managed, by computerizing its operation, to place the same number of calls with a 50% reduction in staff. And the productivity of an insurance company is described as increasing fivefold when the number of policies issued per employee per year increased by a factor of five (Bowen, 1986).

Studies of individual human productivity in specific job situations have often focused on the performance of individuals relative to the performance of other individuals on the same task. It is possible to say that A is more productive than B without saying anything very precise about how productive either individual is relative to a larger frame of reference. Measures of white-collar productivity typically do not yield absolute quantities, but do permit comparisons among similar

organizations (Drucker, 1986).

In the Space Station program, attention will probably be focused primarily on differential productivity (the cost of attaining some production objective in space relative to that of obtaining it on earth; or the cost at one time relative to that at another). While it would be interesting to be able to relate productivity to some theoretical maximum in this context (e.g. by relating production to some measure of capacity), it is not clear how to do that. Fortunately, it is not necessary to be able to quantify maximum productivity in order to determine whether one is moving toward or away from it.

That is not to suggest that assessing differential productivity is likely to be an easy task. Several investigators have commented on the variability of measurements of productivity, especially those that relate to individual human productivity, and on the resulting need to make many measurements over a considerable period of time if reliable numbers are to be obtained (Muckler, 1982). It is especially difficult to measure productivity in intellectual tasks, inasmuch as methods for assessing cognitive performance are not well developed. When a person is staring out of his office window, it may be impossible to tell whether he is idly daydreaming or is engrossed in "productive" thought. And even if he were known to be daydreaming, it would not follow necessarily that that time was lost from a productivity point of view. One widely held view of problem-solving distinguishes an "incubation" period in the problem-solving process during which progress is made on a problem in spite of--perhaps because of--the fact that the individual is not consciously focusing on the problem to be solved--and there are numerous examples of scientists and other thinkers reporting insights that have occurred when they were not actively engaged in working on the problem.

Whatever methods are developed for measuring productivity must take quality--as well as quantity--of output or work into account in some way. In manufacturing operations, product quality affects measures of productivity to the degree that items that fail to meet a preset standard become rejects. The importance of quality control in this sense is obvious and the difficulties that some industries (e.g. the manufacturing of computer microchips) have had are well known. This type of linkage between quality and quantity is a fairly gross one however. Differences in quality tend to be ignored so long as the quality is not sufficiently low to necessitate rejection. In nonmanufacturing activities the relationship between quality and quantity is even more tenuous, in spite of the fact that here one might expect qualitative differences in output to be both large and important. Quality will certainly be an important consideration in the Space Station context. The quality of the experiments that are done, for example, will be at least as important as the number.

Workload and Its Assessment

In a complex system the operation of which depends on functions performed by both people and machines and, especially, by people and machines in interaction, high productivity will require that workloads be at or near optimal level. Significant overload will reduce productivity through increases in the frequency of human error; significant underload will mean wasted resources at best and possibly direct negative impact on productivity resulting from boredom, inattentiveness or other difficulties arising from feelings of being underutilized or unimportant to the operation. Workload and its assessment will be important considerations, therefore, in efforts to understand, measure, or control productivity in space.

As in the case of efficiency, the workload carried by an individual is much easier to measure when the task is primarily physical than when it has major cognitive components. As Wierwille et al. (1985) point out, a major consequence of the increasing automation of modern systems is a shift in the role of the human operator away from manual control and toward monitoring and performance evaluation, and this has complicated considerably the problem of quantifying the operator's workload. How can we hope to determine how hard--how close to capacity--an individual is working when most of what he is doing is mental activity that is not directly observable?

The measurement of mental workload has been recognized by human factors researchers as a major challenge to the field and this recognition has stimulated considerable activity (Chiles and Allu'ssi, 1979; Eggemeier, 1980; Kalsbeek, 1968; Moray, 1979; Parks, 1979; Sheridan and Simpson, 1979; Singleton et al., 1971; Williges and Wierwille, 1979). Work in the area is still in the exploratory and formative stages, however, and there has not yet emerged a theory or even a widely agreed upon set of concepts and measurement procedures that are needed to provide a sense of stability and coherence.

An indication of the magnitude of the problem and of the current status of work on it is provided in the Proceedings of a NATO Conference on Mental Workload published in 1979. Johannsen (1979:3) opened the conference with the observation that "there exist too many conflicting ideas about the definition and measurement of workload", and expressed the hope that the conference would produce a consensus among participants on a definition and on a procedure for workload assessment. In his preface to the conference proceedings, Moray (1979:VIII), the organizer, acknowledged that these hopes were not realized, but noted that participants from various disciplines did come to "very similar conclusions about the validity, usefulness, and promise (or lack of each) for a wide variety of methods for approaching the assessment of workload in the human operator". It is unfortunate that the proceedings does not contain a summary of these conclusions. It does contain, however, a report from each of five participant groups, classified as experimental psychology, control engineering, mathematical modelling, physiological psychology and applications.

The experimental psychologists summarized their conclusions this way: "The concept [mental workload] reflects a genuine dimension or

dimensions of human experience in daily work...it is a concept absolutely required for the adequate analysis and description of such tasks [tasks that are not necessarily physically demanding but that are experienced as exhausting and stressful nonetheless] and for predicting, at the design stage, the future performance of such [automatic and semi-automatic man-machine] systems... On the other hand the concept is at present very ill-defined with several probably distinct meanings... There is no satisfactory theory of 'mental workload'" (Johannsen et al., 1979:101). Johannsen et al stress the multidimensional nature of workload, and deny the appropriateness of trying to quantify it as a scalar variable. They specifically rule out the possibility of meaningfully comparing different tasks with respect to workload, except when the tasks are very similar in structure.

The conclusions drawn by the experimental psychologists in the NATO workshop clearly caution against any expectation that the problem of workload measurement will be resolved soon. They are equally clear, however, in supporting the view that workload is an essential concept if we are to understand the role of human beings in modern systems and design tasks that impose reasonable demands on their capabilities. It could prove to be an especially important concept in the context of the Space Station because of the unusual cognitive demands that that environment will represent. A detailed understanding of those demands--insofar as possible in anticipation of the deployment of the station--surely must be a primary objective of the human factors effort in this program.

One of the approaches that has been used to identify performance measures that are sensitive to workload has been to take a variety of candidate measures in situations in which workload is intentionally varied and see which of them vary with workload manipulation (Casali and Wierwille, 1983; Hicks and Wierwille, 1979; Wierwille and Connor, 1983; Wierwille et al., 1985). Much of this work has been done in flight simulators. Candidate measures that have been studied include opinion scales (subjects' ratings of the task in terms of specified descriptors), physiological measures (heart rate, respiration rate, pupil diameter, eye-blink frequency, eye-fixation fraction), measures of performance on secondary tasks (time estimation, tapping regularity), and measures of performance on the primary task. A limitation of this approach is that viable measures, at best, reflect differences in workload; they do not provide an indication of how hard or how close to capacity one is working in any particular case.

Many of the studies of pilot workload have made use of post flight questionnaires. Because this approach is heavily dependent on memory, Rehmann et al. (1983) explored the possibility of having subjects report how hard they are working periodically while performing a task. Workload judgements did change in this case with controlled changes in task difficulty, but this measurement technique has the disadvantage that it could interfere with the performance of the primary task, especially when the latter is very demanding.

The intrusiveness of the measurement process has been a major drawback of many approaches to workload assessment, and especially those that make use of a secondary task (Rolfe, 1971; for a summary of

nearly 150 studies using secondary tasks see Ogden et al., 1979). One way to avoid the use of an intrusive task and also dependence on the subject's memory is to monitor physiological indicants of workload that can be obtained automatically. Isreal et al. (1980) have argued that some of the physiological measures that have been tried; galvanic skin response, heart rate variability, and pupil diameter reflect changes in autonomic nervous system activity and so are sensitive to changes in emotional state independently of their origin. As a physiological measure that is more likely to be indicative unambiguously of changes in the cognitive demands of a task, they propose the event-related brain potential and, in particular, its late positive or p300 component. Wickens (1979) also has argued for the use of evoked potentials. Isreal et al. (1980) present data from one experiment supporting the idea that this measure does vary with task demands and that obtaining it need not interfere with the primary task. While it would be imprudent to conclude from these data that electro-physiological monitoring of workload will be effective in the Space Station, the possibility deserves further exploration.

Varying workload for experimental purposes is probably not feasible within the Space Station context, or at least the amount of this type of experimentation that can be done will probably be very limited. It will be essential to attempt to have workloads be as close to ideal as they can be made from the very beginning. Of course when evidence indicates that an initially established workload is not ideal, the workload should be changed in the indicated direction, and keeping track of such changes can provide some of the data that would have been obtained from controlled experimentation. The goal must be to minimize the need for such changes, however, which requires being able to predict the effects of different workloads from data obtained in earth environments.

DETERMINANTS OF PRODUCTIVITY

There seems to be a consensus among investigators that productivity is a function of many variables, and that attempts to affect it that focus on one or a small subset of those variables and ignore the others run the risk of doing more harm than good (Muckler, 1982; Sutermeister, 1976). Among the determinants of productivity that would have to be included in any extensive list are the following.

Human Capabilities and Limitations

A great deal of information has been compiled about human capabilities and limitations and is available in various engineering psychology handbooks. What is known in this regard clearly sets bounds on what human beings in general can be expected to do in specific task situations. Individual differences are also germane to the question of human productivity. People differ widely with respect to both physical and mental capabilities, and the productivity of individuals is bound

to vary with the degree to which their individual capabilities match the demands of specific tasks. Aptitude testing and job screening and selection procedures are based on these assumptions.

Task Demands

Evidence supports the intuitively appealing idea that people work best when the demands upon them are neither too great nor too small. This is one form of the "inverted-U hypothesis" regarding the relationship between workload and performance, which holds that performance of a given task is optimal for a workload level that is intermediate between one that is excessively high and one that is so low as to promote boredom (McGrath, 1965; Welford, 1973, 1974). The detrimental effects of overloading are somewhat better documented than are those of underloading (Weiner, 1975; Weiner et al. 1984). The possibility that underloading can affect performance negatively takes on special significance, however, in the context of systems in which humans function primarily as supervisors of automated processes.

Motivation

One can hardly doubt that motivation affects performance. It is clear in particular that performance suffers when motivation is very low. What is less clear is how performance is affected when motivation becomes extremely high. Modest increases in motivation that is relatively low at the outset will almost certainly lead to improved performance, but what happens when motivation that is already very high is increased still further? Is there such a thing as trying too hard? Wanting too badly to succeed? Some investigators believe there is, and that when motivation is extremely high it has a debilitating effect. This is another form of the inverted-U hypothesis mentioned above; except that in this case the performance determinant of interest is motivation rather than task demands. It may be that the detrimental effects associated with motivation becoming too high are better attributed to anxiety over the possibility of failing; fear, especially when it becomes panic, undoubtedly can cause performance to deteriorate. According to this view, if motivation becomes arbitrarily high but is not accompanied by such fear, we would not necessarily expect performance to fall off. The distinction between very high motivation and fear of failure may be an important one in the Space Station context; it would be helpful to have a better understanding of the roles of these variables as determinants of productivity and performance.

Physiological State

Fatigue has long been recognized as a factor in reducing productivity in many settings (Simonson and Weiser, 1976). Indeed it has been

defined operationally as a decrease in performance as a consequence of prolonged activity (Kalsbeek, 1971). Much of the research on this topic has focused on the problem of scheduling rest breaks in such a way as to minimize fatigue (Bechtold et al., 1984; Ganaro and Bechtold, 1985). The tasks involved in these studies have often been physically strenuous and the results are of limited applicability to tasks that are primarily cognitive in nature. Exceptions include studies of the performance of aircrews over extended periods (Cameron, 1971, 1973). A major question of relevance to productivity in the Space Station is how productivity might be affected by the various physiological effects that can be expected from prolonged living in the Space Station environment. Little is yet known about the physiological consequences of living in such environments for longer than a few weeks at a time.

Training

Performance, especially of complex tasks, obviously improves with training and practice. An aspect of the relationship between training and performance that is especially important relative to the Space Station context has to do with the obscuring of differences by ceiling effects. The fact that one has, through practice, gotten to the point of being able to perform a task without error is not compelling evidence that one has really mastered the task. The true test may come when that task must be performed under stress or in concert with competing demands on ones resources. To make the point another way, the fact that two people perform a given task equally well under accommodating conditions is not good evidence that they will perform it equally well under stress.

Capabilities and Limitations of Machines

Just as the capabilities and limitations of the humans in a complex system help determine the productivity of the system as a whole, so do the capabilities and limitations of the machines involved. Unlike the capabilities of human beings, those of the machines that are available for use in the Space Station can be expected to evolve even over the next few decades. Initial plans for the use of technology in the Station take this fact into account. Plans to use artificial intelligence, for example, explicitly note the unlikelihood that this technology will be used extensively for operational purposes during the initial years of the program. However, provision is being made for its incorporation as the technology matures to the point of being reliably applicable. We would expect that as machine capabilities are extended and improved, a major consequence would be increased productivity of the Space Station as a whole. Whether this proves to be the case and, if so, exactly how remain to be seen.

Person-Machine Function Allocation

An important determinant of system productivity, as distinct from both human productivity and machine productivity, must be the way in which system functions are assigned to people and to machines. Several methods for function allocation have been developed (for a review, see Price et al., 1982); but none of them is widely used by system designers (Montemerlo and Cron, 1982; Price, 1985). Investigators have argued that it is not realistic to expect it to be feasible to allocate function by formula anytime soon, if ever, because the problem is too complex and situation-dependent (Price and Pulliam, 1983). Allocations typically are made in an ad-hoc fashion on the basis of human judgment, aided perhaps by efforts of engineering psychologists, beginning with Fitts (1951), to distinguish between generic functions that machines do better than people and those that people do better than machines. While the number of functions that people can perform and machines cannot is likely to grow ever smaller with continuing advances in machine intelligence, it is likely to be some time before machines can match people in their ability to integrate information in so many forms from so many sources; to respond as effectively and adaptively to such a wide range of unanticipated situations; to make judgments of relevance, reliability and importance; to draw upon a large store of common sense, as well as technical, knowledge; and to follow imprecise instructions and work toward high-level goals. And if machines acquire such capabilities in time it does not follow that they should assume these functions in all cases. The question of what functions can be automated and that of what functions should be automated may have different answers. This fact has not received the attention it deserves. There may be reasons not to automate functions that are automatable from a technological point of view. These include reasons of cost effectiveness, human preference, and the need to maintain human skills at a high level in case they are needed in the event of system failure. One function that we can presumably assume will be a human one indefinitely is that of high-level goal-setting. Value judgments, including judgments of what goals are worth working toward, will hopefully remain the purview of human beings, no matter how clever the machines become. This probably means also, at least for the foreseeable future, retaining the role of deciding to what extent the behavior of the clever machines is consistent with those top level goals.

Design of Person-Machine Interfaces

In very complex systems like the Space Station, many functions are performed neither by people nor by machines independently, but by people and machines interactively. This being so, the adequacy of the designs of the interfaces through which information passes between the machines and their users will be a major determinant of productivity of the people, the machines, and the Space Station as a whole. The design challenge for the Space Station is complicated by the fact that the

intent is to accommodate a large fraction of the anthropometric spectrum. It is here, in the design of interfaces, that human factors researchers and engineers are likely to have the greatest impact on productivity. A great deal has been learned about interface design as a consequence of human factors research in other contexts (Nickerson, 1986). A significant general conclusion to be drawn from this research is that designers' intuitions uninformed by human factors research are often wrong. A second similarly general conclusion is that small differences in interface design can often have very large effects. This area deserves a great deal of emphasis in the Space Station program.

Organizational Factors

Gunn (1982:115) has claimed that, in the case of manufacturing, the major opportunity for improved productivity is not to be realized by mechanizing the work of making or assembling products, but rather "in organizing, scheduling, and managing the total manufacturing enterprise. The most important contribution to the productivity of the factory offered by new data processing technology is its capability to link design, management, and manufacturing into a network of commonly available information". Gunn's emphasis on the importance of a single integrated information system, serving various needs of a manufacturing operation, applies with as much, if not greater, force to the Space Station context. Information will be the life blood of the Station. How the information that supports the various functions will be organized and accessed will be a critical aspect of the Station's design. Problems of organization, access, updating, protection, and representation abound. How these problems are addressed is certain to have implications for productivity, which is not to suggest that those implications will be easy to make explicit.

Scheduling Factors

Scheduling is a particularly important problem for any operation that involves numerous interdependent processes that proceed partly in series and partly concurrently. The problem is exacerbated by the fact that an unanticipated delay in the onset or completion of any given process may have implications for the timing of other processes. Small perturbations can ripple and grow into major problems producing inefficiencies at best and sometimes serious difficulties. Dynamic rescheduling of multi-process operations of any complexity usually requires computer involvement. Producing the scheduling algorithms, however, is still a human activity and one that requires a great deal of ingenuity, if major inefficiencies are to be avoided.

Social and Interpersonal Factors

The linkage between social or interpersonal factors and productivity may be indirect, but there can be no doubt of its importance. Interpersonal difficulties among people who must work cooperatively as a group can seriously impair the smooth functioning of the group; conversely, when the members of the working group genuinely like each other and enjoy working together, there can be equally substantive positive influences. Interrelationships outside the working situation, and sudden changes in them, can also have profound effects. A new emotional relationship, illness or death of a loved one, an unresolved dispute with a friend or acquaintance are obvious cases in point. Such factors can affect performance not only through changes in morale or motivation, but also by diverting attention from the demands of one's job.

The above list of determinants of performance could easily be extended, but it is representative of factors that have been studied. Much is known about how these factors relate to performance and thus to productivity in earth environments. Much remains to be learned too, however, and while the themes may seem familiar, the new context of space gives the problems new dimensions. While all of these factors are likely to prove to be important in space, none represents a greater opportunity, and need for research than those that involve the way people will relate to and interface with machines, especially in view of the rapidity with which the capabilities of the latter are changing.

THE SPACE STATION

Anticipated Functions and Uses

The Space Station is expected to serve a variety of functions. These include serving as a laboratory for scientific experimentation and data gathering, manufacturing and processing of materials (e.g., crystals and pharmaceuticals), servicing of satellite and other space vehicles, providing a staging platform for other space missions, and serving as a base for constructing large structures for use in space. The station is viewed as being important not only to scientific and commercial enterprises but to the further development of space technology. Eventually the station is expected to serve as an extraterrestrial control and service center for numerous unmanned satellites orbiting in a variety of inclinations and altitudes. Serving as a control and maintenance center would include deploying, retrieving, repairing, and reconfiguring other satellites or spacecraft (JSC, 1979, NASA-ASEE, 1985). Considerable interest has also focused on the role the Space Station could play as a development and evaluation platform for automation, robotics, knowledge-based systems and other emerging technologies that make intensive use of computer-based resources.

Preliminary Design and Operation Considerations

The station is expected to evolve in at least two ways. As a physical plant it will increase in size and become more complex as modules are added and desirable modifications are identified. Operating procedures will also change as a consequence both of experience gained in operating it and of technological developments. In the interest of facilitating the evolution of the physical plant as new desiderata are identified, design plans call for modularity and expandability.

The living-working modules are an interconnected set of 4 pressurized cylinders, each of which measures 35-40 feet in length and 15 feet in diameter. The sizes of the modules are constrained by the requirement that at least the initial ones be prefabricated to fit in the cargo bay of the space shuttle. Two of these modules are to be living quarters and two are to be laboratories. Each living module will accommodate 6-8 people. A fifth module similar in design is called a logistics module and will be used for transporting equipment and supplies between earth and the station. Each of the modules is equipped with detachable units to facilitate reconfiguration, servicing and replacement.

Safety is, of course, a major concern. And this problem has the added dimension that mishaps that would have relatively minor consequences on earth could be disastrous in space. The possibility of fire in the spacecraft is a major worry for obvious reasons. This concern dictates many aspects of spacecraft design. Among the safety provisions that have been specified in preliminary design documents are: safe exit from any of the pressurized modules; isolatability of each module from the others; sufficient food, waste management, control and communications, and life support facilities in any three-module cluster to sustain crew and make rescue possible. Concern for safety also dictates that much of the training regimen focus on possible malfunctions.

In addition to the issue of safety, that of habitability is receiving considerable attention (Clearwater, 1985; Clearwater and Kasper, 1986). This issue becomes much more important for missions of extended periods than for those of a few days' duration (Wise 1985, 1986). The question is how to use color, texture, lighting, spatial arrangements, window placements, and other design features, within the constraints of other requirements, to make the various Space Station modules, and especially the living modules, pleasant places in which to spend long periods of time.

It is intended that the Space Station be as self-contained as possible. Consequently, much attention is being given to recycling of supplies, such as water, and to on-orbit maintenance and repair. Because the kind of constant and extensive ground control monitoring that has characterized short duration missions is not feasible for a permanent station, much attention is also being given to the objective of giving the station crew a high degree of autonomy and independence in its day-to-day operation. And because the intent is to make the station attractive to the private sector and useful for commercial ventures, the operating policies will have to take account of the

desires of the station's clientele. There is a strong interest in assuring human productivity in the Space Station environment, which stems in part from the anticipated high cost of manned flight.

Uniqueness of the Space Station

Newell, in the preceding paper, has highlighted thirteen "hard constraints" that may be expected to hold independently of the specifics of the station's design. The list makes clear the enormous challenge the Space Station program represents. It also points up the fact that the uniqueness of the space station environment stems not so much from any given constraint or small subset of them, but from the set as a whole. For any given constraint, one can point to one or more other environments or situations with which we have some experience that shares it (e.g. nuclear submarines, submersible laboratories, off-shore oil platforms, polar excursions, scuba and deep-sea diving, incarceration--prisoners of war--and time spent at sea by shipwreck survivors). Some of these environments or situations share several of the constraints in Newell's list, but none of them shares the entire set. This is an important point. Suggestive evidence regarding the expected effects of some specific constraints in the Space Station may be found in the results of studies of other environments that share those constraints; and situations that have been studied include extended submarine patrols (Weybrew, 1961; 1963) and wintering-over parties in the Arctic and Antarctic (Gunderson, 1963, 1974; Gunderson and Nelson, 1963). But extrapolating what is known about the effects of any given constraint or even small subsets of them may overlook important effects of interactions. It is not prudent to assume, in the absence of supportive evidence, that the effects will simply add. It is easy to imagine conditions under which constraints that individually would have minor effects would, in combination, produce major ones.

Many of the constraints in Newell's list have implications for productivity, either directly or indirectly through, say, morale. Multi-month crew residences and infrequent physical communication outside the station, for example, could result in feelings of isolation, deprivation or boredom, or interpersonal tensions among the personnel. Limited resources and space could become uncomfortably restrictive in time. Weightlessness can produce nausea, headache, stuffiness and other physical discomforts, as well as spatial disorientation.

If challenged to extend Newell's list of constraints to incorporate other characteristics of the Space Station environment that are likely to be especially important from the point of view of productivity, my candidates for consideration would include the following:

- High degree of interactivity, and especially cognitive coupling, between crew and equipment.
- Computer mediation of control actions and displays.
- Criticality of information systems.
- Need for aiding or augmenting of human thinking for

- troubleshooting and decision making.
- Importance of human-machine interface designs.
 - Need for continual concern for safety.
 - Need for ability to deal with unanticipated contingencies.
 - Shared responsibility of flight-control decisions between ground and flight crews.
 - Need for some operating procedures and principles to be negotiated with customers; in some cases, perhaps, while in orbit.
 - Heterogeneity of Space Station inhabitants (different languages, different cultures, different professions, different amounts of technical training and flight experience).
 - Importance of satisfying ways for inhabitants to spend free time.
 - Stress.

Each of these characteristics deserves attention as a variable that could have significant implications for productivity. Consider, for example, the second one. In the Space Station most of the control actions that are identified by humans will actually be effected by computers and most of the information provided to the human operators will be provided via computer-generated displays. Focusing only on displays, for the moment, it is easy to see how the ubiquitous computer mediation represents an important departure from more conventional displays. A major concern in the operation of any high performance vehicle is that of keeping the operator(s) aware of those aspects of the system's state that are critical to its operation. In conventional aircraft most indications of system state (altitude, bearing, airspeed, fuel reserve, etc.) are indicated by fixed displays each of which is dedicated to a particular indicant; when the pilot wants to check the plane's altitude, he looks at the altimeter, which is always in the same spot and displays always and only altitude: a little area of the cockpit is totally dedicated to the objective of keeping the pilot aware of how far off the ground he is. In the Space Station, most of the information that crew members receive will be delivered on computer driven displays that are used for more than one purpose. Display functions that were once implemented in hardware will now be implemented in software, and the type of information that is available in a specific spot on a control console will vary from time to time, depending on what piece of software is controlling the display device at the moment. This shift from hardware to software implementation of display functions has some implications for the problem of keeping crew members aware of system state.

Productivity in the Space Station

Productivity can have several connotations relative to the Space Station. It can refer to the impact of the Space Station program as a whole on the GNP or GWP. It can refer to the use of the Space Station by industry in production and manufacturing. It can refer to the

performance of individual humans or person-machine complexes. Also, there may be a diversity of goals relating to the measurement and control of productivity in the Space Station. It may be desirable, for example, to measure the productivity of an individual, a person-machine system, a team, or an entire station over some specified period of time. One goal might be to achieve some targeted productivity on average over extended durations. Another might involve being able to achieve peak productivity for short periods when needed.

Impact on National or Worldwide Productivity

Considerable emphasis is being put on the potential commercial uses of the Space Station and the assumption that it will have beneficial long-range effects on the economy of the participating nations. The 1986 report of the National Commission on Space, Pioneering the Space Frontier, proposes that the space program have three mutually-supportive thrusts:

- Advancing our understanding of our planet, our Solar System, and the Universe;
- Exploring, prospecting, and settling the Solar System;
- Stimulating space enterprises for the direct benefit of the people on Earth (p. 5).

The third of these thrusts is directly relevant to the idea that the space program could have implications for national and international productivity.

Whether productivity gains will be realized will depend, of course, on whether the savings due to better quality control more than offset the cost of getting materials to and from space and any other increases resulting from conducting the operations in a space environment. To have a significant impact on national or international productivity will require a continuing operation of considerable size. The impact on certain industries could be significant relatively quickly, however, if the cost effectiveness of space-based manufacturing is conclusively demonstrated.

The space program could also affect productivity on earth in a variety of ways. Exploration of the earth's atmosphere and surface with photography (e.g. LANDSAT) and other sensors can produce information that can affect productivity by producing a better understanding of weather patterns, energy sources, climatic trends, and so on.

Industrial Productivity in Space

The combination of zero-G and vacuum in space is expected to facilitate production processes for which it is critically important to control for convection forces or airborne impurities. Among the materials and products that are of interest in this context are "shaped crystals,

semi-conductors, pharmaceuticals, biologicals, strategic materials, plastics, films, oils, alloys and mixtures, ultra pure metals, composites, glasses, membranes, metal foam, fibers, microspheres, ceramic/metal, and matrix materials" (NASA-ASE, 1985:9). A major industrial interest in space is the prospect of growing superpure crystals (e.g. gallium arsenide) for semiconductors in an environment free of convective turbulences. Interest in conducting such operations in space stems from the assumption that the quality of the products will be much easier to control (Chaudhari, 1986). It is expected to be possible to grow much larger crystals, for example, and to have a much smaller reject rate.

Individual Productivity in Space

Individual productivity the effectiveness and efficiency with which the individual participants in the Space Station program carry out their assignments is of special interest to the human factors community, inasmuch as the other types of productivity are contingent to no small degree on how well individuals function in their various roles. All of the determinants of productivity mentioned earlier in this chapter represent important considerations for the Space Station, as they do for any complex system. The following are also among the more significant issues relating to individual productivity that are very likely to arise in this context.

- Redundancy and backup: Many of the functions performed by the Space Station crew will be sufficiently important that provision will have to be made for backup in case an individual becomes incapacitated. The necessity to rely on backup capabilities could have implications for productivity, depending on the adequacy of the backup procedures and the extent to which reliance on them has a ripple effect to other functions.
- Use of aids, intelligent and otherwise: There will be a need in the Space Station to augment human cognitive abilities in various ways. Decision-making aids, troubleshooting aids, memory aids, will be needed in various contexts.
- Error recovery: It must be assumed that in a manned Space Station human errors will occur. The standard approach to minimize disastrous consequences arising from such errors is (1) to attempt to build in fail-safe procedures so as to make it difficult to commit the errors in the first place and (2) to buffer operator actions--postponing their effects--so that when an error is made, there is an opportunity to correct it. There is an obvious tradeoff here between safety and short-term productivity. Fail-safe procedures and provisions for failure recovery are likely to slow operations down. In the long run, however, their costs may be more than offset by what they save if they prevent errors with serious consequences.

- Information accessibility: The operation of the Space Station is expected to be highly procedurized. While crew members may be assumed to have had extensive training in how to deal with various contingencies that may arise, it is not safe to assume that all the information they will ever need is stored in their heads. Availability of precisely the right information at specific moments could prove critical not only to productivity, but in some instances to safety or even survival. A recent report from a NASA sponsored workshop identifies a system that explains or assists in the use of other tools as perhaps the single most important tool from the standpoint of EVA autonomy and recommends the development of a real-time maintenance information retrieval system that could provide astronauts information on demand relating to EVA tasks as they are being performed (NASA-ASEE, 1985).
- Life-support systems: Although very great progress has been made in improving the design of space suits over the years of the space program, the suits currently in use for extra-vehicular activity still greatly restrict the wearer in various ways.
- Morale: Excepting complications arising from motion sickness, morale has not been a major problem affecting performance of crews in flight in the space program thus far. But the publicity surrounding the flights and the relative brevity of their durations have probably sufficed to keep the morale generally high. When people are in space for months at a time and the work becomes less of an adventure and more of a job, it will not be surprising if morale becomes an issue, and one that could affect productivity, from time to time.

In addressing these and related issues, it is useful to bear in mind that while the Space Station will differ from other environments in numerous ways, many of the issues that relate to productivity in this environment are of more general interest because of their relevance to earth environments as well. The question of how various types of information are best represented on computer-driven displays is a very general one, for example. And it takes on considerable practical significance in view of the fact that 40 to 50 percent of all American workers are expected to be using electronic terminal equipment on a daily basis by 1990 (Giuliano, 1982). Unquestionably designers of Space Station displays should benefit from the many ongoing efforts to package information more effectively for use in office, industrial, and other earthbound contexts; we expect also, however, that efforts to get the Space Station displays just right--because being almost right may not be good enough in this context--will yield knowledge about display design that will advance the state of the art in a general way.

The Evolving Role of Humans in Space

There has been and continues to be a debate about the advantages and disadvantages of a space program that includes manned spacecraft as opposed to one that does not. That debate will not be rehearsed here, beyond noting that opponents of a manned program have argued that having humans in space is unnecessary for many aspects of space exploration and providing for their safety delays the program and increases its costs (e.g. Van Allen, 1986a, b) whereas proponents of a manned program have presented a variety of arguments in favor of it, among them our inability to provide machines with some human abilities that are seen as critical, especially in responding to unanticipated events. Of particular relevance in the present context is the argument that has been made that the presence of humans in space will contribute positively to the productivity of the program as a whole. In this paper a manned program is taken as given. The problem then becomes that of designing a Space Station environment and operating procedures that will insure both the safety of the crew and the success of its missions.

The human's role in space has expanded and diversified over the life of the space program (Loftus et al., 1982). In the earliest flights the role was primarily that of passenger in a highly automated or ground-controlled vehicle. As experience was gained and the flights became more ambitious the crews took on more of the responsibility of piloting the spacecraft. Still later, the crew's role was enlarged to include functions unrelated to piloting, such as performing scientific experiments and repairing malfunctioning equipment.

Specific tasks that have been performed by crew members include monitoring of the various spacecraft subsystems (guidance and control, propulsion, environmental control, and life support); guidance and control during rendezvous and docking; landing and taking-off of lunar module (about 10,000 key strokes are required to complete all elements of a lunar landing mission, according to Loftus et al., 1982); assembly, maintenance and repair (especially of scientific instruments); aiming of scientific instruments and conducting of experiments; monitoring of data quality; and housekeeping.

The ability of the crew to perform maintenance and repair operations and to handle unexpected subsystem failures of various types has been demonstrated in several missions, including Gemini, Apollo 13, Skylab, and Spacelab (Garriott, 1974; Garriott et al., 1984). Especially in the Skylab and Spacelab programs crewmen on numerous occasions were able to repair malfunctioning equipment that was essential to the planned experiments. As Garriott et al. (1984) have suggested, the importance of the function should be reflected in the training of the crew designed to familiarize them with the equipment and how to repair it.

The ways in which the crews participated in the research activities of the Skylab and Spacelab programs have also been reviewed by Garriott (1974) and Garriott et al. (1984). An important idea emerging from these reviews is the following one. To the extent that crew members are to act in behalf of scientific investigators located on the ground,

this function may go more satisfactorily if there has been more opportunity for the crew members to work with the scientists prior to the space mission.

As the human's role has expanded and diversified, the need for specialized capabilities and talents on space crews has increased, and consequently the crew members are less and less interchangeable. In the Shuttle program, specialization is recognized explicitly in the terminology, which distinguishes between mission specialists and payload specialists. In prolonged flights, like those anticipated for the Space Station, there will be an even greater need for certain specialized skills than has been the case heretofore. It may be necessary, for practical reasons, to have specialists who are also able to function effectively as crew members outside their area of specialty.

An important problem in planning the crew requirements for the Space Station will be that of assuring that collectively the crew has all the knowledge and skills that success and safety will require. What is difficult about this task is specifying the knowledge and ingenuity that will be required to deal with whatever unexpected contingencies arise. While it is not possible, of course, to anticipate everything that could happen, one step that can be taken in this direction is to attempt to identify the major types of problems that could arise (e.g. problems in the station's electrical system, medical problems among the crew, etc.) and to make sure that there is expertise within the crew to deal with problems in those areas.

Some of the activities the Space Station's crew will perform will take place outside the spacecraft. Such extravehicular activities (EVAs) may include the changing of focal planes and other servicing of the Hubble Space Telescope (HST), the Gamma Ray Observatory (GRO), the Advanced X-Ray Astronomy Facility (AXAF), and the Shuttle Infrared Telescope Facility (SIRTIF). (For a tabular summary of extravehicular activity on spaceflights through the Skylab III, see Loftus et al., 1982.) A major component of the cost of EVA activity stems from the large amount of time required to make the transition from the environment inside a pressurized space capsule to that outside it (Howard et al., 1982). Pressure inside the Space Station is 14.7 psi; that in the pressurized suit is 4.3 psi (King and Rouen, 1982). Because of the magnitude of this difference it is necessary for astronauts, in order to avoid the bends, to clear out the nitrogen in their body tissues by breathing pure oxygen for 3 or 4 hours before exiting the spacecraft. This procedure could be eliminated if the pressure maintained by the suit were above approximately half that maintained inside the cabin; thus immediate exit upon donning a space suit would be possible if either suits were designed to maintain 8 psi and cabins were kept at 14.7 psi as they currently are or cabin pressure was maintained at about 8 psi and suits at 4.3 psi, as they now are (NASA-ASEE, 1985).

Extravehicular activity represents a special challenge with respect to productivity for a variety of reasons, including the following:

- Severe constraints on mobility and dexterity imposed by the

- pressurized space suit.
- Limited visibility due in part to restrictions on head movements from the helmet and space suit.
 - Greatly reduced tactile feedback to the hands because of pressurized gloves.
 - Free floating or tethered (and easily tangled) tools.
 - Limited voice communication with in-station crew.
 - Problems associated with personal hygiene and comfort; most serious perhaps are the problems of defecation for males and defecation and urination for females, but the general problem surfaces in numerous other, perhaps less serious, guises as well: it is very difficult to scratch one's nose or any other itch in an EVA suit.
 - Problems of eating and drinking.

To the degree that the Space Station is an automated system that is monitored by human beings and dependent on manual override in case of subsystem malfunctions, it will pose the same kinds of challenge as other systems of this type. One such challenge is that of assuring that the human monitors are adequate to the task. The monitoring and controlling of dynamic systems are quite different tasks, and there is some evidence that people who have not had experience as manual controllers are less effective at detecting small changes in system dynamics than are those who have (Kessel and Wickens, 1983; Wickens and Kessel, 1979, 1980; Young, 1969). Another challenge relates to the dependence on human monitors for back up in case of system failure, and that is the problem of maintaining the human skills needed to perform complex functions that are very seldom performed under normal operating conditions. How does one keep crew members highly skilled at complex tasks that they seldom, if ever, have to perform? According to Jones et al., (1972), the most important functions aboard present spacecraft involve diagnosis and decision making, and retention of diagnostic and decision making skills represents our greatest gap in knowledge about task retention at the present time.

A major challenge for extended space missions, especially those involving long periods of time simply getting to a destination (e.g. interplanetary travel) will be to keep a crew and other inhabitants of the space vehicle occupied in meaningful ways when there is little essential work relating to piloting or maintenance of the vehicle to be done. Work that is invented just for the sake of killing time is unlikely to be very satisfying. It will be important for individuals to perceive their tasks as serving some useful purpose. Some time will have to be spent in doing housekeeping chores and some will be viewed as leisure, but it will undoubtedly prove to be necessary to have significant fractions of most days occupied with activities that are perceived as important to the mission or to other valued goals. Scientific experimentation and research could occupy much of this time, at least for those individuals who are scientists by profession or who would derive satisfaction from participating in scientific work.

The problem of leisure time is considerably more complicated for extended missions than for those of short duration. In the former

case, one must be concerned not only with provision of short periods of free time at frequent intervals (e.g. daily) but also with the need for something analogous to holidays or weekends and vacations on earth, and with the question of how to ensure that individuals find it possible to spend that time to good advantage both from their point of view and that of the mission.

The Close Coupling of Humans and Computers

In 1983, the Space Systems Division of the NASA Office of Aeronautics and Space Technology convened a summer workshop (co-sponsored by the American Society for Engineering Education) at Stanford University to study the role of autonomy in space. The workshop report was issued in 1985, and has been referenced here as NASA-ASEE, 1985. Participants in the workshop included professors from universities across the country. "The workshop sought to generate recommendations on autonomous systems and human functions as well as on a space technology program directed toward symbiotic use of machines and humans"....."The principle objectives of the 1983 summer study were to examine interactions of humans and highly automated systems in the context of specific tasks envisioned for the space station, to search for optimum combinations of humans and machines, and to develop methodologies for selecting human-machine systems" (NASA-ASEE, 1985:2).

Participants in the workshop concluded from their study "that machines will not replace humans in space and that artificial intelligence (AI) systems will not have major impact on initial station design." To be sure, some aspects of the operation of the Space Station--maintenance of orientation, control of in-station environment, pointing of antennas and solar panels--will be done completely automatically, at least under normal circumstances. Moreover, the role of automation and artificial intelligence will increase as these technologies mature. But for the foreseeable future, and perhaps indefinitely, a great many aspects of the operation of the Station and of the performance of various missions will require the interaction of people with machines.

An increasingly common mode of interaction will involve supervisory control, which is viewed by some as intermediate between the use of teleoperators on the one hand, and robots on the other (Thiel and Kurtzman, 1983). In the case of teleoperators, the human has a "virtual" hands-on relationship but at a distance, as it were. In the case of robots, the relationship is of a qualitatively different type and may be remote both with respect to distance and time. The robot is given a capability by its designer to function relatively autonomously, albeit in accordance with principles incorporated in its design. In the case of supervisory control, the human is linked to the machine in real-time, but controls its operation only at a relatively high level. The human provides generic commands, which the system then translates into lower-level commands to the effectors that will, if all goes well, get the job done. How generic the commands are that the human operator

provides depends on the system. The higher the level, the closer one comes to robotics, and at some point the distinction between the two modes disappears.

The fact that so many of the functions in the Space Station will be performed by people and machines in interaction means that the design of the various workstations and person-machine interfaces will be of central importance. There exists a substantial literature, much of it in design-guide form, that is highly relevant to this problem and that should be a major resource for designers of Space Station workstations and displays. But because the Space Station will be extending the frontiers of technology in several ways, designers will also have to consider questions for which answers have not yet found their way into design guides, and in some cases may not have yet been asked. Moreover, as Loftus et al. (1982) point out, the ultimate design objective of any manned spaceflight program is never that of optimizing the crew-to-spacecraft interface, but rather that of achieving overall program effectiveness; and given the numerous constraints within which such programs must function, this may mean that compromises will be necessary in various interface designs. Decisions about such compromises, and selections among various possible tradeoffs, should be made with the best understanding possible of their implications.

Among the issues relating to workstation and interface design that will be of special concern in the Space Station context are the following:

- How to design multifunction input-output devices so as to preclude confusion among functions.
- How to lay out the various display and input devices so as to ensure ease of location, interpretation and use.
- How to design the control and feedback interfaces for teleoperator systems.
- How to design the various input-output procedures (command and query languages, menus, abbreviations, symbols) so as to maximize their usefulness and minimize human error.

Many of the human factors issues relating to the design of workstations and interfaces will center on the question of how to get information--precisely the right information in a useable format and at the appropriate time--from a person to a machine or from a machine to a person. So in addition to the important questions of the physical designs of displays and input devices, there will be many issues relating to the design of methods and procedures for interacting with information per se. When will it make sense to use query languages as opposed to menus for searching a data base? Query languages put a greater learning burden on the user than do menus, but probably are faster for experienced users, because menus typically force one to go all the way down a tree step by step even when one knows precisely what one wants to ask at the beginning.

When menus are used, how should they be structured? This question subsumes a host of others, and although the lower-level questions sometimes seem to have intuitively obvious answers, research often

reveals them to be more complicated than they appear. Consider the apparently simple matter of deciding how many items to show on a single node of a menu hierarchy. For a system with a given number of possible end points, there is a tradeoff between the number of options one sees at a given node in the hierarchy and the number of nodes required to get from the start to the finish. This breadth-versus-depth tradeoff has been the focus of some research (Dray et al., 1981; Miller, 1981; Seppala and Salvendy, 1985). While the results have not led to an unequivocal conclusion, there seems to be some agreement that menus that have very few items per level (say less than four) tend generally to be inefficient (Lee and MacGregor, 1985; Seppala and Salvendy, 1985). The situation is complicated, however, by the fact that how much breadth one can handle effectively will probably depend on how much experience one has had with the system. This may be an argument in favor of permitting a menu structure to modify itself to match the experience level of its user.

Much research effort is currently being devoted to the development of natural-language front ends for information systems. It seems likely that natural language systems with limited but useful capability will be available by the time the Station is operational. This is not to suggest that the reality of natural language capability will make other modes of interaction obsolete. The assumption that natural language would be the preferred mode of interaction with a data base in all cases is not beyond question; there is some evidence that more structured and constrained query languages may give superior performance in certain instances (Small and Weldon, 1983; for a review of human factors considerations that pertain to the design of query languages, see Ehrenreich, 1981).

Speech is also becoming increasingly feasible as a mode of communication between people and machines and could find at least limited use in the Space Station. The technology for synthesizing speech is improving steadily and although the best synthetic speech is still noticeably different from human speech and typically somewhat less intelligible, people get quite good at understanding it with only modest amounts (a few hours) of listening (Schwab et al., 1985). Speech understanding by computer is not so far along, but progress there is also being made. The technology for isolated word recognition probably is sufficiently mature to be used in a Space Station context, and more ambitious uses of speech understanding technology may be feasible by the time the Station becomes operational.

Stress and Performance In Space

Efforts to anticipate how humans will perform on extended space missions have focused on certain ways in which the space environment differs from more familiar environments on earth and on various types of stressors that could have either acute or cumulative long-term effects. Some of the characteristics of the Space Station environment may themselves be stressors, if not continuously, at least under certain conditions. It will be convenient, therefore, to begin this

section with a brief discussion of stress in general terms and then to consider specific environmental characteristics or stressors that might be expected to affect performance and hence productivity significantly.

Effects of Stress on Performance

Stress is likely to be a factor in the Space Station and to affect productivity in several ways. First, under the best of circumstances the Station and its personnel are always at risk. While we would not expect individuals to spend every waking moment worrying about safety, it would be surprising indeed if there were not a constant underlying sensitivity to the tenuousness of the situation; this might be considered a type of chronic stress. Second, from time to time, individuals or the entire occupancy of the Station may be stressed acutely as a consequence of an unanticipated event or situational change. Third, stress may also be caused by factors that are relatively long lived, but not necessarily chronic. These include confinement and social isolation, sensory-motor restriction, interpersonal frictions, dissatisfactions with certain aspects of one's duties or the Station's operating procedures, and anxieties about events or situations on earth. The list of possibilities is easily extended.

According to Sharit and Salvendy (1982) most of the definitions of stress that one finds in the literature reflect biases related to the scientific orientation of the writers and fail to capture the many-faceted nature of the phenomenon. Fidell (1977) has noted that some authors who have written about stress have avoided defining the term (e.g. Broadbent, 1971; Welford, 1974) presumably on the assumption that the word is intuitively meaningful: most of us know what it means to be stressed from personal experience.

In his review of effects of stress on performance, Fidell (1977) classified stressors as physical, physiological, psychological, and social. (Lazarus and Monat (1977) used the last three of these categories but not the first.) In the first category Fidell included thermal (heat, cold, humidity) mechanical (vibration, acceleration, fluid pressure) and sensory (noise, glare, odor, deprivation) and ingested or inhaled substances (drugs, noxious fumes, insufficient oxygen). As physiological stressors he listed musculoskeletal fatigue, sleep deprivation, age, disease, and illness. As psychological stressors he distinguished between cognitive (information or perceptual under/overload) and emotional types (fear, anxiety, insecurity, frustration). The social stressors in his list were occupational factors (e.g. career pressures) organizational structures, major life events, crowding, and solitude. Fidell also pointed out that stress is sometimes thought of as an effect and sometimes as a cause. It is assumed to be an effect, for example, of a perceived threat to one's safety or the imposition of a task that exceeds one's ability to perform. On the other hand, it is sometimes identified as the cause of poor performance or of otherwise inexplicable behavior. It is also sometimes viewed as the cause of certain types of medical problems such

as ulcers, colitis, and cardiac arrhythmias.

Effects of stress on performance are not easy to summarize. Mild to moderate stress for short durations can have a beneficial effect in many situations, possibly as a consequence of increased alertness and the energy spurt that comes with the greater-than-normal production of adrenaline and other hormones. Excessive stress can produce deterioration of performance. Frequent experience of stressful events tends to be accompanied by atypically high incidence of illness of various sorts (Norman et al., 1985). A relatively unexplored aspect of effects of stress on performance relates to how performance changes after a temporary stressor has been removed.

The study of effects of stress is further complicated by the fact that people adapt or accommodate to stressors, especially if they are only moderate in degree and relatively invariant over time. Noise, for example, can be a stressor, but people who work in a continuously noisy environment seem to adapt to it so that its effects as a stressor diminish greatly or disappear. Unexpected substantive change in the level or characteristics of the noise, however, may have disruptive effects.

Leventhal and Lindsley (1972) distinguish between danger control and fear control as two types of concern that one may have in a threatening situation. Concern for danger control is focused on the threatening situation and on how to rectify it. Concern for fear control is focused on the fear response and on how to keep it in check. Both are legitimate concerns and training in preparation for extended space missions should take both into account.

Stress is likely to be an important factor in the Space Station and its effects on productivity could be substantial. Moreover several stressors may be operating simultaneously, producing complex interactive effects, and the stressors will be interacting also with other variables in ways that cannot be foreseen. In the remainder of this section, several of the stressors that could be especially important in the Space Station environment are briefly noted. Exactly how these factors, especially in combination, will affect performance and productivity is not known; that their effects will be substantive, however, seems highly likely.

Weightlessness

Weightlessness has been emphasized as a major feature of a spacecraft environment that could give rise to physiological problems such as altered fluid and electrolyte balances, and deconditioning of specific systems such as the cardiovascular, musculoskeletal, metabolic, and neuroendocrine systems (Lindsley and Jones, 1972). Problems of these types have not yet been shown to be sufficiently severe to preclude prolonged space missions; on the other hand, how they will manifest themselves in long duration missions remains to be seen.

In retrospect many, perhaps most, of the observed short-term effects of weightlessness on human functioning probably were predictable, but many of them were not predicted. In thinking about what it would be

like in a weightless environment, one may find it easy to imagine being able to float freely in space and fail to realize that it will also be difficult to stand on the floor, sit in a chair, or maintain any fixed position without restraints. Who would have thought to ask whether it would be possible to burp? Or whether it would be difficult to bend down to tie one's shoes?

Unfamiliar Motion

Closely related to weightlessness are the various types of motion that can produce motion sickness (Kennedy, 1972). Even astronauts who have had training intended to reduce the probability of motion sickness have experienced such sickness during space flight, usually during the first few days of a mission, although nausea has typically not precluded crew members from carrying out essential activities (Garriott, 1974). There is some indication that dizziness is more likely to be induced in situations that permit individuals to move around in large spaces than in those in which they are more confined (Berry, 1969, 1970). When severe, motion sickness can be debilitating.

Motion Restriction

On the opposite end of the spectrum from the concern for unfamiliar motion is that for motion restriction. A variety of restrictive conditions on earth have been studied with a view to determining their physiological and psychological effects. These include immobilization from a plaster cast, bed rest, and prolonged confinement in submarines, space cabin simulators or other chambers (Fraser et al., 1972). Among the most apparent physiological effects of long-term restriction of activity appear to be cardiovascular and musculoskeletal deconditioning, including some bone decalcification. Other possible effects include electrolyte imbalances and hemolytic anemia.

As measures that can be taken to prevent or counter the deconditioning effects of motion restriction, Fraser et al. (1972) list the following: adequate free living space (200-250 cubic feet per person at a minimum, up to 600-700 cubic feet per person as the "optimal, maximizing habitability in the light of other requirements"), adequate exercise, applied pressure (to control for fluid volume loss and orthostatic intolerance of deconditioning), artificial gravity (seen as expensive and therefore less practical than other approaches), and hormones and drugs (primarily to control fluid loss).

Sensory and Perceptual Restriction

What is known about the effects of sensory and perceptual deprivation or restriction on human performance has been summarized by Schultz (1965) and Zubek (1973). Eason and Harter (1972) have also reviewed the literature on this topic through 1972 and attempted to extract from

it information that would be relevant to the prediction of human performance in prolonged space flight. (Sensory deprivation or restriction connotes an absence or marked attenuation of sensory input to one or more modalities; perceptual deprivation or restriction suggests reduction in patterned stimulation.) Eason and Harter noted that the studies available for their review did not include any in which the period of confinement or isolation exceeded a few weeks. Russian investigators have done studies on effects of confinement in which subjects spent as long as one year in relatively isolated environments but details have not been available.

The data from these studies are fragmentary at best and do not constitute a coherent set of findings. Results of individual studies are often mutually contradictory, some showing negative and some positive effects of deprivation on subsequent perception or performance. As they relate to long duration space missions, Eason and Harter (1972:101) see the findings as "rather heartening, for they suggest that the effects of severe sensory or perceptual restriction, isolation, and confinement are so minor, except in a few instances, that they are difficult to demonstrate with any degree of consistency not only from one laboratory to another but often within the same laboratory".

Eason and Harter caution against making predictions about an astronaut's sensory, perceptual and motor functions during long-range missions on the basis of experiments involving relatively short-term isolation. The results of such studies do provide a basis for raising questions and suggesting directions for research that can be relevant in the space flight context, and had they yielded solid evidence of large effects of isolation on sensory or motor functions, they would have raised some concerns about potential effects in the Space Station program. "As it turns out, the results of studies summarized in this paper suggest that only minimal and relatively insignificant changes in sensory and motor function are likely to occur during long-duration missions" (Eason and Harter:103). Eason and Harter point out that in extended space flight, boredom from repetition of stimulation may turn out to be a more important determinant of performance than sensory deprivation as such. They note, however, that past studies have been too limited in various respects to provide a basis for confident predictions about possible effects of confinement and isolation in space flight and urge further study of these variables under conditions that will assure the applicability of the results.

Sleep Interference

Sleep disturbances and irregularities take many forms. The most obvious departure from a typical sleep-wake cycle is total sleep deprivation--going for extended periods of time without any sleep. Other types of irregularity include unusual cycles (e.g. 4 hours of sleep, 4 hours of work), change in phase in the normal sleep-wake pattern (e.g. shifting from a work-in-the-day-sleep-at-night pattern to a sleep-in-the-day-work-at-night pattern), disruption of the quality of

sleep (fitful or shallow sleep; decrease in stage-3 and stage-4 sleep) resulting from environmental disturbances, psychological stress or other unusual factors. Studies of shift workers have shown that changing from day to night shift typically results in a reduction (1 to 2 hours) in the duration of the main sleep period, an increase in average total amount of sleep per 24 hour period--due to naps taken outside the main sleep period and extra sleep on rest days--and a change in the quality of sleep (Akerstedt and Gillberg, 1981; Tilley et al., 1981; Tilley et al., 1982). Indicators of quality include time to sleep onset, number of awakenings, number of body movements, and number of changes in sleep stage (Johnson et al., 1972).

How sleep disturbances affect performance is not understood well. Data suggest that sleep loss is likely to have deleterious effects on tasks for which sensory stimulation tends to be low and the rate of data handling is not under the individual's control (e.g. monitoring or vigilance tasks) and to have less effect on the performance of complex intellectual tasks involving problem solving and logical analysis (Johnson et al., 1972). Somewhat independent of the question of the effects of sleep disturbances on performance is that of their effect on moods and attitudes. Insomnia is often linked to depression, tension, and irritability. Whether there is a cause-effect relationship and, if so, which way it goes, are not known for certain.

Determination of optimal work-rest cycles will involve consideration of a variety of factors, technological, psychological and social. How often and how long people will need (or want) to sleep will depend in part on the demands of their jobs, and in part on the conditions of the sleeping environment. Requirements for sleep are likely to differ from person to person. With respect to social factors, there is some evidence that crews prefer to be on the same work-rest cycle insofar as possible, and work and get along better when this is the case.

The importance of rest periods interspersed among work tours has been known at least since Taylor's (1947) early studies. Exactly how rest breaks should be scheduled, however, or how this should depend on the nature of the work being done, has not been established very precisely. It is not even clear that it is always optimal for work breaks to occur on a fixed periodic schedule.

Any attempt to understand the relationship of sleep disturbances and stress will illustrate the problem of distinguishing cause from effect. Sleep disturbances, such as those caused by unusual work-rest cycles or the need for prolonged wakefulness to deal with an emergency situation are seen as sources of both physiological and psychological stress. On the other hand, stress originating from other sources can be the cause of insomnia or other sleep-related difficulties.

Boredom and Other Motivational Problems

It is somewhat paradoxical that one of the major concerns about such a risky venture as extended space flight should be a concern about boredom. However, boredom and various attendant complications could be among the most serious problems that have to be faced. Although

surprisingly little empirical work has been done on boredom (Smith, 1981), it has been identified as a significant problem for people living in restrictive environments with monotonous schedules for weeks or months at a time. It is believed to have detrimental effects on motivation and morale and to lead to increased frequency of complaints of headache and other physical problems. The tendency for motivation to decrease over a period of extended confinement is a common report from studies of small groups in isolated environments (Smith, 1969).

Behavioral evidences of a loss in motivation include diminution of one's ability or willingness to engage in sustained purposeful activity. There is some evidence that declining motivation has a physiological correlate in a decreasing frequency of alpha rhythm in the EEG wave (Zubek et al., 1969). This is an interesting finding because it suggests the possibility of using alpha rhythm as a means of monitoring individuals' momentary cognitive state and of predicting how productive they are likely to be in specific work situations.

Many studies have failed to find a decrement in ability to perform some types of cognitive tasks--and in some cases have even found an improvement in that ability--as a consequence of spending substantial amounts of time in confined environments. However, Johnson et al. (1972) note the possibility that studies that measure performance under the circumstances in which motivation might be expected to be low often risk artifactual results by virtue of the possibility that the experimental task itself, if unusual within the context, may be sufficiently arousing and rewarding to improve temporarily the subjects' motivational state. After reviewing the pertinent literature, Johnson, Williams and Stern concluded that very little is known about how to reduce monotony and boredom during long periods of group confinement.

Social Isolation

Isolation can mean a variety of things. Brownfield (1965) identifies four: spatial confinement; separation from persons, places, or things that one values highly; reduction or restriction of sensory stimulation; and reduction in the variability and structure of stimulation. The first, third and fourth of these connotations have already been mentioned. Unfortunately, effects of isolation often cannot be distinguished from those of confinement, motion restriction and social crowding, because these conditions typically occur together; nevertheless, it is believed that social isolation could prove to be among the most important stressors in the context of prolonged space missions. Some concern has been expressed that it, combined with some of the other characteristics of the space environment such as weightlessness, empty time, and distortion of the usual balances among sensory inputs, may lead to an increased frequency of daydreaming and fantasizing and a progressively more subjective orientation (Leventhal and Lindsley, 1972). Studies of groups that have spent extended periods (months) in relative isolation have shown that individuals tend over time to withdraw and become more psychologically remote from other

members of the group (Haythorn et al., 1977). According to Sells and Gunderson (1972:204), extended isolation and confinement of small groups on earth (e.g. at scientific stations in Antarctica) can increase the probability of "irritability and depression, sleep disturbances, boredom, social withdrawal, dissatisfaction, and deterioration in group organization and cohesion". Enriching the stimulus environment can counteract this tendency to some degree, but the stimuli must be meaningful and of interest to the people involved. There is some evidence that part of the withdrawal complex is a decreased tendency to avail oneself of whatever opportunities for stimulation the environment provides.

Special problems may arise when an individual especially close to a person on an extended mission becomes seriously ill (e.g. a child, spouse, or parent) and it is impossible for the person to return to earth, or if unanticipated events of major significance occur on earth during a prolonged mission. The effects of such happenings on attitudes and morale could be substantive. It is easy to imagine other examples of events on earth that could prove to be stressors to people in space. Inasmuch as communication between earth and the station will probably be primarily through ground control stations, at least for some time, information that could have a detrimental effect on the morale of members of the Space Station crew could be withheld from them. Consideration of such a policy raises a serious ethical issue, however, and would probably not be tolerated in any case. There are many reasons for maintaining frequent, if not constant, communication with earth. Not least among these is the need for inhabitants of the station to communicate frequently with people other than themselves.

Excessive Workload

Excessive task demands can be a source of stress and can lead to serious performance decrements. When even moderate task demands are coupled with the constant possibility of catastrophic errors, long term exposure to the situation can produce a variety of stress-related symptoms. One inherently stressful job that has been the focus of considerable attention by researchers, and the general public as well, is air traffic control (Cobb and Rose, 1973; Crump, 1979; Finkelman and Kirschner, 1981; Hailey, 1968). The stress in this case probably stems in large part from the facts that errors in performance can result in human fatalities and that most aircraft accidents are due to human error (Danaher, 1980).

Task demands in the Space Station are unlikely to be excessive for sustained periods of time, although they could be high at critical mission junctures and could become excessive during emergencies. Perhaps more important is the ever-present possibility of human error having a catastrophic result. Every attempt will be made, of course, to ensure that the operating procedures are fail-safe and that any errors that can be anticipated are recoverable, but some degree of uncertainty in this regard is bound to remain, and with it some level of task-induced stress.

Acute Medical Problems

With respect to the control of medical problems within a spacecraft, the emphasis has to be first on prevention (Fraser et al., 1972). Having taken all reasonable preventive measures, however, the chance that medical problems will arise on any long-duration mission is high. Within the Space Station there will be the possibility of many of the same types of physical injuries arising from accidents with equipment that might occur on earth. In addition there are certain types of mishaps that are relatively unique to the space environment; these include the aspiration of particles that float in the weightless environment of the station, effects of prolonged exposure to atypical mixes of atmospheric gases or pressures, exposure to high-Z particles--high energy particles of high atomic number--or other forms of radiation, and heat disorders resulting from malfunctioning of a pressure suit during EVA. Fraser et al note also the possibility that some medical problems that would be very easy to treat on earth could become significant in space, either because of inadequate treatment facilities (e.g. acute appendicitis) or because the medical problem has been complicated by virtue of various ways in which the body has adapted physiologically to the weightless environment (e.g. reduction in blood volume due to weightlessness).

Other Sources of Stress

Other features of space flight that could also be problematic include the absence of normal terrestrial time references, and possibly altered magnetic fields (Fraser et al., 1972). Changes in lines of authority that could prove necessary from time to time could pose challenges for social stability of the spacecraft community. The need for privacy could be an especially important one in extended space flight; the ability to have some time and place wholly to oneself on a fairly regular basis may prove especially important in this environment. Sharing of sleeping quarters and other personal space over long periods of time can increase the frequency and seriousness of interpersonal frictions. Habitability of the spacecraft will increase in importance with increases in the durations of space missions. The difficulty of maintaining a habitable environment will also increase with mission duration.

It will be particularly important that inhabitants of the Space Station be able to resolve, quickly and expeditiously, any interpersonal conflicts that arise. Presumably selection procedures will disqualify from participation in space missions individuals for whom the probability of interpersonal disputes or frictions is determined to be high. It will be important for those who do qualify to receive such training as is available regarding how to avoid various types of interpersonal disputes, and how to resolve them when avoidance proves to be impossible.

Individuals react differently to the same stressors, depending on

motivation, familiarity with the situation, appropriateness of training, degree of confidence in own ability to cope, degree of confidence in supporting colleagues and accessible resources, and other factors. There is some evidence that the magnitude of physiological reaction (e.g. increased pulse rate) to psychological stress is likely to be less for individuals who are aerobically fit than for those who are not (Holmes and Roth, 1985). Tests that provide a reliable indication of how individuals will react to the types of stressors they are likely to encounter in the Space Station environment would be useful both for purposes of selection and for identifying specific training needs. Development and validation of such tests are worthwhile goals. Similarly, development of more effective methods of increasing tolerances to specific stressors and of improving the ability of individuals to function effectively in spite of them should be continuing objectives.

Loftus et al. (1982:II-34) note that stress does not seem to have led to performance degradation so far in the spaceflight program. They attribute failure to observe such degradation "to substantial overtraining of flight crews for the tasks they must perform, diverse and interesting stimuli present in the real environment contrasted with minimum stimulation environment in simulations, and stronger motivation in flight crews compared with test subjects". It would be unwise to extrapolate the relative unimportance of stress as a determinant of performance in the early space program to the future, however; the much longer durations of the missions and the inclusion of participants who are not professional astronauts are two major differences that could make stress of various types much more consequential.

CONCLUSIONS AND RECOMMENDATIONS

The Space Station program is an ambitious undertaking. Establishing a permanently manned facility in space will be expensive and risky, but the long-range benefits for humankind that could result from success in this endeavor are surely very great. Keeping the program moving forward without unpleasant surprises and major setbacks will require intensive planning, continual evaluation of plans, replanning based on the results of evaluations, and compulsive attention to details of countless types.

In the remainder of this paper, I shall identify what appear to me to be some of the major needs, especially research needs, relating to productivity in the Space Station. At the beginning of this paper, it was noted that the term productivity is used in a variety of ways and often without a very precise connotation, and that except in certain highly-structured situations, how to quantify productivity unambiguously is not clear. If high productivity is to be an explicit objective of the space program, some consideration must be given to how it is to be measured or otherwise assessed in this context. Assessment will be desirable at various levels--that of the overall Space Station program, that of specific missions, that of specific crews during designated periods of time, and that of individuals performing specific

tasks.

For present purposes, it is assumed that enhancements--increases in the efficiency, accuracy, reliability--of the performance of humans or human-machine systems are very likely to improve productivity by nearly any reasonable definition and measurement technique. The recommendations that follow are predicated on that assumption. Research that is alluded to in some of these recommendations is already underway, in NASA laboratories and elsewhere. I am aware of some of these efforts, but there undoubtedly are many of which I am not. Inclusion in this list signifies only my opinion that the topic deserves attention; if it is getting it already, so much the better. While all of these recommendations are considered important to the Space Station program, they are not all uniquely applicable to it. Some of them are similar to recommendations that would apply to the design and development of any complex system that will have people interacting with computer-based tools in non-trivial ways (Nickerson et al., 1984).

- There is a need to organize the information that has been obtained from research on earth or from data gathered in previous space flights that is relevant to human performance in space. This information should be organized and indexed so as to make it highly accessible to scientists and engineers in the space program.
- It would be useful also to commission the compilation of an encyclopedia of ignorance about productivity, and performance more generally, in space. The primary objective should be to identify as many as possible of the important unanswered questions about performance in space. Questions should be prioritized with respect to urgency, and classified in terms of the kind of research that could lead to answers.
- What information will be required by specific members of the Space Station team at specific times needs to be determined. This includes determining what information should be presented spontaneously, and in such a way as to capture the intended receiver's attention, what information should be available explicitly on some display all (or most) of the time, and what information should be available but presented only on request.
- Possible and most-likely patterns of communication or information flow both within the Space Station and between the station and earth need to be understood better.
- More effective means of providing EVA access to data-base information pertinent to EVA tasks are needed.
- An inventory of tasks that people will be expected to perform in the Space Station should be compiled.

- Procedure descriptions should be evaluated for accuracy and clarity.
- Criteria need to be established regarding what aspects of the Space Station's operation should be automated. The rule that anything that can be automated (effectively, safely) should be automated is not necessarily a good rule. There may be some functions that can be done acceptably by either people or machines that should be done by people. Issues of morale, perception of control, and skill maintenance must be considered as well as that of technical feasibility.
- More research is needed on the question of how much "intelligence" to build into teleoperator or telerobot systems, and how much to rely on remote control by humans.
- The design of computer-based aids for trouble shooting, problem solving and decision making, and of the protocols for interacting with them deserves considerable attention.
- Efforts to advance the state-of-the-art of aiding human operators through the use of "intelligent", or "expert-system" software should be supported: potential applications in the Space Station program include fault detection, identification, and repair; planning and plan revising; and crisis management.
- The knowledge of astronauts and space professionals must be codified to provide the basis for the development of expert systems and knowledge-based aids.
- The phasing of expert system technology into operational situations as its evolution warrants will represent an ongoing challenge into the indefinite future.
- Possible problems involved in having crew members share responsibility of high-level cognitive tasks with "smart" software or expert systems need to be identified; policies should be established for deciding when to trust a system and when to override it.
- Design of the various interfaces through which Space Station personnel will interact with the numerous systems and subsystems on board is among the most critical problems to be solved, from a human factors point of view. There is a body of literature relating to the design of workstations and displays that should be consulted; however, much remains to be learned about how best to represent and present information in various Space Station contexts. This topic deserves a continuing effort of research focused on the identification of display formats, information coding dimensions, and input techniques that are especially well suited to the Space Station environment and the demands of specific tasks that are to be performed.

- Proposed or planned displays and work stations should be evaluated in terms of conventional human factors criteria: lighting, glare, flicker, contrast, character/symbol legibility/interpretability, functional-positional relationships, clutter, and so on.
- Display configurations and symbology must be designed and evaluated; this includes determination of content and format of specific-purpose displays. Display coding dimensions must be selected so as to minimize confusion arising from multiple functions of a given display space.
- A better understanding is needed of when to use menus and when to use command languages as input methods. The menus and languages to be used must be designed, evaluated and refined.
- There is a need to identify situations in which voice could be used to advantage as an input or output medium, given the probable state-of-the-art of voice recognition and production technology over the next decade or so.
- Further work is needed on the design of control and feedback interfaces for remote manipulators, teleoperators, and semi-autonomous systems. The problem is complicated when the distance between the devices and their operators is great enough to cause significant communication delays.
- The need for high resolution, stereo visual feedback from teleoperator systems should be studied and the feasibility of its use explored.
- More effective helmet-mounted displays for use in EVA should be a continuing research objective.
- The technology for tracking eye fixation and movement, and hand and finger position and movement could have applications in the Space Station, but need to be developed further.
- The technology needed to make a virtual-interface approach to teleoperator control a practical reality requires further exploration.
- Acquisition of anthropometric, range of motion, strength, and force and torque application data, with and without pressurized suits, should be continued.
- The ability to measure and monitor mental workload could be useful, especially for the establishment of crew responsibilities in the Station's day-to-day operation and in high-activity situations. But techniques that are to be used in operational contexts must be unintrusive, and this rules out the

applicability of many of those that have been used to study mental workload in the laboratory.

- A catalog of possible human errors (of both commission and omission) that could have non-trivial consequences in the Space Station should be developed; potential errors should be rated as to seriousness and probability of occurrence, and the results used to develop safeguards and error detection and recovery procedures.
- A detailed study of human errors that are actually made in the Space Station environment will be very useful, as it has been in other contexts (Meister, 1966; Swain, 1970, 1978).
- Methods of assuring the maintenance of critical skills that are typically used only in the event of a system malfunction or failure must be developed.
- Effects of prolonged living in restricted environments on work performance, social behavior and mental state deserve further study. More specifically, attempts should be made to identify aspects of such environments that are the major determinants of behavioral, cognitive or emotional effects.
- Special attention should be given to the types of interpersonal tensions and conflicts that are likely to arise in the Space Station environment and the development of effective techniques for relieving or resolving them.
- The question of how to occupy long periods of time during which the operational demands of the spacecraft are minimal deserves considerable attention. The maintenance of motivation, alertness and social stability during extended stretches of being, in essence, passengers on an automatically piloted craft represents a significant challenge.
- Presumably, productivity in space can be enhanced by factors that contribute to the maintenance of high levels of alertness, motivation and general physical and mental well being. We need to understand better how these variables depend on such factors as appropriate diet; regular physical exercise; the opportunity to engage in interesting and valued activities in free time; frequent communication with earth, not only regarding mission matters, but regarding those of personal interest; adequate variety in job responsibilities; adequate rest; and extensive use of error detection and failsafe procedures (especially those that can be automated).
- We need also to learn more about the relationships among certain performance or psychological variables (attention, vigilance, perception, memory, learning, thinking, and judgement) and

indicants of physiological state (EEG, evoked potential, contingent negative variation, heart rate, blood pressure, respiration, skin temperature, galvanic skin response). To the extent that variables in the latter category can be shown to be reliable indicants of the quality of specific types of human performance, consideration should be given to the development of unintrusive ways of monitoring them, at least at critical times, and using the results of the monitoring to enhance performance in various ways (Johnson et al., 1972). Although techniques exist for doing such monitoring, they tend to be sufficiently intrusive to interfere with the monitored individuals' performance of their primary tasks and to be less reliable than is desired. A continuing goal of research should be the development of less intrusive and more reliable techniques for monitoring cognitive state.

- The ability to monitor--and in particular to detect significant changes in--physiological and psychological states could prove to be especially important in long-term space missions. State changes that could be important to detect include both temporary fluctuations in alertness and long-range changes in general physical condition, motivation and mood.
- Biofeedback technology is still in its infancy, however the evidence is clear that people can learn, within limits, to control certain physiological functions that had been thought to be completely automatic. Further study of biofeedback techniques is warranted with a view to their possible application in the Space Station for purposes of controlling tension, facilitating good quality sleep, and otherwise tuning physiological states to enhance either performance or rest.
- Studies of the mental models that crew members or perspective crew members develop of the Space Station and its hardware and software components could help determine what kinds of models are acceptable for conveyance to future participants in Space Station missions.
- There is a need for better rapid prototyping capabilities especially for prototyping candidate interface designs.
- Procedures and policies must be established for acquiring data in space that can be used to relate productivity and performance to the numerous variables that are believed to affect them in significant ways.
- It is not likely that predictions about performance of humans in space can be very accurate very far into the future. A reasonable goal is the development of a predictive model, based on what is currently known from data collected on earth and from studies of performance in space to date, with the intent of

modifying that model continually as further relevant data are obtained, especially from experience in space. Conditions in space exploration will change and the durations of stays in space will increase, so the model will have to evolve to accommodate those changes. On the assumption that the changes that occur will be evolutionary and relatively continuous, one can hope for a model that is highly predictive of the situation that is current at any given time and reasonably predictive of the situation as it is anticipated to be in the near-term future.

REFERENCES

- Akerstedt, T., and Gillberg, M.
 1981 Sleep disturbance and shift work. In A. Reinberg, N. Vieux, and P. Andlauer, eds., Night and Shift Work: Biological and Social Aspects. Oxford: Pergamon Press.
- Baily, M. N.
 1986 What has happened to productivity growth. Science 234:443-451.
- Bechtold, S. E., Ganaro, R. E., and Summers, D. L.
 1984 Maximization of labor productivity through optimal rest-break schedules. Management Science 30:1442-1458.
- Berry, C. A.
 1969 Preliminary clinical report of the medical aspects of Apollo VII and VIII. Aerospace Medicine 40:245-254.
 1970 Summary of medical experience in the Apollo 7 through 11 manned spaceflights. Aerospace Medicine 41:500-519.
- Bowen, W.
 1986 The puny payoff from office computers. Fortune 5:20-24.
- Broadbent, D. E.
 1971 Decision and Stress. New York: Academic Press.
- Brownfield, C. A.
 1965 Isolation, Clinical and Experimental Approaches. New York: Random House.
- Cameron, C.
 1971 Fatigue problems in modern industry. Ergonomics 14:713-720.
 1973 A theory of fatigue. Ergonomics 16:633-648.
- Casali, J. G., and Wierwille, W. W.
 1983 A comparison of rating scale, secondary-task, physiological, and primary-task workload estimation techniques in a simulated flight task emphasizing communications load. Human Factors 25(6):623-641.

- Casali, J. G., and Wierwille, W. W.
 1983 A comparison of rating scale, secondary-task, physiological, and primary-task workload estimation techniques in a simulated flight task emphasizing communications load. Human Factors 25(6):623-641.
- Chaudhari, P.
 1986 Electronic and ceramic materials. Scientific American 255(4):136-144.
- Chiles, W. D., and Alluisi, E. A.
 1979 On the specification of operator or occupational workload with performance-measurement methods. Human Factors 21:515-528.
- Clearwater, Y.A.
 1985 A human place in outer space. Psychology Today 7:34-43.
- Clearwater, Y. A. and Kaspar, P.
 1986, Humanizing space. Professional Careers. 10/11:5-8.
- Cobb, S., and Rose, R. M.
 1973 Hypertension, peptic ulcers and diabetes in air traffic controllers. Journal of the American Medical Association 224:489-492.
- Crump, J. H.
 1979 A review of stress in air traffic control: Its measurement and effects. Aviation, Space and Environmental Medicine 50:243-248.
- Danaher, J. W.
 1980 Human error in ATC system operations. Human Factors 22(5):535-545.
- Dray, S. M., Ogden, W. G., and Vestewig, R. E.
 1981 Measuring performance with a menu-selection human computer interface. Pp. 746-748 in Proceedings of the Human Factors Society 25th Annual Meeting.
- Drucker, P. F.
 1986 The Frontiers of Management. New York: Truman Talley Books
- Eason, R. G. and Harter, M. R.
 1972 Sensory, perceptual, and motor factors. Human Factors in Long-Duration Spaceflight. Washington, DC: National Academy of Sciences.

- Eggemeier, F. T.
1980 Some current issues in workload measurement. Pp. 669-673 in Proceedings of the Human Factors Society 24th Annual Meeting.
- Ehrenreich, S. L.
1981 Query languages: Design recommendations derived from the human factors literature. Human Factors 23(6):709-725.
- Fidell, S. A.
1977 Effects of stress on performance. The C³ System User Vol I: A Review of Research on Human Performance as it Relates to the Design and Operation of Command, Control, and Communication Systems. BBN Report 3459.
- Finkelman, J. M., and Kirschner, C.
1981 An information-processing interpretation of air traffic control stress. Human Factors 22(5):561-5
- Fitts, P. M. (Ed.)
1951 Human Engineering for an Effective Air-Navigation and Traffic-Control System. Columbus, OH: Ohio State University Research Foundation.
- Fraser, T. M., Busby, D. E., Adey, W. R., and Galambos, R.
1972 Physiological and Medical Factors. Human Factors in Long Duration Spaceflight. Washington, DC.: National Academy of Sciences.
- Ganaro, R. E., and Bechtold, S. E.
1985 A study of the reduction of fatigue impact on productivity through optimal rest break scheduling. Human Factors 27:459-466.
- Garriott, O. K.
1974 Skylab report: Man's role in space research. Science 186:219-225.
- Garriott, O. K., Parker, R. A. R., Lichtenberg, B. K., and Merbold, U.
1984 Payload crew members' view of Spacelab operations. Science 225:165-167.
- Giuliano, V. E.
1982 The Mechanization of Office Work. Scientific American 247:(3)148-164.
- Gunderson, E. K. E. (Ed.)
1974 Human Adaptability to Antarctic Conditions. Washington, DC: American Geophysical Union.

- Gunderson, E. K. E.
1963 Emotional symptoms in extremely isolated groups. Archives of General Psychiatry 9:74-80.
- Gunderson, E. D. E. and Nelson, P. D.
1963 Adaptation of small groups to extreme environments. Aerospace Medicine 34:1111-1115.
- Gunn, T. G.
1982 The mechanization of design and manufacturing. Scientific American 247:(3)114-130.
- Hailey, A. H.
1968 Airport. New York: Doubleday and Company.
- Haythorn, W. W., and McGrath, J. J., and Hollander, E. P., and Latane, B.
1972 Group processes and interpersonal interaction. Human Factors in Long-Duration Spaceflight. Washington, DC.: National Academy of Sciences.
- Hicks, T., and Wierwille, W.
1979 Comparison of five mental workload assessment procedures in a moving-base simulator. Human Factors 21:129-144.
- Holmes, D. S., and Roth, D.
1985 Association of aerobic fitness with pulse rate and subjective responses to psychological stress. Psychophysiology 22(5):525-529.
- Howard, R. D., Miller, R. H., Minsky, M. L., and et al.
1982 Space Applications of Automation, Robotics and Machine Intelligence Systems (ARAMIS). Phase I, Structural Assembly Expansion. Volumes 1-5, Space Systems Laboratory and Artificial Intelligence Laboratory, SSL Reports #25-82 - #29-82, Massachusetts Institute of Technology, Cambridge, MA.
- Isreal, J. B., Wickens, C. D., and Chesney, G. J.
1980 The event-related brain potential as an index of display-monitoring workload. Human Factors 22(2):211-227.
- Johannsen, G.
1979 Workload and workload measurement. Pp. 3-11 in N. Moray, ed., Mental Workload: Its Theory and Measurement. New York: Plenum.

- Johannsen, G., Moray, N., Pew, R., Rasmussen, J., Sanders, A., and Wickens, C.
 1979 Final report of experimental psychology group. Pp. 101-114 in N. Moray, ed., Mental Workload: Its Theory and Measurement. New York: Plenum.
- Johnson, L. C., Williams, H. L., and Stern, J. A.
 1972 Motivation, cognition, and sleep-work factors; central-and autonomic-nervous-system indices. Human Factors in Long-Duration Spaceflight. Washington, DC.: National Academy of Sciences.
- Jones, E. R., Chiles, W. D., and Voas, R. B.
 1972 Human factors and operational requirements. Human Factors in Long-Duration Spaceflight. Washington, DC.: National Academy of Sciences.
- JSC
 1979 Space Operations Center, A Concept Analysis. JSC-16277. NASA/Johnson Space Center.
- Kalsbeek, J. W. H.
 1968 Measurement of mental work and of acceptable work: Possible applications in industry. International Journal of Production Research 7:33-45.
- 1971 Sinus arrhythmia and the dual task method in measuring mental load. In W. T. Singleton, J. G. Fox, and D. Whitfield, eds., Measurement of Man at Work. London: Taylor and Francis.
- Kennedy, R. S.
 1972 A Bibliography of the Role of the Vestibular Apparatus Under Water and Pressure: Content-Oriented and Annotated. M4306.03.1J00F'K9, Naval Medical Research Institute Report No. 1.
- Kessel, C. J., and Wickens, C. D.
 1983 The transfer of failure-detection skills between monitoring and controlling dynamic systems. Human Factors 24(1):49-60.
- King, K. R., and Rouen, M. N.
 1982 Extravehicular Mobility Unit. NASA-JSC-18201. Satellite Services Workshop 1:44-72.
- Lazarus, R. A., and Monat, A.
 1977 Stress and coping--Some current issues and controversies. In A. Monat and R. S. Lazarus, eds., Stress and coping. New York: Columbia University Press.

- Lee, E., and MacGregor, J.
1985 Minimizing user search time in menu retrieval systems. Human Factors 27(2):157-162.
- Leventhal, H., and Lindsley, D. B.
1972 Subjective states. Human Factors in Long-Duration Spaceflight. Washington, DC.: National Academy of Sciences.
- Lindsley, D. B., and Jones, E. R.
1972 Physical factors. Human Factors in Long-Duration Spaceflight. Washington, DC.: National Academy of Sciences.
- Loftus, J. P., Jr., Bond, R., and Patton, R. M.
1982 Astronaut activity. Workshop Proceedings: Space Human Factors, 1, Leesburg, VA, II-28-II-64.
- McGrath, J. J.
1965 Performance sharing in an audio-visual vigilance task. Human Factors 7:141-153.
- Meister, D.
1966 Human factors in reliability. In W. G. Ireson, ed., Reliability Handbook. New York: McGraw-Hill.
- Miller, D. P.
1981 Depth/breadth tradeoff in hierarchical computer menus. Proceedings of the Human Factors Society 25th Annual Meeting. Santa Monica, CA: Human Factors Society.
- Montemarlo, M. D., and Cron, A. C., eds.
1982 Space Human Factors Workshop, Volumes 1 and 2. NASA/Office of Aeronautics and Space Technology OAST, Code RTH-6, Washington, DC.
- Moray, N.
1979 Preface. Pp. V-VIII in N. Moray, ed., Mental Workload: Its Theory and Measurement. New York: Plenum.
- Muckler, F. A.
1982 Evaluating productivity. In M. D. Dunnett and E. A. Fleishman, eds., Human Performance and Productivity, Vol. 1, Human Capability Assessment. Hillsdale N.J.: L. Erlbaum Associates.
- NASA-ASEE
1985 Autonomy and human element in space. Proceedings of the 1983 NASA/ASEE Summer Faculty Workshop at Stanford University.
- National Commission on Space
1986 Pioneering the Space Frontier. New York: Bantam Books.

- Nickerson, R. S.
1986 Using Computers: Human Factors in Information Systems. MIT Press (Bradford Books).
- Nickerson, R. S. and others
1984 Research Needs on the Interaction Between Information Systems and their Users: Report of a Workshop. Washington, DC: National Academy Press.
- Norman, G. R., McFarlane, A. G., and Streiner, D. I.
1985 Patterns of illness among individuals reporting high and low stress. Canadian Journal of Psychiatry 30(6):400-405.
- Ogden, G. D., Levine, J. M., and Eisner, E. J.
1979 Measurement of workload of secondary tasks. Human Factors 21:529-548.
- Parks, D.
1979 Current workload methods and emerging challenges. In N. Moray, ed., Mental Workload: Theory and Measurement. New York: Plenum.
- Price, H.E.
1985 The allocation of functions in systems. Human Factors 27:33-45.
- Price, H. E., Maisano, R. E., and Van Cott, H. P.
1982 The Allocation of functions in Man-Machine Systems: A Perspective and Literature Review. NUREG-CR-2623, Oak Ridge, TN: Oak Ridge National Laboratories.
- Price, H. E., and Pulliam, R.
1983 Control room function allocation - A need for man-computer symbiosis. Proceedings of the 1982 IEEE Computer Forum. Denver Co: Institute of Electrical and Electronic Engineers.
- Rehmann, J. T., Stein, E. S., and Rosenberg, B. L.
1983 Subjective pilot workload assessment. Human Factors 25(3):297-307.
- Rolfe, J. M.
1971 The secondary task as a measure of mental load. In W. T. Singleton, J. G. Fox, and D. Whitfield, eds., Measurement of Man at Work. London: Taylor and Francis.
- Schultz, D. P.
1965 Sensory Restriction: Effects on Behavior. New York: Academic Press.

- Schwab, E. C., Nusbaum, H. C., and Pisoni, D. B.
1985 Some effects of training on the perception of synthetic speech. Human Factors 27(4):395-408.
- Sells, S. B., and Gunderson, E. K.
1972 A social system approach to long-duration missions. Human Factors in Long-Duration Spaceflight. Washington, DC.: National Academy of Sciences.
- Seppala, P., and Salvendy, G.
1985 Impact of depth of menu hierarchy on performance effectiveness in a supervisory task: Computerized flexible manufacturing system. Human Factors 27(6):713-722.
- Sharit, J., and Salvendy, G.
1982 Occupational stress: Review and reappraisal. Human Factors 24(2):129-162.
- Sheridan, T. B., and Simpson, R. W.
1979 Toward the Definition and Measurement of the Mental Workload of Transport Pilots. Flight Transportation Man-Machine Laboratory, Technical Report No. DOT-OS-70055. Cambridge, MA: MIT.
- Simonson, E., and Weiser, F. C., eds.
1976 Psychological Aspects and Physiological Correlates of Work and Fatigue. Springfield, IL: Charles C. Thomas.
- Singleton, W. T., Fox, J. G., and Whitfield, D., eds.
1971 Measurement of Man at Work. London: Taylor and Francis.
- Small, D. W., and Weldon, L. J.
1983 An experimental comparison of natural and structured query languages. Human Factors 25(3):253-263.
- Smith, S.
1969 Studies of small groups in confinement. In J. P. Zubek, ed., Sensory Deprivation: Fifteen Years of Research. New York: Appleton-Century-Crofts.
- Smith, R. P.
1981 Boredom: A review. Human Factors 23(3):329-340.
- Sutcliffe, A.
1976 People and Productivity. 3rd. edition, New York: McGraw-Hill.
- Swain, A. D.
1970 Development of a Human Error Rate Data Bank. Sandia Corporation Report SC-R-70-4286, Albuquerque, NM: Sandia Corporation.

- 1978 Estimating Human Error Rates and their Effects on System Reliability. Sandia Corporation Report SAND77-1240, Albuquerque, NM: Sandia Corporation.
- Taylor, F. W.
1947 Scientific Management. New York: Harper Brothers.
- Thiel, E. D., and Kurtzman, C. R.
1978 Space Applications of Automation, Robotics and Machine Intelligence Systems (ARAMIS). Phase II, Telepresence. [1-3], SSL Reports #30-83 - #32-83, Space Systems Laboratory and Artificial Intelligence Laboratory. Cambridge, MA: MIT
- Tilley, A. J., Wilkinson, R. T., and Drud, M.
1981 Night and day shifts compared in terms of the quality and quantity of sleep recorded in the home and performance measured at work: A pilot study. In A. Reinberg, N. Vieux, and P. Andlauer, eds., Night and Shift Work: Biological and Social Aspects. Oxford: Pergamon Press.
- Tilley, A. J., Wilkinson, R. T., Warren, P. S. G., Watson, E., and Drud, M.
1982 The sleep and performance of shift workers. Human Factors 24(6):629-641.
- VanAllen, J. A.
1986a Myths and realities of space flight. Science 4754:1075-1076.
1986b Space science, space technology and space station. Scientific American 254(1):32-34.
- VanCott, H. P., and Kincaid, R. G., eds.
1972 Human Engineering Guide to Equipment Design. Revised edition. Sponsored by the Joint Army-Navy-Air Force Steering Committee, U.S. Department of Defense. Washington, DC.: U.S. Government Printing Office.
- Weiner, E. L.
1975 On simultaneous monitoring and tracking. Journal of Applied Psychology 60:100-105.
- Weiner, E. L., Curry, R. E., and Faustiar, M. L.
1984 Vigilance and task load: In search of the inverted U. Human Factors 26:215-222.
- Welford, A. T., ed.
1974 Man Under Stress. New York: John Wiley and Sons.
- Welford, A. T.
1973 Stress and performance. Ergonomics 16:567-580.

- Weybrew, B. B.
 1961 Human factors and work environment. II. The impact of isolation upon personnel. Journal of Occupational Medicine 3:290-294.
- 1963 Psychological problems of prolonged marine submergence. Pp. 87-125 in N. Burns, R. Chambers, and E. Hendler, eds., Unusual Environments and Human Behavior. New York: MacMillan.
- Wickens, C. D.
 1979 Measures of workload, stress, and secondary task. In N. Moray, ed., Mental Workload: Its Theory and Measurement. London: Plenum.
- Wickens, C. D., and Kessel, C.
 1979 The effects of participatory mode and task workload on the detection of dynamic system failures. IEEE Transactions on Systems Man and Cybernetics. SMC-13:24-31.
- 1980 Processing resource demands of failure detection in dynamic systems. Journal of Experimental Psychology: Human Perception and Performance 6(3):564-577.
- Wierwille, W. W., and Connor, S. A.
 1983 Evaluation of 20 workload measures using a psychomotor task in a moving-base aircraft simulator. Human Factors 25(1):1-16.
- Wierwille, W. W., Rahimi, M., and Casali, J. G.
 1985 Evaluation of sixteen measures of mental workload using a simulated flight task emphasizing mediational activity. Human Factors 27:489-502.
- Williges, R. C., and Wierwille, W. W.
 1979 Behavioral measures of aircrew mental workload. Human Factors 21:549-574.
- Wise, J. A.
 1985 The Quantitative Modeling of Human Spatial Habitability. NASA Grant NAG 2-346. Moffett Field, CA: NASA-Ames Research Center.
- 1986 The space station: Human factors and habitability. Human Factors Society Bulletin 29(5):1-3.
- Woodson, W. E.
 1981 Human Factors Design Handbook. New York: McGraw-Hill.

- Young, L. R.
1969 On adaptive manual control. IEEE Transactions of Man-Machine Systems. MMS-10:292-331.
- Zubek, J. P.
1973 Behavioral and physiological effects of prolonged sensory and perceptual deprivation--a review. Pp. 8-83 in Man in Isolation and Confinement. Chicago: Aldine Publishing Co.
- Zubek, J. P., and Bayer, S. M., and Milstein, S., and Shephard, J. M.
1969 Behavioral and physiological changes during prolonged immobilization plus perceptual deprivation. Journal of Abnormal Psychology 74:230-236.

DISCUSSION:
COMMENTS ON SYSTEM PRODUCTIVITY: PEOPLE AND MACHINES

Robert C. Williges

Nickerson's paper provides an excellent review of human factors implications when considering productivity in the space station. In an attempt to amplify some of his points, I will restrict my comments to the ramifications of productivity as espoused in modern industrial engineering. As a point of departure, I will use the recent text by Sink (1985) on productivity management to discuss topics related to defining, measuring, and improving productivity.

WHAT IS PRODUCTIVITY?

In the most general form, productivity in industrial engineering is defined as a simple ratio of some quantity of output divided by some quantity of input. From a systems point-of-view, input quantities (e.g., labor, capital, energy, materials, etc.) go through some transformation (e.g., manufacturing, information processing, etc.) to yield an output (e.g., goods, services, waste, etc.) as shown in Figure 1. By comparing the output quantity to the input quantity, one can assess system productivity as a simple ratio.

Two implications are readily apparent from this operational definition of productivity. First, productivity is a metric that represents more than just output performance. It is a measure of output performance relative to input resources. Consequently, productivity is but one component of performance and should not be equated with overall performance. Other related system performance components might include efficiency, effectiveness, innovation, quality, profitability, etc. From a human factors point of view, productivity has the potential to serve as one metric for evaluating humans as components in complex space systems.

A second implication of the operational definition of productivity is that the ratio metric is based on some defined unit of analysis. Just as the Bureau of Labor statistics measure of overall national productivity (i.e., Gross National Product, GNP, divided by labor input) is of limited value, an overall measure of space station productivity is limited. Care must be taken to choose a meaningful level of analysis in assessing productivity in space systems. From a human productivity point-of-view, it may be difficult to distinguish




FIGURE 1 Basic configuration of the productivity metric.

productivity from human performance in cognitive tasks until better measures of input resources, cognitive processes, and output measures are available.

Productivity does, however, seem to be a viable metric to evaluate larger units of analysis of space-related missions in which the astronaut is considered one component of the unit of analysis. These larger units of analysis should be considered in terms of the human/machine interface level and above. For example, the human component could be considered in assessing the productivity of a space station or in assessing productivity of working environments such as intravehicular activities (IVA) at workstations, extravehicular activities (EVA) outside the space station, and combined IVA and EVA operations such as telerobotic activities (Gillan et al., 1986). In each case, the ratio metric of productivity includes human components along with hardware and software components, and these productivity assessments can be used to evaluate the relative contributions of various components.

HOW IS PRODUCTIVITY MEASURED?

Traditionally, both the time domain and the number of component factors measured are considered in calculating the productivity ratio. In the time domain, both static and dynamic measures of productivity are used. Static measures are used to calculate the productivity ratio for a particular point in time; whereas, dynamic measures are used to evaluate changes in productivity across a designated time unit. Both measures appear to be useful in evaluating the productivity of the human component in space. Static ratios can be used to assess the relative effect of the astronaut in terms of training investment and performance on a particular space mission. Dynamic productivity indices can be used to evaluate changes in team size, allocation of

tasks/functions, and return on investments in automation for space missions.

Both static and dynamic measures of productivity can vary in their level of complexity depending upon the number of components measured. Sink (1985), for example, suggests three levels of complexity determined by the number of factors used to construct the productivity ratio. He refers to partial-factor, multifactor, and total-factor measures. Partial factor measures include only one component class (e.g., mission specialist); multifactor measures include several component classes (e.g., mission specialist and computer interface); and total-factor measures include all component classes (e.g., mission specialist, computer-interface, test equipment, documentations, etc.) included in any particular productivity unit of analysis. Obviously, the simple productivity ratio quickly explodes into a complex, multivariate measurement problem once the unit of analysis and number of factors of measurement increases. Research is needed to build and evaluate complex productivity measurement systems for assessing human components of productivity in space missions.

HOW CAN PRODUCTIVITY BE IMPROVED?

In that productivity is a ratio metric, increased productivity must be considered in terms of both input and output quantities and not merely in terms of improving output. Consequently, productivity improvement can be achieved in five ways, as shown in Table 1, depending upon the relationship of the input and output conditions. Although these conditions are somewhat restricted when considering the human component, all appear to be possible if the unit of productivity analysis includes human, hardware, and software components related to space missions. Mostly, one considers human productivity improvement in terms of human performance improvements as Nickerson suggests in his paper. But the implication of the conditions listed in Table 1 suggests that these potential human performance improvements (in output) must be evaluated relative to the input changes (e.g., increased training, cost of automation, etc.) in order to evaluate the real impact on productivity.

TABLE 1 Conditions for Improving Productivity (after Sink, 1985)

Increasing Output

1. Output increases; input decreases
2. Output increases; input remains constant
3. Output increases; input increases at a lower rate

Constant Output

4. Output constant; input decreases

Decreasing Output

5. Output decreases; input decreases at a more rapid rate
-

RESEARCH ISSUES

Productivity from an industrial engineering point-of-view provides an important metric for assessing human performance as a systems component in space missions. Human productivity *per se* needs to be considered in a systems context, and any evaluation of productivity must assess both input and output quantities in order to establish a ratio metric. Two general areas of productivity research in space-related missions appear to warrant increased attention.

Measuring Productivity

Several measurement issues must be addressed before human productivity assessments of space missions can be made. The appropriate units of analysis for productivity measurement must be specified. Criteria for partial-factor, multifactor, and total-factor measures need to be established and verified. Automated human performance assessment schemas (Williges, 1977) need to be constructed which could then be used for embedded performance measurement, evolutionary operation, empirical modeling, multivariate criteria, and realistic data bases from which theoretical extrapolations could be made to the design of a variety of future space-related tasks. Improved productivity measurement models with sophisticated human productivity parameters need to be developed and validated. Many of these measurement issues can be addressed by current multivariate measurement procedures, but each of them will require validation during actual space missions.

Improving Productivity

Most of the research issues presented in the Nickerson paper dealing with performance enhancements can relate to improving human productivity if the antecedent input quantities are evaluated in order to establish appropriate productivity indices. The unit of analysis at the human-machine interface level or above seems to provide the best opportunities for improved productivity given the characteristics of the productivity metric. Research issues raised by Nickerson dealing with workstation design, human input modes, decision aids, and automation are particularly relevant. In fact, many of the remaining topics to be discussed during this symposium are candidate issues that could be evaluated in terms of productivity improvement metrics.

CONCLUSION

Productivity is an often used and abused term. By accepting the rather straightforward operational definition of productivity as a ratio of output quantity divided by input quantity, I believe productivity holds promise as an important component metric of space station performance which include human, hardware, and software parameters. Before such a

metric is useful, several productivity measurement and productivity enhancement research issues must be addressed.

REFERENCES

- Gillan, D. J., Burns, M. J., Nicodemus, C. L., and Smith, R. L.
1986 The space station: human factors and productivity. Human Factors Society Bulletin 29(11):1-3.
- Sink, D. S.
1985 Productivity Management: Planning, Measurement and Evaluation, Control and Improvement. New York: Wiley.
- Williges, R. C.
1977 Automation of performance measurement. Pp. 153-168 in Symposium Proceedings of Productivity Enhancement: Personnel Performance Assessment in Navy Systems. San Diego: Navy Personnel Research and Development Center.

SYNOPSIS OF GENERAL AUDIENCE DISCUSSION

Areas of Concern

Two aspects of the space station provide different concerns for evaluating human productivity. Housekeeping activities may prove to be an important candidate for productivity improvement in terms of reducing the amount of time required to perform these functions. Another major component of the space station is the conduct of scientific activities. Improving productivity related to space research activities appear to be more difficult to measure. In addition to 'on orbit' space station concerns, the integration of ground-control and on-board activities is a prime candidate for productivity improvement studies.

Productivity Metrics

Several of the components related to human productivity in space will be difficult to quantify. Consequently, the accuracy and viability of these measures may be somewhat questionable at certain units of analysis. This underscores the appropriate choice of the unit of analysis. In addition, qualitative measures may need to be substituted for quantitative measures in certain instances.

Lessons Learned

Analysis of other isolated, long duration missions such as early warning systems and sea lab may be useful in making assumptions and generating initial models of key parameters related to productivity for space-related missions. For example, isolations may be a catalyst to trigger stress factors affecting productivity. Caution, needs to be exercised in extrapolating from these analogs, because clear differences exist. Nonetheless, evaluation of these related systems may be useful in isolating a common thread of critical variables affecting human productivity.

SESSION II
EXPERT SYSTEMS AND THEIR USE

AI SYSTEMS IN THE SPACE STATION

Thomas M. Mitchell

INTRODUCTION

Among the technologies that will help shape life in the space station, Artificial Intelligence (AI) seems certain to play a major role. The striking complexity of the station, its life support systems, and the manufacturing and scientific apparatus that it will house require that a good share of its supervision, maintenance, and control be done by computer. At the same time, the need for intelligent communication and shared responsibility between such computer programs and space station residents poses a serious challenge to present interfaces between man and machine. Hence, the potential and need for contributions from AI to the space station effort is great.

The purpose of this paper is to suggest areas in which support for new AI research might be expected to produce a significant impact on future space station technology. Given the breadth of this task, the approach here will be to sample a few such areas and to rely on the other symposium participants and other sources (e.g., Technical Report NASA-ASEE, 1983; Technical Report NASA, 1985) to fill in the picture. More specifically, we will address here (1) the use of knowledge-based systems for monitoring and controlling the space station, and (2) issues related to sharing and transferring responsibility between computers and space station residents.

Before focussing on the specifics of these two problem areas, it is useful to understand their significance to the development of the space station (and to other advanced projects such as development of a lunar base and interplanetary probes).

In his keynote address at this symposium, Allen Newell provides an analysis of the general characteristics and constraints that define the space station effort. Those of particular relevance to this paper include the following:

- The station is an extraordinarily complex system with an extremely high premium to be placed on reliability, redundancy, and failsafe operation. In past space efforts, a large share of astronaut training has gone into acquiring the knowledge needed to supervise, control, and troubleshoot various spacecraft

subsystems. The increased complexity of the space station argues for computer-based assistance in the supervision of many station subsystems, and it is no surprise that the history of the space program is a history of increasing automation and computer supervision. Furthermore, the high premium on failsafe operation places strong demands on the flexibility and adaptability of such computer-based supervisors. Such systems must be flexible enough to recognize and adapt to unanticipated events, and to communicate such unanticipated events clearly to the humans who help choose the response to these events. The flexibility demanded here goes well beyond that associated with present-day computer based supervisory systems.

- The space station is intended to be a highly evolutionary system, which will be continually reconfigured and upgraded over the course of its lifetime in space. The highly evolutionary nature of the station will make the task of crew training even more difficult than if the station were a static system. The problem of updating operating and troubleshooting procedures will be greatly exacerbated. In general, there will be greater demands on maintaining and updating the external documentation of the space station subsystems, and on prompt, thorough updating of procedures for monitoring, controlling, and troubleshooting the evolving space station. Computer-based methods for automatically updating such procedures, given updates to the description of the space station, would greatly enhance the ability to manage the evolving station.
- The crew of the space station will possess differing levels of expertise regarding different space station subsystems, and will live in the station long enough that their expertise will change over the course of their stay aboard the station. These differences in level of sophistication among various crew members (and between the same crew member at differing times) pose significant problems and opportunities for the computer systems with which they will interact. For naive users, computer systems that recommend given actions will have to provide a fairly detailed explanation of the reasoning behind the recommendation. For more expert users, less explanation may be needed. For advanced users, there will be an opportunity for the computer system to acquire new problem-solving tactics from the users. Furthermore, as a particular user becomes familiar with the competence and limitations of a particular computer-based supervisor, his willingness to allow the system to make various decisions without human approval may well change. The ability to interface effectively with a range of users, acting as a kind of tutor for some and acquiring new expertise from others, would allow the computer to act as the "corporate memory" for the particular aspect of the space station that is its domain and for which it will house a continually evolving set of expertise.

MONITORING, DIAGNOSING, AND CONTROLLING THE SPACE STATION

Given the above characteristics of the space station effort, it is clear that the use of computer-based assistants for supervising various space station subsystems could have a major impact on the overall reliability and cost of space station operations. In order to develop such computer-based supervisors, basic research is needed in a number of areas such as representing and reasoning about complex designed artifacts, inferring the behavior of such systems from schematics showing their structure, and automatic refinement of supervisory procedures based on empirical observation as well as the known system schematics.

Since the space station will itself be a large, well-documented artifact, it is reasonable to expect a significant number of opportunities for applying computers to the task of supervising, controlling and diagnosing the space station. For example, one might well expect that a computer could monitor various space station subsystems such as the parts of the navigation system, to detect behavior outside their expected operating ranges, take remedial actions to contain the effects of observed errors, diagnose the likely causes of the observed symptoms, and reconfigure the system to eliminate the error. Of course, limited applications of computers to this kind of problem are fairly common in current-day space systems. But present methods for automated monitoring, diagnosis and control are far from the levels of generality, robustness, maintainability, and competence that one would desire. AI offers a new approach to the problem of automated supervision. With appropriate research support, NASA might expect to significantly accelerate the development of AI methods for dealing with this class of problems, and thereby provide important new technology to support the space station.

A number of recent AI systems have addressed problems of monitoring, diagnosing, or controlling designed artifacts such as computer systems (Ennis et al., 1986), electro-mechanical systems (Pazzani, 1986), chemical processes (Scarl et al., 1985), and digital circuits (Davis, 1984; Genesereth, 1981). From this work, an initial set of techniques has emerged for building computer programs that embody a model (often in qualitative terms) of the behavior of the system under study, and which use this model to reason about the diagnosis, control, or reconfiguration of the system. While much remains to be understood, the initial approaches have shown clearly the potential for supervisory computer systems that combine judgemental heuristics with reasoning from a concrete model of the systems under study.

An Example

As an example of an AI system that deals with monitoring and troubleshooting a designed artifact, consider Davis' circuit troubleshooting system (Davis, 1984). This system troubleshoots digital circuits, given a schematic of the misbehaving circuit together with detected discrepancies between predicted and observed signal

values. Its organization is typical of several troubleshooting systems that have been developed for electronic, mechanical, and other types of systems.

The basic idea behind this troubleshooting system is that it uses the schematic of the system, together with its knowledge of the expected behaviors of system components, in order to reason backward from observed incorrect output signals to those upstream circuit components that could have produced the observed error. This process is illustrated in Figure 1, taken from Davis (1984).

In this figure, if the circuit inputs are given as shown, the system will infer the expected outputs as shown in round parentheses, based on its knowledge of the behaviors of multipliers and adders. If the two observed outputs are as shown in square parentheses, then a discrepancy is found between the expected and observed values for signal F. The system will then enumerate candidate fault hypotheses by considering that the error may be due to a failure in Add-1, or to incorrect values for one of its inputs (either X or Y). Each of these last two hypotheses might be explained further in terms of possible failures of the components or signals on which it, in turn, depends. Thus, candidate fault hypotheses are enumerated by examining the structure of the circuit as well as the known behaviors of its components.

In addition to enumerating fault hypotheses in this fashion, the system can also prune these hypotheses by determining other anticipated consequences of presumed faults. For example, the hypothesis that the error in signal F is caused by an error in signal Y, carries with it certain implications about the value of signal G. The value of 10 for signal F can be explained by a value of 4 for signal Y, but this would in turn lead to an expected value of 10 for signal G (which is observed to hold the value 12). Hence, this hypothesis may be pruned, as long as one assumes that the circuit contains only a single fault.

The above example illustrates how a computer system can reason about possible causes of observed faults, by using knowledge of the schematic of the faulty system as well as a library describing the expected behaviors of its components. There are many subtleties that have been

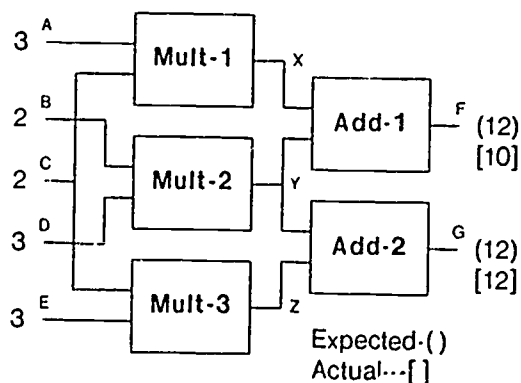


FIGURE 1 Troubleshooting example. Source: Davis (1984).

glossed over in this example, such as reasoning about the possibility of multiple system faults, interactions between faults, intermittent errors, utilizing statistical knowledge of likely faults and the resulting faulty behavior, scaling this approach to more complex systems, and the like. Basic research is still needed to develop more realistic diagnostic systems of this sort, and many of these issues are under study at this time. In addition, a good deal of research has been devoted to developing similar troubleshooting systems for artifacts other than digital circuits (e.g., mechanical electromechanical, and chemical processes). The topic of reasoning about the expected behavior of designed artifacts of many types is an active research area within AI (see, for example, the recent special volume of *Artificial Intelligence on qualitative reasoning about physical systems* (North-Holland, 1984).)

"Hands-On" Supervisory Systems

The above example is meant to suggest how a program can utilize an internal model of the system it is monitoring in order to localize the cause of anomalous behavior. Since the space station will be heavily instrumented with sensors and with computer-controlled effectors, the real opportunity here lies in developing a technology for "hands-on" AI supervisory systems: systems that have the means to directly observe and control the behavior of systems that they monitor, and that possess an explicit model of the system under supervision to guide their reasoning about monitoring, controlling, and troubleshooting this system. Figure 2 illustrates the general organization of such a hands-on supervisory system.¹

One instantiation of the scenario characterized in the figure could be an electronically self-sensing, self-monitoring space station. Here the system under supervision is the space station, sensors may observe the temperatures, pressures, and electrical behavior of various subsystems of the space station, and effectors may correspond to electrically controlled devices such as signal generators, heaters, compressors, and alarm systems. The goal of such an intelligent, self-monitoring space station would be to observe its behavior through its sensors, comparing these observations to the behavior anticipated by its internal model, and utilizing its effectors to maintain stable operation, reconfigure subsystems, and control the trajectory of states of the system. A number of observations are apparent about such a system: To a limited degree it is already possible to build such partially self-monitoring systems. The theoretical possibilities for computer monitoring and control in such systems far exceed the capabilities of our present techniques. The effectiveness of such a system will depend on continuing fundamental research in AI, especially in areas such as qualitative reasoning, diagnosis, control, and learning. To allow for such a future, the initial design of the space station must allow for flexible introduction of new sensors and effectors in all subsystems of the space station, and over the entire life of the station.

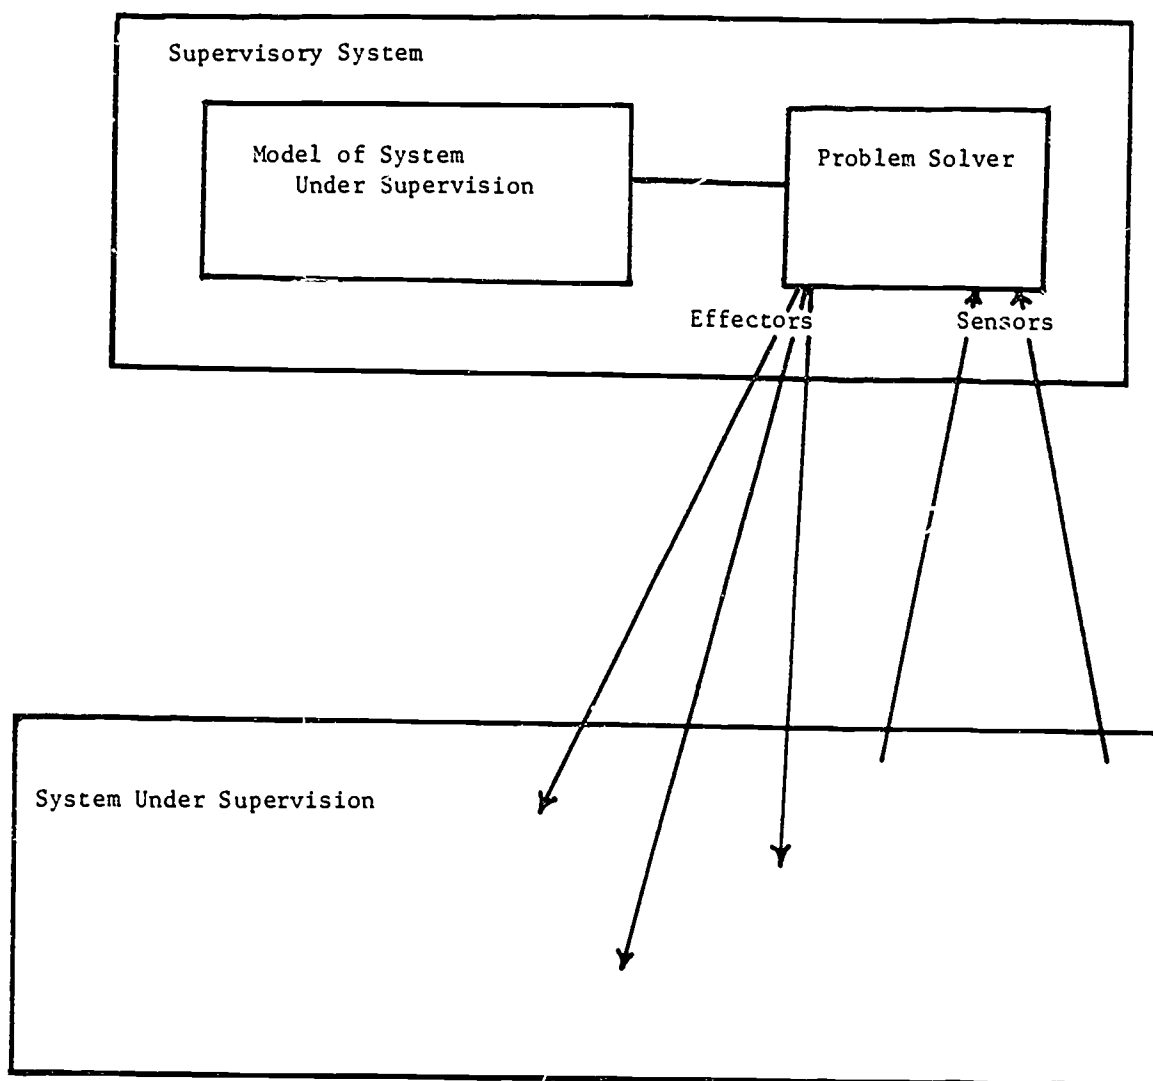


FIGURE 2 Hands on supervisory system.

A very different instantiation of the scenario of Figure 2 is obtained by introducing mobility in the sensors and effectors of the computer monitor. In this case, the supervisor could take the form of a collection of mobile platforms whose sensors include cameras, range finders, touch sensors, and oscilloscope probes, and whose effectors include wheels, rocket engines, manipulators, signal generators, and arc welders. Such a system might be expected to monitor the physical plant of the space station, checking for wear, and repairing the station as necessary, both interior and exterior. Several observations follow from considering this scenario: The leverage gained by adding

mobility to sensors and effectors is large--especially in situations such as troubleshooting where the system parameters in question might not be directly observable or controllable by statically positioned sensors and effectors. A number of difficult issues arise in representing and reasoning about three dimensional space, navigation, and the mechanics of physical systems. Given previous experience with robotics, it is clear that the difficulty of the technical problems can be considerably eased by designing a well-engineered work environment (e.g., by including easy grasping points on objects that are to be manipulated) in the space station.

In fact, we would like our supervisor to possess a combination of mobile and stationary sensors and effectors, including the union of those in the above scenarios. Thus, these two scenarios illustrate different aspects of the class of hands-on supervisor problems summarized in Figure 2. The two scenarios suggest a number of common technical problems, including problems of integrating human judgement with computer judgement, planning a sequence of control operations based on only an incomplete model of the system under supervision, and utilizing sensory input to refine the model of the system under supervision. At the same time, each scenario carries its own technical problems which overlay those generic issues. For example, a mobile supervisor for monitoring and repairing the exterior surface of the space station must face issues such as representing and reasoning about three dimensional space and navigation, interpreting a rich set of perceptual data taken from a changing (and incompletely known) vantage point, and using tools to manipulate the space station. Thus, NASA should consider supporting research on the generic problems of hands-on supervisory systems, as well as research on selected instances of the problem which it expects would yield significant practical gains.

Nature of the Problem

A fundamental defining characteristic of the system supervisor problem is uncertainty in the supervisor's knowledge of the system under study. A supervisor can almost never have complete and certain knowledge of the exact state of the system, of the rules that determine how one system state will give rise to the next, or of the exact effects of its control actions on the system. This characteristic alters dramatically the nature of diagnostic and control tasks. For example, given a perfect model of the system under study, a program might derive an open-loop control sequence to place the system in some desired state. However, in the absence of a perfect model, controlling the system requires interleaving effector actions with sensory observations to detect features of the system state.

The types and degrees of uncertainties faced in system supervision problems vary, of course, with the specific task. For instance, the task of monitoring a digital circuit might correspond to an extreme point in the spectrum of possibilities, since circuits schematics do, in fact, provide a very detailed model of the system, and since observing digital signal values is (by design) a relatively unambiguous

task. It is probably no accident that several of the earliest attempts to construct AI troubleshooting aids were conducted in the domain of digital circuitry. However, that work showed that even in this domain it was very difficult to troubleshoot circuits based only on the knowledge available from the circuit schematic (Davis, 1984). The problem is that circuit behavior can depend on thermal effects, physical proximity of components, and other factors which are not typically reflected in a circuit schematic. Furthermore, it is precisely in troubleshooting situations that such effects become significant to determining the system's behavior. The problem of incomplete knowledge in modeling subsystem behaviors is even more difficult when one considers systems with combinations of electrical, mechanical, chemical, and biological subsystems.

In addition to uncertainty in modeling the expected behavior of the system under study, the difficulty of interpreting sensory input adds another kind of uncertainty in many domains. In the digital circuit world, it is fairly straightforward to observe the value of a desired signal, though it is rare that circuits are constructed so that every signal is brought outside the circuit for troubleshooting purposes. If the system under study is a chemical process rather than electrical, detecting relative concentrations of chemicals can often be a more complex task. In mechanical systems, detecting exact locations and forces is generally out of the question. If the system is the exterior of the space station and the sensors are video cameras, then the difficulty of sensing the exact location and physical condition of each subcomponent can itself become such an overwhelming task that the observations themselves must be treated as uncertain.

Yet another dimension of uncertainty arises from the effectors that are utilized by the supervisor to alter the system under study. Again, in the circuit domain effectors such as signal generators are relatively reliable. But in the robotics domain, in which the system being supervised is the physical world, effectors such as artificial limbs may be fairly unreliable in executing actions such as grasping. In such cases, the problem of planning a sequence of actions to bring the system to a desired state must take into account nondeterminism in the effect of actions it performs.

In a sense, the ability to observe and affect the system under study and the ability to predict its behavior provide redundant sources of knowledge so that one can be used to make up for uncertainty in the other. For instance, feedback control methods utilize sensory information to make up for an incomplete model of the next-state function. On the other hand, one can make due with observing only a small proportion of the signal values in a circuit and use the model of subcomponent behaviors to infer additional signal values upstream and downstream of observed signals.

Given the various uncertainties that must be faced by a supervisory system, it is unlikely that purely algorithmic methods can be mapped out for dealing with all eventualities (although the vast NASA troubleshooting manuals indicate the degree to which this might be possible). A supervisory system will do best if it possesses redundancy to make up for the uncertainties that it must face:

redundancy in the sensors that give it information about the world, in the effectors with which it controls the world, and in the behavioral models that it uses for reasoning about the system under study. While such redundancy can help reduce uncertainty, it will not be eliminated, and the supervisor must therefore employ problem solving methods designed to operate under incomplete information. All of these needs suggest the importance of combining heuristic methods with deductive methods for reasoning about the system under study. Finally, these same problem characteristics that suggest the utility of employing AI methods (the need for flexibility in solving problems despite uncertainty) also suggest the importance of including humans in the problem-solving process. Even by optimistic estimates, it seems unlikely that AI systems will be able to completely replace human judgement in many supervisory tasks, though they may well augment it in many tasks. Thus, in many cases we envision cooperative problem solving involving computer systems and humans. Section "Sharing and Transferring Expertise in Man-Machine Problem Solving" discusses issues related to man-machine cooperation in this regard.

Research Recommendations

What research should be supported by NASA in order to maximize the future availability of hands-on supervisory systems of the kind described above? This section lists some areas that seem especially important, though the list is certainly not intended to be complete.²

- Modeling system behavior at multiple levels of abstraction. At the heart of the ability to supervise a system lies the ability to model its behavior. Systems theory provides one body of (primarily quantitative) techniques for describing and reasoning about systems. AI has developed more symbolic methods for describing and reasoning about systems, given a description of their parts structure. A good deal of research is needed to further develop appropriate behavior representations for a variety of systems at a variety of levels of abstraction, and for inferring behavioral descriptions from structural descriptions. In addition, work is needed on automatically selecting from among a set of alternative models the one most appropriate for the task at hand. For example, one useful research task might be to develop a program which can be given a detailed schematic of a large system (e.g., a computer) as well as a particular diagnostic problem (e.g., the printer is producing no output), and which returns an abstract description of the system which is appropriate for troubleshooting this problem (e.g., an abstracted block diagram of the computer focussing on details relevant to this diagnostic task).
- Planning with incomplete knowledge. The planning problem is the problem of determining a sequence of effector actions which will take the external system to a desired state. This problem has

been studied intensely within AI, especially as it relates to planning robot actions in the physical world. However, current planning methods make unrealistic assumptions about the completeness of the robot's knowledge of its world, and of its knowledge of the effects of its own actions. New research is needed to develop planning methods that are robust with respect to uncertainties of the kinds discussed above. One useful research task here would be to develop methods that produce plans which include sensor operations to reduce anticipated uncertainties in the results of effector actions, and that include conditional branches in the plan to allow for "run-time" decisions based on sensory actions.

- Integrating methods from control theory with symbolic control methods. Problems of system control, diagnosis (identification), and monitoring have been studied for some time in fields such as system control theory. Such studies typically assume a quantitative, mathematical model of the system under supervision, whereas AI methods model the system in a symbolic, logical formalism. System theory has developed various methods for using sensory feedback to make up for uncertainty in the model of the system under supervision, but these methods are difficult to apply to complex planning problems such as determining a sequence of robot operations to repair a failed door latch. Still, both fields are addressing the same abstract problems. Very little attention has been paid to integrating these two bodies of work, and research on both vertical and horizontal integration of these techniques should be supported.
- Automatically refining the supervisor's theory of system behavior through experience. As discussed in the previous subsection, a major limitation on the effectiveness of a supervisor lies in its uncertain knowledge of the system under supervision. Therefore, methods for automatically refining the supervisor's knowledge of the system would be extremely useful. In AI, research on machine learning and automated theory formation should be supported as it applies to this problem. The integration of this work with work in systems theory on model identification should also be explored. Possible research tasks in this area include developing robot systems that build up maps of their physical environment, and systems that begin with a general competence in some area (e.g., general-purpose methods for grasping tools) and which acquire with experience more special purpose competence with experience (e.g., special methods for most effectively manipulating individual tools).
- Perception from multiple sensors. One method for reducing uncertainty in the supervisor's knowledge of the system's state is to allow it to use multiple, redundant sensors. Thus, a robot might use several video cameras with overlapping fields of view, placed at different vantage points, together with touch

sensors, range finders, infrared sensors, etc. Or a supervisor for monitoring a power supply system might utilize a set of overlapping voltage and current sensors together with chemical sensors, heat sensors, etc. The benefits of using multiple sensors is clear--they provide more information. However, in order to make use of the increasing amounts of data available from multiple sensors, research is needed to develop more effective sensory interpretation/perception methods for individual sensors, and for fusing data from several sensors. An example research task here might be to develop a system that employs a number of video cameras, and which determines the correspondence between image features of the various images. A more ambitious project might try to predict image features likely to be found by one camera, based on information from other touch, video, and heat sensors.

- Representing and reasoning about 3D geometric properties. For supervisors that possess mobile sensors or effectors, a variety of problems exist in reasoning about navigating through space, and in reasoning about 3D mechanical linkages such as those that couple a robot arm to a screw via a screw driver. Research is needed on representing 3D objects (including empty space) in ways that allow for efficient computation of relations among objects, such as intersections (collisions), unions, possible packings, etc. Furthermore, since manipulating the world involves constructing temporary mechanical linkages among objects (e.g., among a robot arm, screw driver, screw, and wall), research is needed on efficiently representing and reasoning about such linkages so that effector commands can be planned that will achieve desired effects. While special-purpose robots operating in special-purpose environments can sometimes avoid using general methods for reasoning about 3D geometry, general purpose systems expected to solve unanticipated problems will require this capability.
- Designing systems to minimize difficulty in observing and controlling them. Given the great difficulties in the supervisory task that are introduced by uncertainty, one obvious reaction is to try to design the space station to reduce the uncertainties that automated supervisors will face. In short, the station should be designed to maximize the observability and controllability of those features which the supervisor will need to sense and effect. In the case of a supervisor with immobile sensors and effectors, such as a system to monitor the power supply, this requires that a broad and redundant set of sensors and control points be built into the power supply at design time. In the case of mobile supervisors, the observability of the station can be engineered, for example, by painting identifying marks on objects which will ease problems of object identification and of registering images obtained from multiple viewpoints. Similarly, the controllability of the physical

space station can be enhanced, for example, by designing all its parts to present the same simple grasping point. While a good deal of anecdotal experience has been obtained on designing robot workstations to maximize their controllability and observability, little exists in the way of a science for designing such easily-supervised systems. Research in this area, if successful, could significantly reduce the number of technical problems that automated supervisors in the space station will face.

- Feasibility of replacing hardware subsystems by software emulations. For immobile supervisors which monitor subsystems such as power supplies, navigation systems, etc., one intriguing possibility is that they might be able to substitute additional computation in place of failed hardware. For example, consider a subsystem, S, with a failed thermostat, T1. If S is being supervised by a computer system with a good model of the subcomponents of S, then this supervisor might be able to keep S working acceptably by substituting its own simulated output of T1 for the output of the failed thermostat. The degree to which this is possible will depend, of course, on (1) the veracity of the supervisor's model of S, (2) the access the supervisor has to other sensors in S (the more redundant, the better), and (3) the ability of the supervisor to control the point in S corresponding to the output of T1. While a software simulation might be slower and less accurate than a working thermostat, the advantage of substituting software for failed hardware is clear. Perhaps a small number of high-speed processors (such as parallel processors that have been developed for circuit simulations) could be included in the space station precisely for providing high-speed backup for a wide range of possible hardware failures. While the feasibility of adding robustness to the space station by adding such computational power is unproven, the potential impact warrants research in this direction.

SHARING AND TRANSFERRING EXPERTISE IN MAN-MACHINE PROBLEM SOLVING

As noted in the previous section, the same problem characteristics that argue for flexibility and adaptability in computer supervisory systems also argue for allowing humans to participate in problem solving and decision making processes. As the complexity of computer support for the space station grows, the need for communication and shared responsibility between the computer and space station residents will grow as well. If ever we reach the stage of a fully automated, self-supporting space station, we are likely to first spend a significant period of time in which computer assistants will provide certain fully-automated services (e.g., simply monitoring station subsystems to watch for unexpected behavior), but will require interaction with their human counterparts in responding to many novel

events. Effective methods for such man-machine interaction will encourage the introduction of computer assistants for many more tasks than possible if totally automated operation were demanded. This section considers some of the research issues related to developing effective communication between AI systems and their users. Since several other symposium participants will address the issue of man-machine communication in general, I will try to focus this section on issues specific to sharing problem solving responsibilities and to transferring expertise from humans to their computer assistants.

Shared responsibility is a desirable characteristic whenever one is faced with a multifaceted task for which humans are best suited to some facets and machines to others. Humans use mechanical tools (e.g., wrenches) and computational tools (e.g., pocket calculators) for exactly such reasons. In the space station, we may find it desirable to share responsibility in motor tasks, as in a human controlling the mechanical robot arm in the space shuttle, in cognitive tasks, as in a human and computer system working jointly to troubleshoot a failed power supply, or in perceptual tasks, in which a human may assist the computer in finding corresponding points in multiple camera images so that the computer can then apply image analysis and enhancement procedures to the images. In each case, shared responsibility makes sense because the machine has certain advantages for some aspects of the task (e.g., physical strength and the ability to operate in adverse environments) while the human possesses advantages for other aspects (e.g., motor skills and flexibility in dealing with the unanticipated).

Sharing in the process of problem solving also raises the prospects for transfer of expertise. In many fields, humans learn a great deal by acting as an apprentice to help a more advanced expert solve problems. As the medical intern assists in various hospital procedures, he acquires the expertise that eventually allows him to solve the same problems as the doctor to whom he has apprenticed. One recent development in AI is a growing interest in constructing interactive problem solving systems that assist in solving problems, and that attempt to acquire new expertise by observing and analyzing the steps contributed by their users. This section argues that research toward such learning apprentice systems is an important area for NASA support.

An Example

In order to ground the discussion of shared responsibility and learning apprentices, we briefly summarize a particular knowledge-based consultant system designed to interact with its users to solve problems in the design of digital circuits. This system, called LEAP (Mitchell et al., 1985), is a prototype system which illustrates a number of difficulties and opportunities associated with shared responsibility for problem solving.

LEAP helps to design digital circuits. Users begin a session by entering the definition of some input/output function that they would like a circuit to perform (e.g., multiply two numbers). LEAP provides

assistance in designing the desired circuit, by utilizing a set of if-then rules which relate desired functional characteristics to classes of circuit implementations. For instance, one rule in this set dictates that "IF the desired function requires converting an input serial signal to an equivalent parallel signal, THEN one may use a shift register." LEAP utilizes these rules³ to suggest plausible refinements to the abstract circuit modules that characterize the partial design at any given stage.

Figure 3 depicts the interface to LEAP as seen by the user. The large window on the right contains the circuit abstraction which is presently being designed by the user/system. As shown in the figure, the circuit consists at this point of two abstract circuit modules. For each of these circuit modules, LEAP possesses a description of the function to be implemented. At any point during the design, the user selects one of the unimplemented circuit modules to be considered, and LEAP examines its rule set to determine whether any rules apply to this module (i.e., rules whose preconditions match to the specifications of the circuit module). If LEAP determines that some of its rules apply to this situation, it presents the recommendations associated with these rules to the user. The user can then examine these options, select one if he wishes, and LEAP will refine the design accordingly. Figure 4 depicts the result of such an implementation step. Should the user decide that he does not want to follow the system's advice, but instead wishes to design this portion of the circuit manually, he can undo the rule-generated refinement and use LEAP as a simple, graphics-oriented, circuit editor.

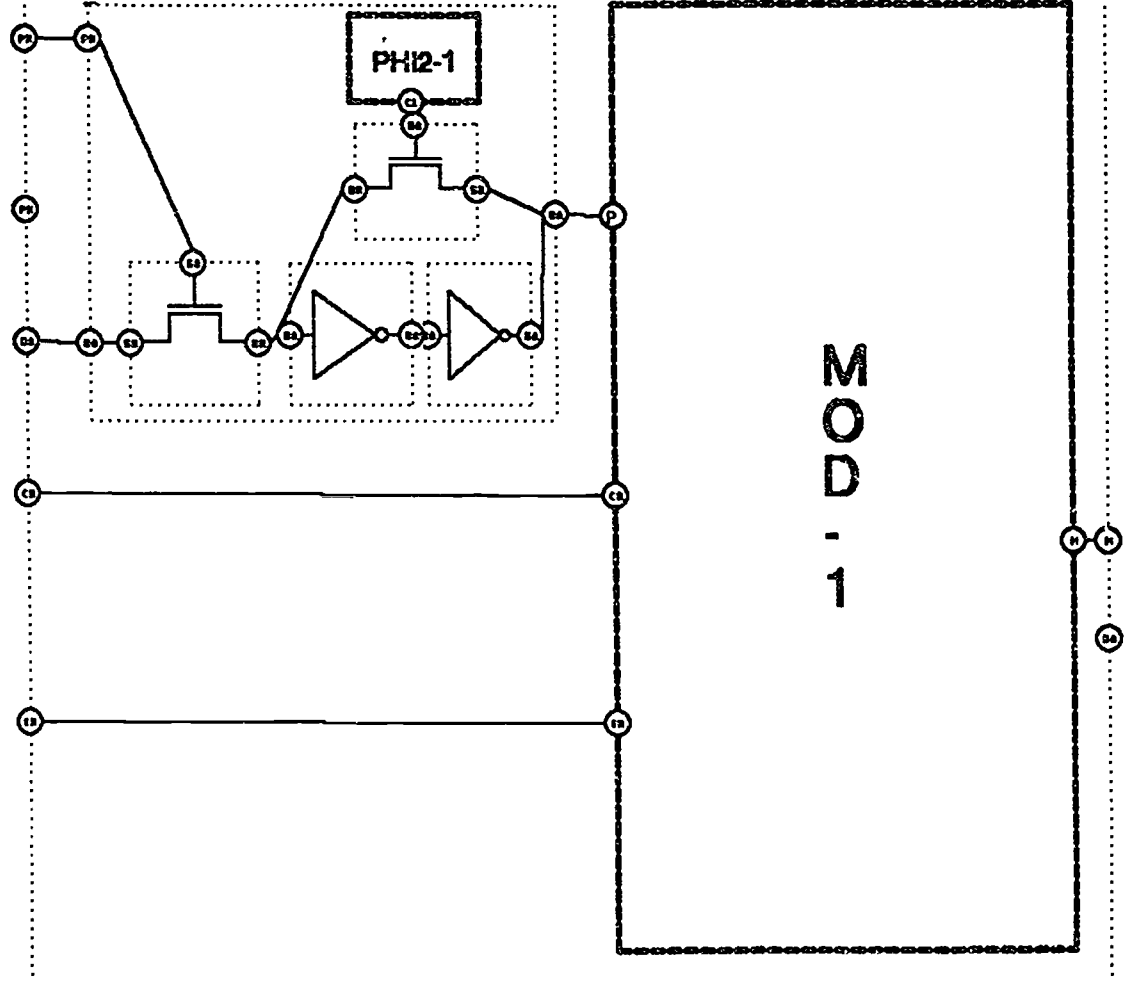
LEAP provides a simple example of shared problem solving between man and machine. The user directs the focus of attention by selecting which circuit module to refine next. LEAP suggests possible implementations of this module, and the user either approves the recommendations or replaces them with his own. LEAP thus acts as an apprentice for design. For design problems to which its rule base is well-suited, it provides useful advice. For circuits completely outside the scope of its knowledge it reduces to a standard circuit editing package, leaving the bulk of the work to the human user. As the knowledge base of LEAP grows over time, one would expect it to gradually take on an increasing share of the responsibility for solving design problems.

LEAP also illustrates how such knowledge-based apprentices might learn from their users (Mitchell et al., 1985). In particular, LEAP has a primitive capability to infer new rules of design by observing and generalizing on the design steps contributed by its users. In those cases where the user rejects the system's advice and designs the circuit submodule himself, LEAP collects a training example of some new rule. That is, LEAP records the circuit function that was desired, along with the user-supplied circuit for implementing that function. LEAP can then analyze this circuit, verify that it correctly implements the desired function, and formulate a generalized rule that will allow it to recommend this circuit in similar subsequent situations. The key to LEAP's ability to learn general rules from specific examples lies in its starting knowledge of circuit operation. Although it may not

Vlsi Expert Editor
©AI/VLSI 1983

QUIT
UP
DOWN
JUMP
BACKTRACK
SHOW HIERARCHY
EDIT DISPLAY
EDIT STRUCTURE
INSPECT TASK
DO FULL TASK
DO TASK

(REFINE INVB-1)
(REFINE INVA-1)
(REFINE MOD-1)



Building display for CAN-CELL
Attempting to match rules to
MEM-1BIT-1
Rule MEM-RULE fails to match
MEM-1BIT-1
Rule PASS-PAIR-RULE fails to match
MEM-1BIT-1
Rule PASS-TRANSISTOR-RULE
fails to match MEM-1BIT-1
Rule INVERTER-LOOP-MEM-RULE matches
MEM-1BIT-1
Rule XOR-NET-RULE fails to match
MEM-1BIT-1
Executing rule
INVERTER-LOOP-MEM-RULE
Module INVA-1 created.
Module INVB-1 created.
Module PTA-1 created.
Module PTB-1 created.
Module PHI2-1 created.
Data path DP-7 created.
Data path DP-8 created.
Data path DP-9 created.
Data path DP-10 created.
Data path DP-11 created.
Data path DP-12 created.
Building display for MEM-1BIT-1

Visi EXpert Editor

AI/VLSI 1983

QUIT
UP
DOWN
JUMP
BACKTRACK
SHOW HIERARCHY
EDIT DISPLAY
EDIT STRUCTURE
INSPECT TASK
DO FULL TASK
DO TASK

(REFINE MOD-1)
(REFINE MEM-1BIT-1)

MEM-1BIT-1

MOD-1

Attempting to match rules to
CAM-CELL
Rule MEM-RULE matches CAM-CELL
Rule PASS-PAIR-RULE matches CAM-CELL
Rule PASS-TRANSISTOR-RULE
fails to match CAM-CELL
Rule INVERTER-LOOP-MEM-RULE
fails to match CAM-CELL
Rule XOR-NET-RULE fails to match
CAM-CELL
Executing rule MEM-RULE
Building display for CAM-CELL

initially have the expertise to generate a particular implementation of the desired function, it does have the ability to recognize, or verify, the correctness of many of its users' solutions. In general, it is easier to recognize a solution than to generate one. But once a solution can be recognized and explained, then LEAP can generalize on it by distinguishing that certain features of the example are critical (those mentioned in the verification), whereas others are not (those not mentioned in the verification).

LEAP is still a research prototype system, and has not yet been subjected to testing on a large user community. While there are no doubt many technical issues still to be solved, it serves as a suggestive example of how a knowledge-based consultant might be useful as an apprentice even before its knowledge base has been fully developed. It also suggests how its interaction with the user might lead it to extend its knowledge base automatically. The methods for collecting training examples and for formulating general rules appear generic enough that similar learning apprentice systems might be developed for many supervisory tasks of the kind discussed in the previous section. Other current research is exploring the feasibility of such learning apprentices in task domains such as signal interpretation (Smith et al. 1985), proving mathematical theorems (O'Rorke, 1984), and planning simple robot assembly steps (Segre and DeJong, 1985).

Nature of the Problem

The LEAP system suggests one kind of shared responsibility between computer and human, as well as a mechanism for the gradual accretion of knowledge by the system so that over time it can take on a progressively greater share of responsibility for problem solving. The ability to acquire new rules by generalizing from the users' actions follows from LEAP's starting knowledge of how circuits work. That is, it begins with enough knowledge of how circuits operate, that it is able to explain, or verify, the appropriateness of the users' actions once it observes them. Once it has verified that the user's circuit correctly implements the desired function, then it can generalize on this action by retaining only those features of the specific situation that are mentioned in this explanation. Similarly, if one tried to construct such a learning apprentice for troubleshooting power supply faults, one would want to include sufficient initial knowledge about the power supply (i.e., its schematic) that the system could verify (and thus generalize on) users' hypotheses about the causes of specific power supply malfunctions.

Thus, in order for a system to learn from observing its users, it must begin with sufficient knowledge that it can justify what it observes the user do. It seems that for supervisory tasks of the kind discussed above, the primary knowledge required⁴ to construct such explanations is a description of the structure and operation of the system under supervision. Since AI has developed methods for

representing such knowledge, supervisory tasks seem like good targets for further research on learning apprentices.

In addition to cognitive tasks such as monitoring, designing, and debugging, one might consider learning apprentices for robotics tasks such as using tools (see Segre and DeJong, 1985 for one example). Given a new tool for the robot to use, one way to train it might be to use a teleoperator to guide the robot through several uses of the tool. For example, given a new type of fastener, a user might guide the robot to grasp the fastener and use it to fasten two objects together. If the system could start with enough knowledge to explain which features of its trajectory and other motions were relevant to accomplishing the given task, then it might be able to generalize accordingly. Research on such robotic learning apprentices seems worthwhile and highly relevant to the goals of the space station program.

To understand the issues involved in sharing information and responsibility between human and machine, it is instructive to consider the issues involved in sharing responsibility strictly among humans. In both cases there are certain subproblems that are best dealt with by individual agents, and others where shared responsibility makes best sense. Successful interaction requires arriving at an agreement on which agent will perform which task. In LEAP, the user makes all such choices. But in more complex scenarios the user may not want to spend the time to approve every suggestion of the apprentice. In such cases, there must be ways to agree upon a policy to determine which decisions are worth having the human approve. Of course there are many other issues that follow from this analogy as well: the cooperating agents eventually need accurate models of their relative competence at various subtasks. And there will be questions of social and legal responsibilities for actions taken.

Here we have tried to suggest that one class of computer assistants on the space station be viewed as dynamic systems that interact with their users and work toward extending their knowledge and competence at the task they perform. Preliminary results from AI suggest that this is a worthwhile research task. The nature of the space station suggests that such self-refining systems are exactly what will be needed. The continually changing configuration of the station itself, the continually changing crews and types of operations that will be conducted aboard the space station, the evolving technology that will be present, all dictate that the computer assistants aboard must be able to adjust to new problems, new procedures and new problem solving strategies over the life of the space station.

Research Recommendations

Here we suggest several areas in which NASA might support research toward advanced interfaces for interaction between humans and intelligent consultant systems.

- Architectures that support graceful transfer of expertise and responsibility. Research toward developing learning apprentice systems for space station applications is warranted based on recent AI results and on the importance of such systems to the space station program. A prudent research strategy at this point would be to support development of a variety of learning apprentices in various task areas (e.g., for troubleshooting space station subsystems, for monitoring and controlling subsystems, for managing robot manipulation of its environment). Such a research strategy would lead to experimenting with alternative software architectures for learning apprentices, as well as an increased understanding of the feasibility of constructing learning apprentices for specific space station task areas.
- Evolution of grainsize and initiative of interaction. As the expertise of the apprentice grows, and as the human becomes more familiar with the competence and communication capabilities of the computer, one expects that the optimal style of communication should shift. Changes may occur, for example, in who takes the initiative in controlling the direction of problem solving, and in the grainsize of the tasks (e.g., initially small subtasks will be discussed, but later it may be sufficient to focus only on larger grain subtasks). Research on interfaces that support these kinds of changes over time in the nature of the interaction, and which support explicit communication about such issues, should be encouraged. Such flexible interfaces are important whether the apprentice learns or not, since the user will certainly go through a learning period during which his understanding of the system's competence and foibles, and his willingness to trust in the system will change.
- Task-oriented studies of cooperative problem solving. In order to understand the kinds of knowledge that must be communicated during shared problem solving, it may be worthwhile to conduct protocol studies in which a novice human apprentices with an expert to assist him and to acquire his expertise (e.g., at a task such as troubleshooting a piece of equipment). Data collected from such experiments should provide a more precise understanding of the types of knowledge communicated during shared problem solving, and of the knowledge acquisition process that the apprentice goes through.
- Transferring knowledge from machine to man. Given the plans for a frequently changing crew, together with the likely task specialization of computer consultants, it is reasonable to assume that in some cases the computer consultant will possess more knowledge about a particular problem class than the human that it serves. In such cases, we would like the system to communicate its understanding of the problem to the interested but novice user. Certain work in AI has focused on using large

knowledge bases as a basis for teaching expertise to humans (e.g., Clancey and Letsinger, 1984). Research advances on this and other methods for communicating machine knowledge to humans would place NASA in a better position for crew training and for integrating intelligent machines into the human space station environment.

SUMMARY

This paper presents a sampling of recommended research directions which NASA may wish to support in order to accelerate the development of AI technology of particular relevance to the space station. We feel that recent AI research indicates the potential for a broad range of applications of AI to space station problems. In order for this potential to become reality, significant support for basic AI research is needed.

Research toward developing a wide range of "hands-on" supervisory systems for monitoring, controlling, troubleshooting and maintaining space station subsystems is strongly recommended. Such research is important both because of its potential impact on reliability and safety of the space station and because the technical development of the field of AI is at a point where a push in this area may yield significant technical advances. Such hands-on supervisory systems could include both physically stationary supervisory systems that monitor electronic subsystems, power supplies, navigation subsystems and the like, as well as physically mobile supervisors that monitor and repair the exterior and interior physical plant of the space station. Important technical challenges remain to be addressed in both areas.

In support of developing and deploying such knowledge-based supervisors, it is recommended that research be conducted leading toward interactive, self-extending knowledge-based systems. Such systems may initially serve as useful apprentices in monitoring and problem solving, but should have a capability to acquire additional knowledge through experience. The evolutionary nature of the space station together with the turnover of crew assure that a continually changing set of problems will confront onboard computer systems. This feature of the space station, together with the need to continually extend the knowledge of problem solvers onboard, argue for the importance of research toward interactive, self-extending knowledge based systems.

There are certainly additional areas of AI research which would also benefit the space station program. The goal of this paper is to point out a few such areas, in the hope of stimulating thought about these and other possible uses of AI in the space station.

ACKNOWLEDGEMENTS

My thanks go to Allen Newell and Oren Etzioni for providing useful comments on earlier drafts of this paper. This work was supported in part by NSF grant DCR-8351523.

NOTES

1. In fact, initial AI systems for troubleshooting and control have generally been restricted to dealing with typed-in observation inputs and to typing out their recommendations rather than exerting direct control over the system. However, there are exceptions to this, such as the YES/MVS system (Ennis et al., 1986) which directly monitors and controls operations of a large computer system.
2. The research recommendations listed here represent solely the opinion of the author, and should not necessarily be interpreted as recommendations from the symposium as a whole.
3. LEAP also utilizes knowledge about behaviors of individual circuit components, plus knowledge of how to symbolically simulate digital circuits.
4. Other relevant knowledge includes the goals of the user (e.g., a decision must be made to act within 15 seconds), and empirical data on the frequencies of various types of faults.

REFERENCES

- 1984 Artificial Intelligence, Special Volume on Qualitative Reasoning About Physical Systems. North-Holland.
- Clancey, W. and Letsinger, R.
1984 NEOMYCIN: reconfiguring a rule-based expert system for application to teaching. Pp. 361-381 in Clancey and Shortliffe, eds., Readings in Medical Artificial Intelligence. Addison-Wesley.
- Davis, R.
1984 Diagnostic reasoning based on structure and behavior. Artificial Intelligence 24:347-410.
- Ennis, R. L., et al.
1986 A continuous real-time expert system for computer operations. IBM Journal of Research and Development 30:1.

Genesereth, M.

- 1981 The Use of Hierarchical Models in the Automated Diagnosis of Computer Systems. Technical report Stanford HPP memo 81-20, Stanford University, Stanford, CA.

Mitchell, T. M., Mahadevan, S., and Steinberg, L.

- 1985 LEAP: a learning apprentice for VLSI design. Pp. 573-580 in Proceedings of the Ninth International Joint Conference on Artificial Intelligence. August.

National Aeronautics and Space Administration

- 1983 Autonomy and the Human Element in Space. Technical report NASA-ASEE 1983 summer faculty program final report. National Aeronautics and Space Administration, Stanford, CA.

National Aeronautics and Space Administration

- 1985 Advancing Automation and Robotics Technology for the SpaceStation and for the U.S. Economy. Technical report NASA technical memorandum 87566. National Aeronautics and Space Administration, Springfield, VA.

O'Rorke, P.

- 1984 Generalization for explanation-based schema acquisition. Pp. 260-263 in Proceedings of the AAAI. AAAI, Austin, TX.

Pazzani, M. J.

- 1986 Refining the knowledge base of a diagnostic expert system: An application of failure-driven learning. Pp. 1029-1035 in Proceedings of the Fifth National Conference on Artificial Intelligence. AAAI, August.

Segre, A. M., and DeJong, G. F.

- 1985 Explanation-based manipulator learning: acquisition of planning ability through observation. Pp. 555-556 in Proceedings of the IEEE Conference on Robotics and Automation. IEEE, St. Louis, MO.

Scarl, E. A., Jamieson, J. R. and Delaune, C. I.

- 1985 A fault-detection and isolation method applied to liquid oxygen loading for the space shuttle. Pp. 414-416 in Proceedings of 1985 International Joint Conference on Artificial Intelligence. International Joint Conference on Artificial Intelligence, Los Angeles.

Smith, R. G., Winston, H. A., Mitchell, T. M., and Buchanan, B. G.

- 1985 Representation and use of explicit justifications for knowledge base refinements. Pp. 673-680 in Proceedings of the Ninth International Joint Conference on Artificial Intelligence. August.

EXPERT SYSTEMS: APPLICATIONS IN SPACE

Bruce G. Buchanan

INTRODUCTION

Artificial intelligence is one of the most important trends in computing because making computers behave intelligently is at least as important as manipulating data efficiently. Opportunities for using intelligent programs in NASA space station environments are numerous and obvious. But many of those opportunities require substantial research in artificial intelligence before they can be realized. This paper looks at the technology of artificial intelligence, especially expert systems, to define "from the inside out" what capabilities exist that are relevant for applications and environments in the space station, and what research needs to be promoted in order to achieve systems better able to interact symbiotically with a variety of persons for long times in space.

Anderson and Chambers (1985) mention a number of characteristics of systems in a human-centered space station. These include:

- symbiosis with humans: human and machine capabilities may complement one another
- autonomous,
- continuing operation for a period up to 20 years,
- operating in an information-rich environment,
- consequences of interactions with humans not entirely predictable,
- maturation of system implies flexibility to accommodate operational growth and minor upgrades,
- evolution of system implies flexibility to accommodate new and enhanced functionality,
- system may include electronic crew members (ECMs),
- humans may have to learn new skills to interact productively with computers,
- computers may learn from humans,
- autonomous agents may serve a variety of roles with varying degrees of decision making power and authority.

These are some of the relevant considerations in a top-down design of systems for the space station. Each of these points implies a

research and development program of some intensity. This paper takes a bottom-up view of the same considerations--i.e., starts with what exists today and asks how we can achieve these design goals. By doing so, I hope to introduce some relevant details into the design of systems and the planning of research.

Expert systems are now being used in many decision-making situations of direct relevance to NASA's mission, spanning manufacturing, engineering, medicine, and science. At present, they are used more as "intelligent assistants" than as replacements for technicians or experts. That is, they help people think through difficult problems and may provide suggestions about what to do, without taking over every aspect of the task.

Computer programs that reason autonomously are also of extreme importance in space, but they, too, must be integrated into an environment that is centered around people. They are extensions of present technology along several dimensions discussed here, that involve all of the same principles of design as the intelligent assistant programs.

One primary consideration is why intelligent systems are necessary in space. Although there are many reasons to build an expert system, they are all based on the premise: "Expertise is a scarce resource." The corollary (by Murphy's Law) is: "Even when there is enough expertise, it is never close enough to those who need it in a hurry." Because this is true--almost by definition of the term 'expertise'--constructing expert systems that reason at the level of NASA's, or their contractors', specialists may have several benefits. These are summarized in Table 1.

WHAT IS AN EXPERT SYSTEM?

The general nature of expert systems is familiar to everyone within NASA. A reiteration of the four major characteristics is provided below to help define the most important dimensions for research and development efforts.

An expert system is a computer program with expert-level problem solving abilities, which also fits some other criteria: it is a

TABLE 1 Some Perceived Benefits of Expert Systems: Responses of 86 Users of Knowledge Engineering Tools

-
1. Replicate expertise
 2. Preserve expertise
 3. Increase productivity and cost savings
 4. Free human experts for more demanding problems
 5. Provide expert consultations to inexperienced staff
-

SOURCE: Bauman (1984)

symbolic reasoning program that uses heuristics, its reasoning and knowledge base are understandable, and--most importantly--it is flexible. These characteristics are discussed below. All are important for applications in the space station, and all define research topics that will enhance current capabilities.

Example

One well-known expert system that has become a classic, although not actively used, is MYCIN. It was developed at Stanford by E. H. Shortliffe and others in the mid-1970s. Its task is two-fold: (a) diagnose the cause(s) of infection in a patient and (b) recommend appropriate drug therapy. From a medical perspective, MYCIN's knowledge base is now dated; from the perspective of expert systems it represents much of the kind of reasoning that is captured in today's systems. MYCIN's conclusions were demonstrated to be equal in quality to those of infectious disease specialists at Stanford Medical Center.

The sample typescript shown in Appendix A illustrates MYCIN's requesting information about a case and reasoning to conclusions about the best treatment.

Performance

Naturally we want computer programs to solve problems without error. But that is not always possible--in fact, outside of mathematics and logic we don't have flawless methods we can put into programs. Specialists in engineering, science, education, the military--and every area outside of pure logic--must solve problems with less than perfect methods. How do they do it? Mostly by building up specialized knowledge through extra years of training and experience and by reasoning carefully with that knowledge in situations they have learned to recognize. They are not infallible, though. Specialists' decisions are challenged frequently--most noticeably in the courts. So it is also unreasonable to expect computer programs to reason infallibly in all of these areas. Occasionally new methods are discovered that provide much better results than the established methods of the old practitioners. But these improvements can then be put into programs, thus raising the overall standard of performance while still keeping the same relative standard of comparison with the best specialists.

Reasoning

When we say that expert systems are reasoning--and not just calculating with numbers--we are saying that they belong to a class of programs using the methods of artificial intelligence (hereafter AI). In the 1940s, computers were used almost exclusively for large mathematical problems. At Los Alamos, for instance, scientists had to solve complex mathematical equations in order to calculate elements in the design of

the atomic bomb. These applications are usually referred to as large-scale scientific computation, or "number crunching" for short. In the 1950s, IBM and other computer manufacturers, realized the enormous value in helping business solve problems of record keeping, payroll and the like. These applications extended the concept of computer-as-calculator to computer-as-data-manager.

In both of these classes of applications, the method of computation is error-free. There is no question that the result is correct, providing of course that the computer has been programmed correctly. A mathematical equation is solved correctly; an employee roster is sorted correctly--if the methods are followed precisely. And computers are better able to follow complex instructions than people are. In computer science, logic and mathematics we call these procedures algorithms. They are procedures that can be guaranteed to provide a correct answer in a finite time, if there is one, and otherwise will provide a statement that the problem is not solvable.

Some algorithms are too expensive to use, however, even in computers. A classic example is finding the shortest route a travelling salesman can take to visit many cities once and end up at home. With more than a handful of cities, algorithmic methods will not finish in time to be useful. For this reason, alternative methods have been developed.

Around the mid 1950s and early 1960s an alternative style of computing came to be recognized as important. Instead of always using algorithms, a computer may use heuristics--rules of thumb that aid in finding plausible answers quickly without guaranteeing the correctness of the results. Sometimes these rules of thumb are introduced into large numerical simulations in order to get the simulations to crank out answers more quickly. Or approximate methods may be substituted for more precise ones for the same reason. The assumptions may not all be correct; thus the results of the simulation may not be correct.

When heuristic (non-algorithmic) methods are combined with symbolic (non-numeric) data, we are dealing with that part of computer science known as artificial intelligence.

Understandability

When someone truly knows something, he or she can "give an account" of what he knows. In our terms, good performance is not enough to call a person (or program) an expert--he/she (it) should also be able to explain why the solution is plausible, what features of the situation were noted to be important, what knowledge and problem solving methods were used. Otherwise we label a person as "consistently (but unaccountably) lucky", or maybe "psychic". Each field has its own standards of what a reasonable explanation is. A surgeon who recommends amputation of a leg generally talks about the process of disease or extent of injury and what will happen if it is not amputated. A broker who advised liquidation of one's stock portfolio may explain the advice with respect to technical charts, historical trends, or some economic principles that point to a stock market

collapse. In their own communities, both the surgeon and the broker can usually justify--in court if necessary--the advice they give. And we regard them as experts partly because they have the knowledge that lets them do this.

Flexibility

We expect experts to be flexible in their thinking. And we regard persons as amateurs, not experts, when we encounter opinions that are rigid, locked-in ways of dealing with problems, or an inability to deal with new situations.

In particular, there are two situations in which we want expert systems to be flexible:

1. At advice-giving time we want the program (or a person) to provide good advice about situations that have never been encountered before. Novices with good memories may be able to provide the "textbook" answers for classic situations. Experts however, should, in addition, be able to reason about novel situations.
2. At the time a program is being constructed or modified (or a person is learning), we want it to be flexible enough to assimilate new bodies of information. There should be a capacity for growth of knowledge, not a rigidity that freezes either the depth or breadth of the program's knowledge.

SOME APPLICATIONS

Some of the types of problems for which expert systems have been constructed are shown in Table 2. Many of these, such as small troubleshooting assistance programs, are relatively straightforward. Although the state of the art is difficult to quantify, the programs in the table represent the kinds of commercially robust systems that can be built for NASA today, provided adequate resources and an appropriate problem. We don't have an adequate taxonomy of problem types. Many of these overlap, in being different forms of data interpretation, for example. Even this brief characterization, however, provides a reasonably good idea of what expert systems can do.

In general, expert systems can reduce costs or increase quality of goods and services--in a single phrase, they can increase productivity in an organization. If you believe either that there is not enough expertise in the world, or that it is not well distributed, then you will be willing to entertain the idea that putting human expertise into an easily-replicated form may answer some productivity problems. Or, at least expert systems may provide a partial answer. Consider medical

TABLE 2 Some Problems for which Expert Systems are Providing Solutions

RISK ASSESSMENT	St. Paul Insurance Co.--assess risk of insuring large commercial clients
EQUIPMENT DIAGNOSIS	General Motors--determine causes of vibration noises and recommend repairs
CONFIGURATION	Digital Equipment Corp.--translate customers' orders for computer systems into shipping orders
LOADING	U. S. Army--design loading plan of cargo and equipment into aircraft of different types
ROUTE PLANNING	U. S. Air Force--plan an aircraft's route from base to target and back to avoid detection and threats when possible
DATA INTERPRETATION	Schlumberger--interpret down-hole data from oil well bore holes to assist in prospecting
SCHEDULING	Westinghouse--plan manufacturing steps in a plant so as to avoid bottlenecks and delays
THERAPY MANAGEMENT	Stanford Medical Center--assist in managing multi-step chemotherapy for cancer patients
MONITORING	IBM--monitor operations of MVS operating system
SCREENING	U. S. Environmental Protection Agency--determine which requests for information fall under the exceptions to the Freedom of Information Act
PORTFOLIO MANAGEMENT	First Financial Planning Systems (Travelers Insurance)--analyze an individual's financial situation and recommend types of investments
TROUBLESHOOTING MANUFACTURING STEPS	Hewlett Packard--diagnose causes of problems in photolithography steps of wafer fabrication
CROP MANAGEMENT	Virginia Polytechnic Institute - <POMME> Assist in managing apple orchards
EQUIPMENT DESIGN	Delco--design special-purpose, low voltage electric motors
TRAINING	Elf Aquitaine Oil Company--demonstrate reasoning to find cause of drill bit sticking in oil well and to correct the problem
SOFTWARE CONSULTANT	Shell Oil Corporation--advise persons on which subroutines in large FORTRAN library to use for their problems and how to use them
EQUIPMENT TUNING	Lawrence Livermore National Laboratory--specify parameter settings to bring a sensitive instrument (triple quadrupole mass spectrometer) into alignment

NOTE: Many more examples are listed in Buchanan (1986), and Harmon (1986).

diagnosis. Specialists at university medical centers generally see more of the unusual disorders than a rural practitioner and thus stand a better chance of diagnosing them correctly. Putting some of that expertise more directly at the service of the rural practitioner could allow more effective treatment, and save patients the time and trouble of travel to the medical center.

Or consider troubleshooting a complex piece of equipment. Persons with the most field experience are often the ones promoted to desk jobs in the central office. When subtle combinations of causes keep a less experienced field service technician from fixing a mechanical failure, someone with more expertise is needed. On earth, depending on travel times and the criticality of the work flow in the central office, calling the experienced specialist out may be a very expensive repair procedure.

The following situations are all cases where it may make good sense to build an expert system:

- too few specialists for the number of problems;
- specialists not at the sites of problems when they occur;
- long training time for a specialist;
- high turnover among technicians;
- combination of complex equipment and poorly trained technicians;
- organization's best (or only) specialist in an area is nearing retirement;
- too many factors for a person to think through carefully in the time available.

KEY CONCEPTS

The four goals that characterize expert systems can be achieved with a few key methodological ideas. In this section, the key ideas will be introduced; in successive sections they will be elaborated on so as to explain a little how they work. The main organizational principle of expert systems is to keep specialized knowledge separate from the logical and heuristic inference methods that use it. This is easy to say but difficult to follow, for reasons that will be described later.

Keep Domain-Specific Knowledge
Separate from General Reasoning Methods

-- KEY IDEA #1 --

Another key concept, which is imported from principled design of software generally, is modularity. (The first key idea is an instance of this, but that instance has taken on more importance than all the other instances of the general concept.) Modularity at the level of knowledge about the problem area implies conceptual separation of elements in the knowledge base. For example, medical knowledge about penicillin, although not totally independent, can often be separated from knowledge of other drugs. It can be modified in major ways, or deleted, without altering the program's knowledge of other drugs. So, this is to say that the concepts used to talk about objects in the domain should be chosen so as to allow talking separately about an individual object, a single

property of an object, or a single relation of one type of object with another. Modularity at the level of programming constructs implies that the program's internal representation of knowledge elements (e.g., objects, properties, relations) is similarly "clean".

Keep independent pieces of knowledge independent.
Keep the rest as nearly-independent as possible.

--KEY IDEA #2 --

A third key concept is uniformity of conceptualization and representation of knowledge. The underlying intuition is that it is easier for a person or a program to build, understand, and modify a body of knowledge if it doesn't mix and merge a variety of different types of things. This is as true at the knowledge level as at the programming level. For instance, one of the most compelling aspects of Newton's Laws is that all physical bodies are treated as quantities with mass. He didn't need one set of laws for planets and another for apples. So it is desirable to build an expert system with a "conceptually clean", well-organized, simple collection of concepts. And it is important to use a simple, well-organized collection of programming constructs as well. Otherwise there are too many different kinds of things to keep track of and reason with.

There is more dispute among AI specialists about this principle. There are good reasons to violate it, as we shall see, in the interest of being able to say more about the objects and relations of interest than can conveniently be said in a single language. We are frequently told by bi-lingual friends, for instance, that there are some concepts that just can't be expressed fully in English. The same is true for programming constructs, but the basic principle for constructing expert systems is to try to maintain uniformity as much as possible.

Strive for uniformity of language
and programming constructs

-- KEY IDEA #3 --

A fourth principle is to design the expert system to mirror the ways experts think about problems in their domains. That means using the same terms and the same rules of reasoning as the experts use. One reason for this is that building and debugging a knowledge base depends necessarily on the expert, and using less familiar terminology or methods will introduce confusion and error before the knowledge base is completed. Also, after it is completed it needs to be comprehensible and unambiguous to the practitioners using the system or else confusion and error will result.

Note that we are assuming here that the expert designing the system knows how to make it understandable to users. Great care must be taken when building a system, however, to insure that this assumption is true.

There are times when this principle will be, or should be, violated. For example, when efficient computer algorithms can solve part of a problem, it doesn't often make good sense to use anything else for that part, even if the experts don't think about it in that way.

As much as possible, use the same vocabulary and methods in the program as the experts and practitioners use.

-- KEY IDEA #4 --

These key ideas help us achieve all of our four goals in the following ways.

- PERFORMANCE--in problems whose solution methods are not already well formalized, which are considerable, much of the effort in building a knowledge base from an expert system lies in building the conceptual framework. Which properties and relations of objects to describe is often not well specified at the beginning. So the knowledge base is built incrementally, where experience with one knowledge base guides future modifications, extensions, or reformulations.
- REASONING--When the solution methods are not well characterized, it is important to encode heuristics that experts say they use. Storing these separately and in a simple form allows them to be changed easily. Since it is nearly impossible for an expert to articulate a complete and consistent set of heuristics at one sitting, it must be easy to add, remove, or modify the heuristics that determine the reasoning.
- UNDERSTANDABILITY--with modularity, individual elements of the knowledge base can be displayed meaningfully in isolation. Moreover, with the separation of knowledge base and inference procedures it is possible to peruse the knowledge base in order to find just those elements that were used to reason about a new case. And with uniformity of data structures, it is possible to build one set of procedures that produce explanations.
- FLEXIBILITY--when the elements of the knowledge base are in separate data structures, and not intertwined with code for inference procedures, we can add more knowledge with considerably more ease. When the individual items in the knowledge base are nearly separate, we have fewer interactions to worry about when we change an items. And when the representation is homogeneous, we can more easily write other programs that act as "editing assistants" or explainers that help us insure correctness of new items and help us understand what is in the knowledge base.

PERFORMANCE ISSUES

Expert systems constitute one class of computer programs. As such, they work the same way as every other program: they process input data to produce output data. But the nature of the processing is different from most conventional programs. The key ideas mentioned earlier are the key differences in the design and implementation of expert systems.

In order to design a reasoning program, we need to provide knowledge to reason with and reasoning methods to use. Both are needed. A

powerful thinker needs something to think about, and a body of facts without methods for using them is sterile. Over the last few decades, research in AI has elucidated programming methods for making inferences and storing knowledge. We briefly characterize these topics below, although with some reservations about oversimplifying, in order to highlight research issues relevant to increasing the performance of expert systems. In addition to research on inference methods and representation of knowledge, several other issues are mentioned briefly as needing more research in order to improve the performance of expert systems.

Inference Methods

Aristotle's theory of the syllogism defined acceptable inference methods outside of mathematics for about 2000 years. His theory has been extended in this century by Russell & Whitehead, and others, in a formal theory that includes methods of reasoning with several statements and several variables in an argument.

Formal logic defines several inference rules which are guaranteed to create true conclusions if the premises of the argument are true. The chain rule (modus ponens) is the single most important inference rule in expert systems. It allows us to chain together a string of inferences:

```

If A then B
If B then C
If C then D
  A
-----
  D

```

Many of the inferences we make in our lives are not guaranteed by the rules of logic, however, nor do we have certain knowledge about the truth of our premises. Whenever we argue that the future will be like the past, as in stock market predictions, we have to be prepared for exceptions. These inferences, labeled "plausible inferences" by George Polya, are the ones of most interest in AI.

One set of programming methods were in AI for making plausible inferences is to assert the facts categorically--as if they were known to be true with certainty--and then reason about exceptions that might force revisions to the conclusion.

Another set of methods deals explicitly with the degrees of uncertainty in the facts and in the associations. MYCIN (see Appendix A) uses this style of reasoning. Usually the degrees of uncertainty implied by words like "often" and "may" are expressed as numbers. And often these numbers are interpreted as probabilities.

A third, and most powerful, set of methods is to introduce heuristic rules, or rules of plausible inference, into the reasoning. These are facts or relationships that are not guaranteed to produce correct conclusions, but will often do so. Moreover, they often produce answers more quickly than their algorithmic counterparts. In the

traveling salesman problem, for example, the problem is to plan a route for visiting each city in a set exactly once and end at the home city. This is an NP-complete problem, that is, the algorithm for solving it takes times that is exponential with the number of cities. One heuristic we may introduce is to go to the nearest city that has not yet been visited. This certainly speeds up the computation of the route, but may (and probably will) miss the route that is shortest overall. Some rules of plausible inference used, with caution, in some expert systems are shown below:

- **Satisficing:** If it will be expensive to find the very best solution to a problem, then stop with the first solution that satisfies easier criteria of being good enough.
- **Inheritance:** (Some specified) properties of a whole are shared by all its parts. E.g., An ice cube is cold and hard. Pieces of an ice cube are cold and hard. [But other properties, like "weight", do not behave the same.]
- **Single Fault:** If a piece of equipment (or any organized system) is malfunctioning, and one hypothesis explains the problem, then there probably is only a single cause of the problem.
- **Compelling Evidence:** If you have gathered a lot of evidence in favor of hypothesis H1, and very little evidence against it, and you have gathered little positive evidence for alternative hypotheses, then H1 is a plausible hypothesis.
- **Decomposability:** If there are many parts to a problem that are nearly independent, assume they can be solved independently. Then adjust the composed solution to take account of known interactions.
- **Parsimony of Design:** Designs or plans with fewer elements are preferred to those with more.

In principle, the rules of inference (both logical and plausible) may be applied again and again to a situation description, in any order, and the resulting conclusions will be the same. This is not always possible in practice, however. There may not be enough time to reason exhaustively about all possibilities and contingencies. For that reason AI researchers talk about controlling the inferences as being a more important, and more difficult, problem than making the inferences.

Controlling inferences breaks down into two subtasks: (a) deciding which rules to apply now, at this stage of the problem-solving process, and (b) deciding which part of the problem to work on now. Since we believe these subtasks require some intelligence, all of the principles for building knowledge-based systems also apply at this level of reasoning. In particular, it is desirable to make this control knowledge explicit and separate from the inference methods

Representation of Knowledge

We have said that a key idea in building expert systems is storing knowledge separately from the inference methods. Another key idea was to avoid, as much as possible, representing it in a low-level computer language. But we have not said how to represent for the computer what an expert wants to tell it. English is too difficult for a computer to interpret unambiguously; FORTRAN and BASIC are too low-level for an expert to deal with efficiently. Clearly we need some stylized representations that are somewhere in between.

AI researchers have developed several different representation methods. There is no single one that is best in every case--they each have strengths and weaknesses. One of the fundamental trade-offs in thinking about the representation of knowledge is between simplicity and expressive power. We want a simple set of conventions for storing knowledge because that makes it easier--for a person or a program--to understand what is in the knowledge base at any moment. It is also easier to write simple statements without error. Aristotelian logic ("All As are Bs", etc.) and arithmetic are simple representations. The difficulty is they lack the expressive power to let us say everything we think is important about a problem. A hundred years ago DeMorgan noted the lack of expressive power is Aristotelian logic (and a weakness in its inference methods): if you know that all horses are animals, he said, you cannot prove that the head of a horse is the head of an animal. This sort of problem led Russell & Whitehead to develop a formalism with more expressive power.

There are two major classes of representation methods, reflecting two different ways of viewing the world: action-centered or object-centered. Different problem areas may focus on one or the other, or different experts in the same problem area may. For example, physicians talk about disease and classes of diseases as entities with expected properties and also talk about clinically relevant actions that determine what to do--e.g., asking questions, measuring things, relating signs and symptoms to possible causes, matching likely causes to acceptable therapies. Neither point of view is wrong, but they focus on medical phenomena quite differently. And an expert system would similarly have one focus or the other.

Action-centered representations focus on conclusions that can be drawn from facts or, more generally, on relations between situations and actions. The formalism of mathematical logic is one popular choice. Another popular formalism is rules.

Object-centered representations focus on the organization of objects in the world, for instance into hierarchies. They still allow conclusions to be drawn when an object is found to have some properties, but those inferences are triggered from "within" an object rather than from outside. That means that objects and their properties--and changes to any of them--drive the inferences. But in an action-centered model, the inference rules drive the creation of new objects and properties. The net effect may be identical, as we said, but the way one thinks about the domain of discourse is distinctly different.

Also, in object-centered representations there is more machinery for saving storage space by using hierarchies. Properties of classes of object, for example, may be implicitly inherited by all of the instances without having to store it with each instance. The manager of a group is the manager of each person in the group, so the program only needs to store (once for each group) the name of the group manager and can use that, plus the class-instance hierarchy, to find the name of any individual's manager.

There are as many different conventions for representing knowledge as there are AI researchers working on this topic. This can be confusing when reading the literature. But they are basically all variations--usually mixtures--of the two different styles just discussed. There are many expert systems built out of these two sets of ideas, but considerably more experience--and analysis--is necessary to understand their strengths and limitations.

Validation and Robustness

It is impossible to prove logically that the contents of an expert system's knowledge base are correct or complete or that the inference procedures will always provide the best answers. Yet persons in a space station whose equipment and lives depend on the expertise of many systems need to know the scope and limits of each system. Or, alternatively, they need tools for determining the scope and limits of the programs they use. These range from better explanation systems to tools for checking knowledge bases.

Spatial and Temporal Reasoning

Many complex problems in a space station require autonomous computer programs that represent and reason about three-dimensional objects. Simpler representations do not allow programs to solve problems involving 3-d shapes and positions, such as problems of fitting parts or of maintaining some equipment. Building expert systems requires attention to making the systems' reasoning understandable to persons onboard the space station and changeable by them. That, in turn, requires a flexible, high-level description language as well as computationally efficient operations that implement the language.

Similarly, reasoning about sequences of inter-dependent actions and about situations that may change at arbitrary times are important aspects of problem solving in space.

Very Large Knowledge Bases

To date expert systems have used knowledge bases of modest size. With the complexity of operations in space, we need to design and maintain expert systems with very large knowledge bases. Although size is difficult to define, most knowledge based mention only a few thousand different facts and relations. Probably the largest today is the

INTERNIST system in which about 250,000 facts are encoded (Miller et al., 1982). Some of this limit results from our own inability to keep in mind the interrelationships among more facts as much from the technology of storing and retrieving them. We must improve the technology to make it easier to build and maintain knowledge bases of much larger scale, which will be necessary in a system as large and complex as the space station.

Shared Knowledge Bases

Today's systems use single knowledge bases that have been built specially for them. As more and more systems are constructed, however, it will be important to use knowledge bases in different contexts and then reuse one system's knowledge base in another system. It is wasteful--and should not be necessary--to duplicate the contents of an old knowledge base in a new application. One should expect, for example, programs in the space station that reason about the function of life support equipment and others that reason about the mechanical structure of the same equipment, both of which must share considerable detail about the equipment itself.

Distributed Databases

Databases exist now on many machines. Yet it is nearly impossible to treat several of them as if they were one logical unit--from any program. Expert systems also need this capability. Current research will allow much broader sharing of data among different databases than is currently available in commercial systems. There will be many computers in the space station. It is much sounder to think of separate specialized databases (with appropriate backup) that can be accessed from various programs than to consider separate copies of every data base on every machine.

Parallel Computation

Computers are fast, but never fast enough. In addition to the immense speed-ups from improvements in the hardware, there are potential speed-ups from software. When a problem can be divided into nearly independent subproblems, it is conceptually easy to see that multiple computers could be used to solve the subproblems in parallel, thus saving considerable time. Work in the research laboratories indicates that this is feasible. Thus it will almost certainly become a commercial reality in the near future if it is cost-effective.

DEVELOPMENT AND MAINTENANCE ISSUES

Building an expert system requires finding out how an expert solves a problem and translating that expertise into a stylized form that can be read by computer. This is no different in principle from building a conventional program in which programmers find out what equations or algorithms experts use and then write FORTRAN or COBOL programs that embody those procedures. The main difference in practice is that expert systems must incorporate knowledge that is much more qualitative and judgmental. In fact, much of the time the expert's "know-how" is not yet written down and what he/she does is regarded as an art.

Because the expert's knowledge is often not already codified and because writing symbolic reasoning programs is itself often regarded as an art, building an expert system requires patience. It generally works best as a team effort involving one or more experts and one or more so-called knowledge engineers. A knowledge engineer is a programmer of knowledge-based systems who understands the conventions of the computing framework and who assists the expert in mapping judgmental knowledge into that framework. The dialogue between expert and knowledge engineer is often called "knowledge engineering".

One of the key ideas in knowledge engineering is to focus on case studies. It is much easier for any of us to tell someone how we would approach a specific situation than to say in general terms how we solve problems of a type. Of course, if we have a set method (sometimes called a "canned procedure") that we always use, we can say that. "Oh yes, I always use the French variation of the Alekhine-Gorbachev wave theory in situations like that", you might say. But then the knowledge engineer wants to know what do you do next and—more interestingly—when would you make exceptions to your set policy. And the best way for you to think about those things is to focus on cases. As long problem solving requires more than the application of set procedures, knowledge engineers will need to go through many cases, and variations on them, to help codify the expert's judgemental expertise.

Steps Involved in Knowledge Engineering

It may take months or years to build an expert systems, with the time depending largely on the complexity of the problem and the extent to which expertise is already codified. One reason it takes so long is that there are many steps involved. And at each step, the knowledge engineer or the expert may decide it is necessary to undo some results of previous steps. Very roughly, the steps are thought of as beginning, middle and end phases in which attention is focused on different aspects of the system, as shown below:

- Beginning—define the problem precisely; understand which concepts are used, what their definitions and inter-relationships are.

- Middle--implement a substantial prototype after choosing a set of representation conventions and writing a small but substantive knowledge base.
- End--fill out the knowledge base to fix errors and extend the scope of the system's problem solving abilities, both of which are generally discovered by testing the systems on many test cases.

Tools to Aid in the Construction of Expert Systems

Just as carpenters can construct houses faster with the right tools, knowledge engineers can build expert systems faster with software tools that boost their productivity. These come in several forms. The main idea, however, is to provide programmers with mechanized intelligent assistants that know about programming convention (including abbreviations and shortcuts), that can help locate and fix errors, that can display the contents and interrelationships in a program or knowledge base, and so forth. These are the kinds of extra capabilities that distinguish system-building environments from programming languages.

Some of the more powerful environments--sometimes called shells--are shown below. One characteristic of a shell is its commitment to a set of representation conventions of the sort outlined previously. See Table 3.

Learning

At present, expert systems do not learn from experience. This is a defect that many research groups are working to remedy. Early prototypes of learning systems promise some automated assistance in maintaining and extending a knowledge base through the experience of routine use, but these are not yet available.

TABLE 3 Some Commercially Available Shells for Building Expert Systems

S.1	Teknowledge
KEE	Intellicorp
Knowledge-Craft	Carnegie Group
ART	Inference Corp.
LOOPS	Xerox
Personal Consultant	Texas Instruments
M.1	Teknowledge
ESDE	IEM

It is possible, however, to learn an initial set of rules from a case library (collected past experience) and use it for classification problems. Induction programs are being used to build simple rule sets for expert systems in which there is little chaining of the rules and little use of uncertain inferences. These are largely marketed in Great Britain where it is better understood that even simple problems may carry great economic leverage. Current research is extending the scope of induction programs to more complex rule sets.

Resources Required

The major cost involved in building an expert system is in personnel time. Shell systems now run on most common computers, so it is not necessary to buy new equipment and, most importantly, it is not necessary to build the complete set of programming tools found in a shell.

Purchasing the shell and some training in how to use it are recommended. The amount of time needed from a team of experts and knowledge engineers is variable--as are their salaries. Table 4 gives some estimates for a hypothetical small system constructed within an existing shell.

It is assumed here that a problem has been precisely defined before beginning, that a case library of at least a half dozen typical and hard cases has been assembled, that a commercial shell has been purchased and runs on an available computer, and that the senior knowledge engineer is very familiar with both the shell and the computer. It is also assumed that the team's primary responsibility is this activity, and that they have the blessing of their management.

In this simple model, the senior knowledge engineer also fills the role of project leader, with as much as half his/her time filled with reports, briefings, budgeting, and other managerial responsibilities. The junior knowledge engineer in this model is responsible for software

TABLE 4 Percent Time Per Quarter

	Q1	Q2	Q3	Q4
expert	75	75	100	75
sr.KE	100	100	100	100
jr.KE	100	100	100	100

NOTE: Approximate percentage of time required from an expert, a senior knowledge engineer, and a junior knowledge engineer to build a hypothetical small system over four quarters of a year. The two main variables in determining the amount of time required are the nature of the problem and the definition of the deliverable.

engineering--that is, integration of the expert system into the run-time environment--as well as for help in building the knowledge base. And the expert, here, is (atypically) also filling the role of "management champion" with some time devoted to securing resources to make the project happen.

One of the main factors that determines the length of time a project will take is, not surprisingly, the nature of the problem. This includes both the scope of the problem and the extent to which a commercially available shell is appropriate for the problem. Another main factor is the definition of the "deliverable", that is the terms of the contractual agreement specifying whether the product delivered is a prototype or is a smoothly polished software package.

There are added gains in building an expert system that offset some of the costs just mentioned. Besides the obvious gains showing up in work performed, there are very noticeable gains in the quality of information available.

Shortening the time required to build systems and increasing our ability to maintain them are thus two of the central issues for putting expert systems in the space station.

ENVIRONMENTAL ISSUES

The environments in which expert systems currently operate are closely constrained. While there is wide variation in the degree of autonomy exhibited across all working systems, most systems in place are interactive, requiring intelligent input from humans. The predominant model of interaction is a consultation model in which an expert systems asks a person for the facts (and interpretations of them) and then provides some advice. A consultation with MYCIN about a medical case is shown in the Appendix.

There are several reasons why the consultation model is appealing, each of which constitutes an opportunity for research. In the first place, a program that asks short-answer questions of a person can finesse the very large problem of understanding free-form English sentences and phrases. The program knows what answers are reasonable in the current context and can have advance expectations about the ways these answers may be framed.

Second, the consultation model provides a strong sense of context which not only helps the program understand a person's answers, but helps the person understand the sense of the questions. This is important because misinterpretation of the program's questions can have serious consequence.

Third, in a consultation it is reasonable to make strong assumptions about the users of an expert system--what they know, what they don't know, what vocabulary they use, what environment they are working in, and so forth. This helps minimize problems in communication. This means also that so-called "common sense" knowledge may be supplied by users and need not all be supplied by the program.

Real Time Monitoring

As expert systems become faster, it will be easier to build systems that monitor other devices or processes with rapid changes. Conceptually a difficult problem is managing time-dependent relations efficiently, which is one of the necessary components of a monitoring system. The large amounts of data received and the speed with which they are received are also critical issues. Integrating AI methods of reasoning about the data with numerical methods for digitizing and filtering is essential.

Richer Input/Output

No one likes to interact with computers by typing. Considerable work on interactive graphics has reduced the need for typing. But it will be even easier when we can communicate with programs by giving voice commands and receiving spoken English output in return.

Models of Users and Situations

No single style of interaction is best for all users at all times. Specialists do not need explanations of the meanings of terms, for example, while less experienced users used considerable help understanding the context of the problem. Also, the criticality of the situation may demand taking shortcuts in data acquisition or reasoning to reduce the risk immediately before taking a more systematic, detailed look at the problem. Expert systems must be sensitive to models of both the user and the situation in order to request appropriate input, reason at an appropriate level of detail, and present conclusions and suggestions in an appropriate way.

CONCLUDING OBSERVATIONS

Expert systems already are saving organizations millions of dollars and performing tasks routinely that ordinarily require human expertise. The number of applications of today's technology is nearly boundless--consider, for example, the number of pieces of equipment in a space station that we don't readily know how to fix. The first commercial shells on the market are robust enough to be used effectively. Integrating intelligent systems with conventional computer programs and with persons in the space station involves new research in many dimensions. The single biggest advantage of AI programs, amply demonstrated in expert systems, is their flexibility. This matches precisely the single biggest design requirement on software in the space station.

What we see now is just the beginning of a wave of intelligent software that can have as great an effect as business data processing software. It is impossible in any area of technology to make accurate

predictions. However, there are many parallels between the growth of expert systems and of computing hardware, with about a 25-30 year lag. When electronic computers became available commercially, businessmen began to ask about applications that would make a difference to them. In 1955, several of these innovators assembled at Harvard to discuss their experiences. Some of the conclusions they drew from their early experience are summarized below (Sheehan, 1955):

1. "The initial overenthusiasm, which inevitably accompanies a project of this scope, can and does make the job harder. Too many people had the impression that this was the answer to all problems. Perhaps it is, but we haven't been smart enough to develop all of them..."
2. "Some of our original thinking has been partly confirmed in that the greatest benefits to be derived from a computer will probably consist of information impossible to obtain previously..."
3. "Our experience has shown that the computer is more adaptable to some projects than others..."
4. "Programmers should be recruited within your own company...It is easier to teach men the required computer and program techniques than to acquaint them properly with the complex procedures and routines of modern-day industry..."
5. "I doubt if it is possible to overemphasize the desirability of providing for convenient corrections or deletion of errors in data..."
6. "The maximum justifiable amount of flexibility for extending or integrating applications must be included in the initial programming..."

The parallels with expert systems are clear: we have just begun to explore the first set of applications and have only a limited view of possibilities in the future. Over the next 5-10 years, however, we do know what extensions in capabilities will be possible because they are the ones that are active research areas in the present.

Future developments will make the technology more broadly applicable and easier to use. Every professional activity--from medicine to manufacturing--contains many problem solving tasks that are not mathematical. The development of artificial intelligence methods that can solve many of these kinds of problems will change the nature of applications of computers in our society as much as the development of transistors changed the architecture and applicability of computers themselves.

Appendix: Annotated Typescript of a Sample MYCIN Consultation

A.1 Instructions for Users at the Start of an Infectious Disease Consultation (Optional)

[Here, as elsewhere, the responses typed by the user follow a double-asterisk prompt]

MYCIN 3-Jun-78

Special options (type ? for help)

**

Instructions? (Y or N)

** YES

This is a computer program named MYCIN that has been designed to advise you regarding an appropriate choice of infectious disease therapy. It is understood that you have a patient who may have an infection. Please answer the following questions, terminating each response with RETURN. To correct typing errors, use the DELETE key to delete single characters, <ctrl>W to delete a word, and <ctrl>Q to delete the whole line.

If you are not certain of your answer, you may modify the response by inserting a certainty factor (a number from 1 to 10) in parentheses after your response. Absolute certainty (10) is assumed for every unmodified answer. It is likely that some of the following questions cannot be answered with certainty.

You may change an answer to a previous question in two ways. If the program is waiting for a response from you (that is, has typed **), enter CHANGE followed by the number(s) of the question(s) whose answers will be altered. You may also change a previous answer at any time (even when the program is not waiting for a response from you) by typing <ctrl>F (Fix), which will cause the program to interrupt its computation and ask what you want to change. (If the response to <ctrl>F is not immediate, try typing the RETURN key in addition.) Try to avoid going back because the process requires reconsidering the patient from the beginning and therefore may be slow.

Note that you may also enter UNK (for UNKnown) if you do not know the answer to a question. ? if you wish to see a more precise definition of the question or some examples of recognized responses. ?? if you want to see all recognized responses, the word RULE if you would like to see the decision rule which has generated the question being asked, the word WHY if you would like to see a more detailed explanation of the question, or the letters QA if you would like to interrupt the consultation in order to ask questions regarding the decisions made so far in the consultation. If you are ever puzzled about what options are available to you during a consultation, enter the word HELP and a list of options will be listed for you.

Sample Response [user input follows the **]

Does the patient have a risk factor for tuberculosis?

**?

Appendix: Annotated Typescript of a Sample MYCIN Consultation

One or more of the following are considered risk factors for TB a) positive PPD (5TU) b) history of close contact with a person having active TB c) household member with a past history of active TB. d) chest X-ray showing apical scarring. e) granulomas seen on biopsy of any organ tissue

expected responses are YES NO
Enter HELP for user options
** YES

SUMMARY

(type ctrl-O to abort printout)

UNK - Answer not known
? - Rephrases the question and gives examples of recognized responses
?? - Prints a list of all recognized responses
RULE - Prints the current decision rule
QA - Program enters question-answering mode
CHANGE # - Go back and re-request answer to question number #
COMMENT - Enter comments about any aspect of the system's performance Your comments will be forwarded to those in charge of the MYCIN program
WHY - Gives high-level explanation of the current reasoning chain that provoked this question
HOW # - Explains HOW the system will achieve a goal referred to by number # in a previous explanation
EXPLAIN - Provides a more detailed explanation of a previous answer given by a WHY command
FORGET - Resets the explanation of the reasoning chain back to lowest level, as if you never used the WHY:EXPLAIN commands
STOP - Halts the program without completing the consultation
SAVE - Halts the program, saving the current patient on a disk file, retrievable at a later date
HELP - Prints this list

A.2 Infectious Disease Consultation Session

[This is a typescript of a MYCIN consultation on an actual patient treated at Stanford. MYCIN was not used to suggest treatment. Spaces have been added for readability, and comments (in italics) have been inserted. Otherwise it is a verbatim transcript. User's responses follow double asterisks.]

-----PATIENT-538-----

1) Patient's name

** PT538

[Name, age, sex, and race are always asked because they are INITIALDATA.]

2) Age

** 34 YEARS

3) Sex

** MAEL

= MALE

[MYCIN corrects misspellings.]

4) Race

** WHY

[The user asks for an explanation. In this instance it is trivial.]

[i.e. WHY is it important to determine the race of PATIENT-538?]

[1 0] The race of a patient is one of the 5 standard parameters that identify a patient.

This may be useful later in the consultation. For example, it is relevant in determining the organisms (other than those seen on cultures or smears) which might be causing the infection.

[back to question 4]

** CAUCASIAN

[This is the answer to the question about race.]

Please supply one answer under each heading. If your response is not

Appendix: Annotated Typescript of a Sample MYCIN Consultation

- 16) Please give the date on which clinical evidence (symptoms, signs, or laboratory tests) of the meningitis first appeared
 ** 29-Jan-77 22:15
- 17) Does Pt538 have an abnormal chest X-ray?
 ** NO
- 18) Does Pt538 have one or more of the following risk factors for TB a) positive intermediate strength PPD, b) history of close contact with a person with active TB, c) a household member with a past history of active TB, d) apical scarring on chest X-ray, e) history of granulomas on biopsy of liver, lymph node, or other organs?
 ** NO
- 19) Does Pt538 have any cutaneous lesions or rash on physical examination?
 ** NO
- 20) Does Pt538 have evidence of ocular nerve dysfunction?
 ** NO
- 21) Is Pt538 a compromised host (e.g. alcoholic, sickle-cell-disease, immunosuppressed)?
 ** YES
- 22) Is Pt538 immunosuppressed (e.g. corticosteroid therapy, cytotoxic drug therapy, radiation therapy, leukemia, lymphoma)?
 ** NO
- 23) For how many days has Pt538 had abnormal neurologic signs?
 ** 72 HOURS *[Note that the answer was given in hours
 MYCIN has to be flexible enough to convert units]*
- 24) For how many days has Pt538 had neurologic symptoms?
 ** 1 DAYS
- 25) Has Pt538 had recent evidence of tender or enlarged salivary glands?
 ** NO
- 26) Has Pt538 been exposed to any contagious disease recently (e.g. meningococcal disease, mumps)?
 ** NO
- Please enter CSF findings in the following table
- | CSF
WBC | Percent
PMN | Protein | CSF
Glucose | Peripheral
Glucose |
|------------|----------------|---------|----------------|-----------------------|
| 27)** 100 | 99 | 280 | 5 | UNKNOWN |
- [Note that some information is missing]*
- 28) Do you know the results of a recent CBC of Pt538?
 ** YES
- Please enter CBC results in the following table
- | WBC
in 1000s | Percent
PMN | Percent
BANDS |
|-----------------|----------------|------------------|
| 29)** 1.9 | UNKNOWN | UNKNOWN |
- 30) Has Pt538 had any injury or insult to or defect in the central-nervous-system (excluding neurosurgery)?
 ** NO
- 31) Has Pt538 ever undergone any type of surgery?
 ** YES
- 32) Does Pt538 have a history of NEUROSURGERY?
 ** NO
- 33) Does Pt538 live in a crowded environment?
 ** NO
- 34) Do you suspect recent alcoholic history in Pt538?
 ** NO
- 35) Is meningitis a hospital-acquired infection?
 ** WHY

[This is a more interesting question requiring presentation of the rule that caused question 35 to be asked. Note that the first five clauses of the rule have been investigated and question 35 is motivated by the last clause. MYCIN first rephrases the question to give the user its interpretation of the WHY.]

Infectious Disease Consultation Session

{i.e. WHY is it important to determine whether the infection was acquired while the patient was hospitalized?}

[2 0] This will aid in determining the organisms (other than those seen on cultures or smears) which might be causing the infection
It has already been established that

[2 1] the infection which requires therapy is meningitis, and

[2 2] organisms were not seen on the stain of this pending CSF culture and

[2 3] the type of the infection is bacterial, and

[2.4] neurosurgery is not one of the relevant items from the history of Pt538, and

[2 5] the age of Pt538 is greater than 10 days

Therefore, if

[2 6] the infection was acquired while the patient was hospitalized then

there is evidence that the organisms (other than those seen on cultures or smears) which might be causing the infection is e coli (.75) staphylococcus-coag-pos (3) pseudomonas-aeruginosa (3) klebsiella-pneumoniae (5)

Author: Yu

[RULE545]

[back to question 35]

*[MYCIN now requests the answer to question 35 again
Instead of answering, the user asks how clause 2.3 was confirmed]*

** HOW 2 3

{i.e. HOW was it established that the type of the infection is bacterial?}

[Rather than printing the contents of the rules, MYCIN shows the rule numbers and the strength of positive and negative evidence from each rule]

The following rules concluded about:

the likelihood that the type of the infection is bacterial

	Rule	Yes	No
[2 7]	RULE148	(19)	
[2 8]	RULE500	(51)	
[2 9]	RULE501	(95)	
[3 0]	RULE502	(97)	
[3 1]	RULE526	(98)	
[3.2]	RULE504	(97)	
[3.3]	RULE524	(98)	

Thus, it has been established that the type of the infection is bacterial

[back to question 35...]

** HOW 2.5

{i.e. HOW was it established that the age of pt538 is greater than 10 days?}

In answer to question 2 you said that the age of pt538 is 34 years

Thus, it has been established that the age of pt538 is greater than 10 days

[back to question 35. ..]

** Y

36) Is Pt538's clinical history consistent with EPIGLOTTITIS?

** NO

37) Is Pt538's clinical history consistent with OTITIS-MEDIA?

** NG

38) Has Pt538 ever undergone splenectomy?

** NO

39) Is Pt538 a burn patient?

** WHY

[Another explanation is requested]

Appendix: Annotated Typescript of a Sample MYCIN Consultation

[i.e. WHY is it important to determine whether PATIENT-538 has been seriously burned?]

[4.0] This will aid in determining the organisms (other than those seen on cultures or smears) which might be causing the infection
It has already been established that
[4.1] the infection which requires therapy is meningitis, and
[4.2] organisms were not seen on the stain of this pending CSF culture.

and
[4.3] the type of the infection is bacterial

Therefore, if

[4.4] pt538 has been seriously burned

then

there is suggestive evidence (.5) that pseudomonas-aeruginosa is one of the organisms (other than those seen on cultures or smears) which might be causing the infection

Author, Yu

[RULE578]

[back to question 39]

** WHY

[In this case, the user asks another WHY question, backing up the subgoal tree]

[i.e. WHY is it important to determine the organisms (other than those seen on cultures or smears) which might be causing the infection?]

[5.0] ... in order to determine the therapeutic regimen of PATIENT-538

if

[5.1] information has been gathered about organisms isolated from the patient, organisms noted on smears of the patient, negative cultures of the patient, suspected infections without microbiological evidence, current drugs of the patient and prior drugs of the patient, and

[5.2] an attempt has been made to deduce the organisms which require therapy, and

[5.3] You have given consideration to organisms (other than those seen in cultures or smears) which might be present, and found at least one organism which requires therapy.

then

determine the best therapy recommendation from among the drugs likely to be effective against the significant organisms

[RULE092]

[back to question 39]

** Y

[MYCIN now prints its intermediate conclusions about the likely causes of the infection]

I have found evidence (RULE545 RULE557) that treatment should cover for E coli causing meningitis

I have found evidence (RULE545 RULE557 RULE578) that treatment should cover for Pseudomonas-aeruginosa causing meningitis

I have found evidence (RULE545 RULE557) that treatment should cover for Klebsiella-pneumoniae causing meningitis

I have found evidence (RULE507) that treatment should cover for Diplococcus-pneumoniae causing meningitis

[Now MYCIN begins to ask questions related to appropriate therapy]

40) Does Pt538 have focal neurological signs?

** UNKNOWN

[Some questions are motivated by the organism under consideration. Thus the following lines sometimes provide useful contextual information to a user]

[Determining which drugs are desirable for use against the E coli]

[Determining which drugs are desirable for use against the Pseudomonas-aeruginosa]

Infectious Disease Consultation Session

[Determining which drugs are desirable for use against the Klebsiella-pneumoniae]

[Determining which drugs are desirable for use against the Diplococcus-pneumoniae.]

My therapy recommendations will be designed to treat for organisms that either are very likely or, although less likely, would have a significant effect on therapy selection if they were present. It is important to cover for the following probable infection(s) and associated organism(s):

INFECTION-1 is MENINGITIS

- ITEM-1 - E. COLI [from clinical evidence only]
- ITEM-2 - PSEUDOMONAS-AERUGINOSA [from clinical evidence only]
- ITEM-3 - KLEBSIELLA-PNEUMONIAE [from clinical evidence only]
- ITEM-4 - DIPLOCOCCUS-PNEUMONIAE [from clinical evidence only]

41) Does Pt538 have a clinically significant allergic reaction to any antimicrobial agent?

** NO

[Considering AMPICILLIN-AND-GENTAMICIN for use against INFECTION-1]

42) Do you have reason to suspect that Pt538 may have impaired renal function?

** NO

43) Pt538's weight in kilograms (or <number> POUNDS):

** 70

[Finally, MYCIN prints the best therapy it has found to maximize the chances of effective coverage of the cause(s) of the infection while minimizing the number of drugs, their toxicity, and side effects.]

[REC-1] My preferred therapy recommendation is as follows

In order to cover for Items <1 2 3 4>.

Give the following in combination.

1) AMPICILLIN

Dose: 3.5g (28.0 ml) q4h IV [calculated on basis of 50 mg/kg]

2) GENTAMICIN

Dose: 119 mg (3.0 ml, 80mg/2ml ampule) q8h IV [calculated on basis of 1.7 mg/kg plus consider giving 5 mg q24h Intrathecal]

Comments: Monitor serum concentrations

Since high concentrations of penicillins can inactivate aminoglycosides, do not mix these two antibiotics in the same IV bottle.

Do you wish to see the next choice therapy?

** YES

[Other therapies can be examined.]

Which, if any, of the above drugs do you wish eliminated from consideration?

** AMPICILLIN

[Considering GENTAMICIN for use against INFECTION-1]

[Considering PENICILLIN for use against INFECTION-1]

[REC-2] Next best therapy recommendation

In order to cover for Items <1 2 3>

Give: GENTAMICIN

Dose: 119 mg (3.0 ml, 80mg/2ml ampule) q8h IV [calculated on basis of 1.7 mg/kg] plus consider giving 5 mg q24h Intrathecal

Comments: Monitor serum concentrations

In order to cover for Item <4>:

Give: PENICILLIN

Dose: 3.5 million units (7.0 ml) q4h IV [calculated on basis of 50000 units/kg]

Infectious Disease Consultation Session

understood. you will be prompted by the standard questioning method
(enter ? for normal prompting, * for questions that are not applicable)

Please enter the results of positive cultures using the following table

site	culture number	date collected	date positive	method of collection [urine,sputum only]
------	-------------------	-------------------	-------------------	---

5)**DONE

*[The next question asks for several items in tabular form
since they are found together in the hospital record]*

Please enter the results of pending cultures in the following table.

site	culture number	date collected	date stain examined	method of collection [urine,sputum only]
------	-------------------	-------------------	-----------------------------	---

6)** CSF 1234 30-Jan-77 YES *

stain of smear	morph	[MYCIN follows up on the stain of the smear and the morphology.]
-------------------	-------	---

[type DONE if no organism was seen on the smear]

7)**DONE

*[No information about stain or morphology is available. MYCIN now
follows up question 6, asking about other positive cultures.]*

site	cul #	date coll.	smear examined?	coll. method
------	-------	------------	--------------------	--------------

8)**DONE

[There are no more positive cultures.]

Please enter negative cultures in the following table.

site	culture number	date collected	method of collection [urine,sputum only]
------	-------------------	--------------------	--

9)**DONE

[There are no negative cultures.]

10) Do you suspect Pt538 may have an infection at a site from which you
have not obtained culture specimens?

** NO

Please enter information on current antimicrobials

drug name	route of administration	date started
--------------	----------------------------	------------------

11)** CEPHALOTHIN

IV

23-Jan-77

12)**DONE

Please enter information on prior antimicrobials

drug name	route of administration	date started	date stopped
--------------	----------------------------	------------------	------------------

13)**DONE

[Considering organisms (other than those isolated from cultures) which might
be present...]

14) Has Pt538 recently had symptoms of persistent headache or other
abnormal neurologic symptoms (dizziness, lethargy, etc.)?

** YES

15) Has Pt538 recently had objective evidence of abnormal neurologic signs
(nuchal rigidity, coma, seizures, etc) documented by physician
observation or examination?

** YES

The CSF cultures will be considered to be associated with MENINGITIS

Appendix: Annotated Typescript of a Sample MYCIN Consultation

RECOMMENDATION-2 is equivalent to RECOMMENDATION-1

Would you like to see details of this comparison?

** YES

Looking at the coverage of the 'most likely' items, we see that:

RECOMMENDATION-1 consists of 1 first choice drug (AMPICILLIN for ITEM-1 and ITEM-4) and 1 third choice drug (GENTAMICIN for ITEM-1, ITEM-2 and ITEM-3)

RECOMMENDATION-2 consists of 1 first choice drug (PENICILLIN for ITEM-4) and 1 third choice drug (GENTAMICIN for ITEM-2, ITEM-1 and ITEM-3)

ORGANISMS	RECOMMENDATION-1	RECOMMENDATION-2
'most likely'	Drug -- Choice Ranking	Drug -- Choice Ranking
ITEM-1	AMPICILLIN-AND-GENTAMICIN--1st	GENTAMICIN--3rd
ITEM-2	GENTAMICIN--2nd	GENTAMICIN--2nd
ITEM-3	GENTAMICIN--3rd	GENTAMICIN--3rd
ITEM-4	AMPICILLIN--1st	PENICILLIN--1st

(The desirability of a drug is defined to be its lowest ranking for the items it covers)

Do you wish to see the next choice therapy?

** NO

REFERENCES

- Anderson, J. L. and A. Chambers.
1985 Human-Centered Space Station Design. Paper presented at Conference on Computers in Aerospace, Long Beach, CA, October 21-23.
- Bauman, B.
1984 MBA Thesis, University of Texas, December.
- Buchanan, B. G.
1986 Expert systems: working systems and the research literature. Expert Systems 3(1):32-51. Also appeared as Stanford KSL Report 85-37 and STAN-CS-85-1075.
- Harmon, Paul.
1986 Inventory and analysis of existing expert systems. Expert Systems Strategies 2(8):1-16.
- Miller, R. A., Pople, H. E. and Myers, J. D.
1982 INTERNIST-1, An Experimental Computer-Based Diagnostic Consultant for General Internal Medicine. NEJM 307(8):468-476. Also appears in 'Readings in Medical Artificial Intelligence: The First Decade', Clancey, W. J. and Shortliffe, E. H., eds., Reading, MA, Addison-Wesley, Pp.190-209.
- Sheehan, G. M
1955 An application to payroll. Proceedings of the Automatic Data Processing Conference. Harvard University Press, Cambridge, MA, September.

DISCUSSION: COMMENTS ON EXPERT SYSTEMS AND THEIR USE

Allen Newell

Bruce Buchanan gave us a broad view of expert systems and showed a rather large collection of aspects across the whole field that need to be worried about to make the advances NASA needs. This leads to a point I want to make, which concerns my own concern about whether research is really needed on some parts of expert systems.

As preparation, Figure 1 shows my current favorite diagram to explain AI. You need to understand about AI that there are two dimensions in terms of which to talk about the performance of systems. The first is the amount of immediate knowledge that they have stored up, that they can get access to. This can conveniently be measured by the number of rules. The second is the amount of knowledge that they obtain by exploring the problem. This can conveniently be measured by the number of situations examined before committing to a response. Thus, there are isobars of equal performance, with better performance increasing up towards the northeast. You can roughly locate different intelligent systems in this space. Expert systems are well up on the immediate-knowledge scale, without much search. The Hitech chess program, which has a little, but not very much knowledge, lies far out on the search dimension. The human being is substantially above the expert systems on the knowledge dimension. Also, most expert systems do less search than humans do. The whole point of this diagram is that, in the current era, expert systems are an attempt to explore what can be achieved without very much search and reasoning, but with a modest amount of immediately available knowledge.

If you accept the characterization of expert systems in the figure, then even without all the research that Bruce was talking about, there exists an interesting class of programs, even though it is very limited in capability. The expert systems of today constitute a class of programs that appears to be very useful if you limit the tasks to the right kinds. Bruce was helping to characterize that. We actually know a modest amount about this type of task. If you have the right knowledge assembled, then you know what to do and how to do it without very much involved reasoning. For such tasks and their expert systems, it is not clear that the big need is to do a lot more research. The big issue is to build lots of these systems for lots of these tasks. What is needed is more like a development effort, to find out which tasks can successfully be done with modest amounts of expertise. The

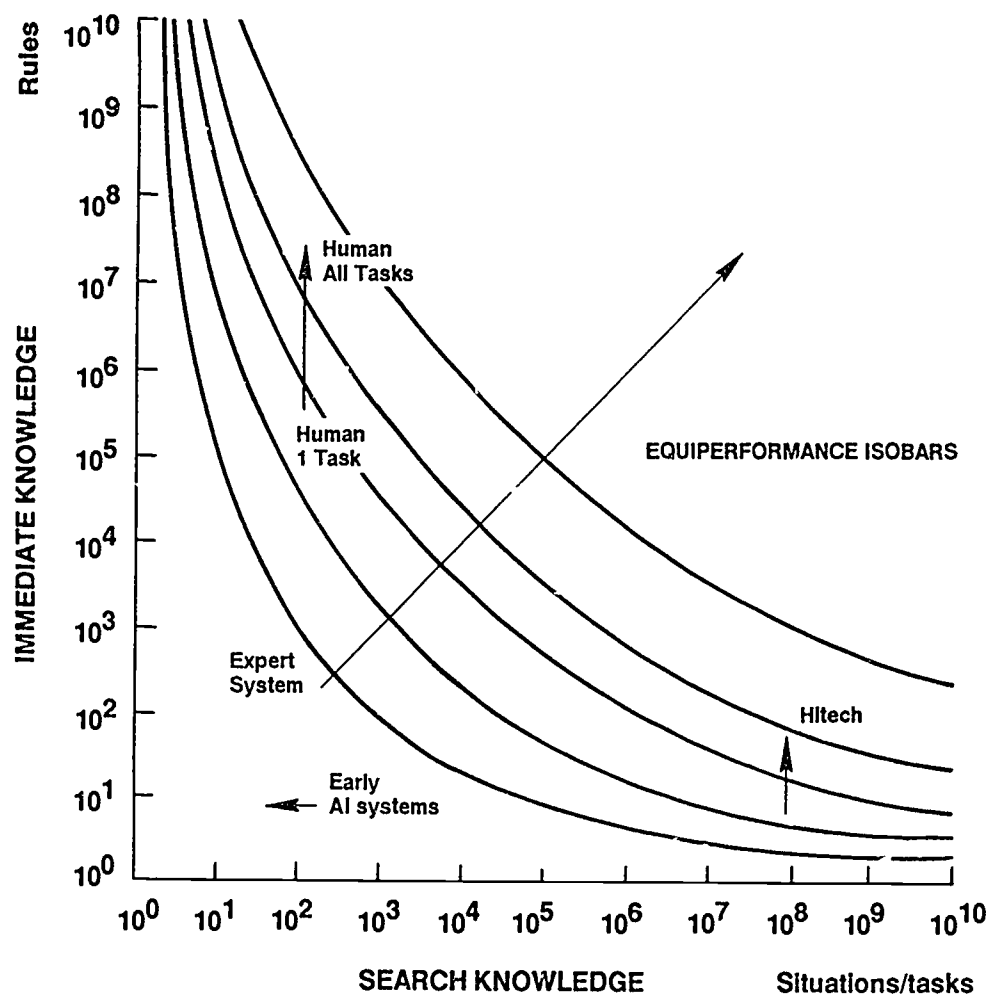


FIGURE 1 Immediate knowledge versus search knowledge trade-off

need is not to build any more expert-system shells, or to build more tools. The need is to pour all of the effort into finding out, in the plethora of space-station tasks, which are the ones that the current level of technology really does provide interesting and useful solutions.

Tom Mitchell talked much more specifically than did Bruce about the fact that the space station is a physical system—that if you want to use expert systems and AI systems, they had better interact directly with physical devices. I agree absolutely that this is a major issue and a very important one for NASA to research. In particular, bringing control theory and symbolic reasoning together so we understand those as a single field is important. What I would like to emphasize is how little we know about that. In some respects we do not even know the units to use to talk about it, or how such symbolic programs ought to interact with control systems.

To bring this point home, let me note that a lot of current effort in understanding the human motor system is directed toward exploring a kind of system which is not controlled in detail. A particular dynamic system that has the right properties is composed, and is sent off to do a motor action. A good example is Hollerbach's model of handwriting, in which the whole system is composed of simply-interacting dynamic subsystems, which continuously draw letter-like curves, which are then modulated for specific letters. These dynamic systems are not cast in concrete. They are created and torn down in seconds, in order to compose and recompose dynamically according to short-term task requirements. The motor units that the cognitive system interacts with are these composed dynamic systems. We know almost nothing about such systems. When we finally understand something about it, I suspect it will change our notion entirely of the interface between the symbolic system and the dynamic system. The point is that there is a lot of research before we even get a clear idea clear about how symbolic systems ought to interact with mechanical and dynamic systems.

Tom made a suggestion about emulating devices. If a device breaks, then the emulation can be plugged in. I think this is an intriguing idea and there may be a whole world of interesting research in it. You might counterargue that, if this is possible, then everything might as well be run in computer mode. But there is a real reason not to do that. Making the emulation work may take a lot of computing power. A principal reason for using real physical devices and not simulating everything is that your system runs faster if you do not simulate it. But that does not imply that, if one device breaks, you cannot bring to bear an overwhelming amount of computational capacity to try to compensate for it. Thus, the system is prepared to emulate everywhere, but only has to do it in one or two places on any occasion. Emulation provides a backup capability. In fact, it is never likely to be as good, but at least it will be better than having to shut down the whole system. I think this is an interesting path of research, which could be pursued a long way. In particular, the feature that Tom mentioned about thinking of ways to construct systems so that they are decomposable and emulatable might yield many interesting possibilities.

Tom also raised the issue of sharing responsibility. However, he did not in fact tell us much about how tasks should be shared. Rather he described a particular aspect of the issue, which suggests that the machine ought to learn from the human, and then, quite properly, that the human ought to learn from the machine. I approve of both of these activities, but they beg the whole question of sharing. They do not elaborate ways of sharing, but both spend a fair amount of their time simply learning to be like each other, and confusing who really has the knowledge and who really knows how to do what. In fact, if one has machines with this kind of capability, the entire question of what it means to share may get transformed. It will become extremely difficult to quantify or be precise about who knows what, who ought to do what, and even who is doing what in the space station. There exists a kind of complementarity, in which the more you spread capabilities around in the system, so that there is a lot of redundancy, the less possible will it be to characterize the role of system components

effectively—to say for instance what the separate contributions are to the productivity of the total station. All I want to observe is that such systems are not clean, and learning and performance get confused. However, even though they are not clean, they may turn out to be the kind of system one has to build in order to get the margins of safety that are needed in space.

Finally, I want to talk about the issue of robustness, although it was not a major focus of either speaker. It is a fact, I believe, that there has been essentially no work on making expert systems robust. There is much attention, of course, to their giving explanations. But fundamentally expert systems are collections of rules, which are ultimately brittle and unforgiving. The lack of attention to robustness arises, in part, because there is a market for programs that are not very flexible or very robust. They can nevertheless, be successful. They will be increasingly successful, especially if the problem is turned around by saying 'I've got this hammer; where are interesting things to hit with it?' As a result, the expert systems field is not focused on solving the problem that I think NASA has to get solved, which is that it cannot use expert systems in space unless we understand how to build robust expert systems.

A research program in robust expert systems could be fielded by NASA, and I would certainly recommend it. Given requirements on robustness, one could explore more redundant rule sets or the provision of greater backtracking and reasoning mechanisms. There are many approaches to robustness and reliability that have their analog in expert systems and could provide guidance.

However, I think something more basic is at stake. What is really wrong here is the whole notion of laying down code—or rules, which play the role of code for existing expert systems. That is, as soon as you lay down code, it becomes an echo from the past, unadapted to the future. You have become subject to a mechanism. Code is blind mechanism, complex perhaps, but blind. The important thing about a blind mechanism is that it does not care. A bullet does not care who it kills. A broken beam does not care on whom it falls. The horror stories about non-robust software almost invariably reflect the fact that code was laid down in the past, in a fantasy land of what was going to be, and something different happened at run time, for which the code was not adapted.

The problem, I believe, is that the unit, the line of code, is wrong. A clue for what might be right comes from the database world, with its adoption of transaction processing. It was concluded that the wrong thing to do was to take a line of code to be the unit. What had to be done was to package the specification of behavior in a hardened form called the transaction, for which some guarantees could be made. This has the right flavor of having changed the nature of the unit to make real progress. It has the wrong flavor because the unit is still just a little mechanism. Somehow, in the area of robustness, the smallest unit of action has got to be, if I can use a metaphor, a caring piece of action. It has to be an action, which has a big enough context, even in its smallest unit, to react in terms of the global goals of the system, so it can care about safety and can care about the

consequences of what it is doing. Somehow we have to find out how to create units that have that property. The units cannot be rules or code and so forth, which are just mechanisms. I think NASA ought to go after that. It would be a great research project. It is my contribution to this symposium of a really basic research goal that has an exceedingly small chance of succeeding, but an immense payoff if it does.

SYNOPSIS OF GENERAL AUDIENCE DISCUSSION

Concerns of several varieties were expressed about the knowledge engineering aspects of expert systems. Members of the audience with direct experience with developing expert systems gave these remarks special cogency. Expert systems seem to work better where good extensive formulations of the knowledge base already exist. Attempting to develop that knowledge base as part of the expert system effort often fails. The domains of expert systems are often exceedingly narrow, limited even to the particularity of the individual case. Given the dependence of the knowledge in expert systems upon the informants, there exists a real danger of poor systems if the human experts are full of erroneous and imperfect knowledge. There is no easy way to root out such bad knowledge.

On this last point it was noted that the learning apprentice systems discussed in Mitchell's paper provide some protection. The human experts give advice for the systems to construct explanations of the prior experience, and what the systems learn permanently is only what these explanations support. Thus the explanations operate as a filter on incorrect or incomplete knowledge from the human experts.

Concern was expressed about when one could put trust in expert systems and what was required to validate them. This was seen as a major issue, especially as the communication from the system moved towards a clipped "Yes sir, will do". It was pointed out that the issue has exactly the same complexity with humans and with machines, in terms of the need to accumulate broad-band experience with the system or human on which to finally build up a sense of trust.

Trust and validation are related to robustness in the sense used in Newell's discussion. It was pointed out that one path is to endow such machines with reasoning for validation at the moment of decision or action, when the context is available. This at least provides the right type of guarantee, namely that the system will consider some relevant issues before it acts. To make such an approach work requires providing additional global context to the machines, so the information is available on which to make appropriate checks.

Finally, there was a discussion to clarify the immediate-knowledge vs search diagram that Newell used to describe the nature of expert systems. One can move along an isobar, trading off less immediate-knowledge for more search (moving down and to the right) or,

vice-versa, more immediate-knowledge for less search (moving up and to the left). Or one can move toward systems of increased power (moving up across the isobars) by pumping in sufficient additional knowledge and/or search in some combination. The actual shape of the equal-performance isobars depends on the task domain being covered. They can behave like hyperbolic asymptotes, where further tradeoff is always possible at the cost of more and more knowledge (say) to reduce search by less and less. But task domains can also be absolutely finite, such that systems with zero search are possible, with all correct response simply known. For these, there comes a point when all relevant knowledge is available, and no further addition of knowledge increases performance.

SESSION III

LANGUAGE AND DISPLAYS FOR HUMAN-COMPUTER COMMUNICATION

CHANGE IN HUMAN-COMPUTER INTERFACES ON THE SPACE STATION:
WHY IT NEEDS TO HAPPEN AND HOW TO PLAN FOR IT

Philip J. Hayes

OVERVIEW

The space station is unique in the history of manned space flight in its planned longevity. Never before have we had to deal with a manned space system that was expected to perform for twenty five years or longer. The implications of this requirement are far-reaching. This paper attempts to explore some of these implications in the area of human-computer interfaces.

The need for hooking (designing software for future extension and modification) is already well established¹ in the space station program as a whole. The paper explores in some detail why hooking is an important requirement for human-computer interfaces on the space station. The reasons are centered around the rapid rate of expansion in the kinds and combinations of modalities (typing, graphics, pointing, speech, etc.) available for human-computer interaction and in the interaction and implementation techniques available for them. Many of these modalities and associated interaction techniques are well-developed, others are in embryonic stages. Different modalities (or combinations of modalities) are appropriate to different situations. The paper therefore also looks at the appropriateness of the modalities according to task, user, and the space station environment. An appropriate matching of interface modalities, task, and user is essential to maximizing the potential of on-board computer systems in their primary goal of supporting and amplifying human abilities.

A second rationale for providing hooking in human-computer interfaces is related to the currently developing possibilities for intelligent interfaces. So the paper discusses methods of achieving intelligence in interfaces, and in what circumstances it is desirable. The issue of intelligence is also related to the distinction between conversational/agent type systems and machine/tool-like systems. The current culture at NASA is highly oriented towards the latter. The paper explores the tradeoffs between the two approaches and discusses the circumstances in which a more conversational/agent style system could fit space station goals and NASA culture.

After examining the need for hooking in human-computer interfaces, the paper turns to the question of how to achieve it. The discussion

here centers around methods of achieving a clean separation between the interface and the underlying application (space station system) it interfaces to. The key advantage of this kind of separation is that it allows the interfaces to be changed independently of the application so that a new interface (possibly employing different modalities from the old one) can be rolled in without altering the application in any way. In an environment such as the space station where the underlying applications may be complicated, mission critical, and highly integrated with other applications, such separation becomes all the more important.

The feasibility of a completely clean separation between interface and application is unclear at the moment. The question is currently being addressed by the major subarea of human-computer interaction that deals with user interface management systems (UIMSs). Unfortunately, it is infeasible to wait for research on this topic to reach full maturity. Unless the original applications and interfaces are built with separation in mind, retrofitting separation is likely to be impossible. So the paper discusses what kind of interface/application separation is feasible for the space station initial operating capability (IOC), and looks at how this will constrain the overall possibilities for human-computer interaction.

Separation of interface from application has two other important advantages in addition to hooking. First, it promotes consistency between interfaces to different applications. Most of the work on UIMSs emphasizes a common set of tools for construction of the separated interfaces, and this inevitably leads to considerable consistency of (at least fine-grained) interface behavior between interfaces. The importance of consistency in interfaces has been appropriately emphasized by Polson in the preceding paper. Secondly, the hooking made possible through separation also makes it easier to alter interfaces during their initial development. The only effective way of developing excellent human-computer interfaces is to build interfaces, see how users perform, and then repeatedly alter them to deal with problems. This process is much more effective if the interfaces are easy to modify. The paper explores these two other aspects of interface/application separation further.

APPROPRIATE INTERFACE MODALITIES

The need for change in human-computer interfaces on the space station and the consequent need for hooking arises out of the rapid development that has occurred and continues to occur in interface modalities (typing, graphics, pointing, speech, etc.) and the interaction techniques used with them. This section discusses what interface modalities (or combinations of modalities) and techniques are appropriate for different kinds of interface tasks. An appropriate matching of interface modalities, task, and user is essential to maximizing the potential of on-board computer systems in their primary goal of supporting and amplifying human abilities.

Interface Requirements for the Space Station

The basic considerations in designing good human-computer interfaces for the space station are the same as for any human-computer interface on Earth. In particular, the interfaces should be:

- easy to learn
- easy to use
- efficient to use

Much has been written, e.g. (Hansen, 1971), about this and similar lists of attributes. For present purposes, we can treat them as self-evident, though of different relative importance in different interface situations. There are, however, some special characteristics of the space station environment that require further discussion before looking at the relative utility of the different available interface modalities. These characteristics include:

- **Weightlessness:** In addition to being the most obvious special characteristic of the space station environment, zero-g causes specific problems for human-computer interfaces. The problem is that movement by humans in a weightless environment induces other movement. This is particularly true if the movement involves pressure against another object, such as in typing or pointing on a touch sensitive screen, but it is also true for any kind of gesture, such as with a non-touch light pen. A person employing such interface modalities will tend to drift away from or change orientation with respect to the workstation he is using. The simplest solution to involuntary movement induced by human-computer interaction is simply to tether the user physically to the workstation. This, however, has the obvious disadvantage of inconvenience, especially if the interaction session will not last long. Also, the tethering would have to be relatively complex and therefore intrusive to solve completely the problem of changing orientation.
- **Analogue/continuous interaction:** Many interactions on the space station require (or could benefit from) command input which can be given rapidly and/or in an analogue/continuous manner. Obvious examples include any kind of docking or remote manipulation activity. Less obvious ones include manipulation of continuous variables in, for instance, systems controlling the life-support environment. Analogue/continuous interactions require different kinds of interaction modalities and techniques from those used in more traditional computer command languages.
- **Varied groups of users:** Although the most mission-critical systems will continue to be operated by highly trained personnel, the sheer number of systems likely to be available in the space station suggests that this will not be true for all systems. Some less mission-critical or time-critical systems

in, for instance, the areas of personal comfort, provisioning, or inter-crew communication, are likely to have to interact with users of varying degrees of sophistication and experience with respect to those systems. To avoid negative transfer effects between different systems, interfaces need to be as consistent as possible across the various systems. To deal with users who are inexperienced (for that system), interfaces also need to be as self-evident, self-explanatory, and self-documenting as possible. The goal should be for experience with some subset of the non-mission critical systems and appropriate knowledge of the domain the system deals with to serve as sufficient experience for the accomplishment of straightforward tasks with any of the other non-mission critical systems.

- Hands-free operation: There are many situations in the space station environment in which hands-free interaction would be useful. An obvious example is extra-vehicular activity, but more frequent examples might arise when it was important to avoid the induced motion problems mentioned above (in the weightlessness bullet) or when it was useful to have an additional I/O channel in the context of a complex hands-on analogue activity such as remote manipulation. The most natural hands-free modality is speech, but other possibilities include control through eye-movement, or in specialized circumstances use of feet or other body parts.

Having looked at some of the space factors which might influence choice of interface style and modality, we now look at the appropriateness and range of applicability of the various modalities. Some of the discussion presupposes certain styles of interface for each type of modality. The presuppositions are not always necessarily valid, but are characteristic of the way the modalities have typically been used.

Character-Oriented Interfaces

The vast majority of human-computer interfaces currently in use are character-oriented. The users of these interfaces provide input by typing on a keyboard, and the systems provide output through a screen with a fixed number of character positions (typically 24 lines of 80 characters). Interfaces of this kind do not have a great deal to commend them for the space station environment. Reasons include:

- The physical pushing motion involved of typing leads to the induced motion problem mentioned above. Typing sessions of any length require some kind of tethering arrangement.
- Typed input is unsuitable for analogue/continuous interaction.
- In character-oriented interaction, the user typically issues commands through expressions in a line-oriented artificial

command language). Such languages generally require significant learning effort, making them difficult to use for initial or casual users. Some command languages, such as the one for DEC's Tops-20 operating system, have shown that it is possible through uniformity and carefully thought out help facilities, to reduce the difficulty of use by non-expert users. However, command line interaction is inherently more limited in its perspicuity than the direct manipulation style described in the section titled "Graphically-Oriented Interaction".

- Although some of the learnability and ease of use problems with command-line interaction can be overcome through selection from menus from the keyboard, this can be seen as an attempt to overcome the limitations of the modality by use of an interaction technique borrowed from another modality, i.e. pointing input. It seems more appropriate to use the pointing modality directly.
- Character-oriented interaction is essentially an old, though very well worked out (see e.g. Martin, 1973), technology.

Graphically-Oriented Interaction

A recently developed and increasingly popular style of interaction is based on the use of a high-resolution graphical display and a pointing device such as a mouse or joystick. A well known system exemplifying this scheme is the Macintosh personal computer (Williams, 1984). Interaction in this style is based on techniques such as menu-selection, icon selection and movement, and other kinds of graphically-oriented operations. This style of interaction is also known as direct manipulation (Hutchins et al., 1986; Shneiderman, 1981), indicating ideally that the user should feel that he is directly manipulating the objects represented by the computer system. An example of this kind of direct manipulation analogy is deleting a file by using a mouse to "pick up" the icon representing the file and move it into an icon depicting a wastepaper basket.

There are many interfaces that are graphical in nature, but fall well short of the ideal of direct manipulation of providing the user with the illusion of operating directly on the "world" of the underlying application. Interfaces that rely on menus, for instance, often do not support such an illusion. Interaction will have more of the flavor of direct manipulation if the user can perform an operation by moving an icon, for instance, as in the file deletion example above, than by selecting the name of the operation from a list in a menu. To the extent that they can be maintained, the metaphors implicit in direct manipulation interfaces make the interfaces more easily learnable, and reduce the need for help systems. This is important for the varied groups of users that will be using non-mission-critical systems. The Xerox Star (Smith et al., 1982) and Macintosh (Williams, 1984) have given some idea of what is possible in this line in the

office and personal computing arena. More research is needed to provide more interaction metaphors on which to build direct manipulation interfaces. The creation of such metaphors will be aided by the existence of new and innovative I/O devices (see section titled "Novel I/O Modalities").

Graphically-oriented or direct manipulation interfaces are in many ways superior to character-oriented interfaces for the space station environment, but there are still some deficiencies. In particular, some of the standard pointing devices used on earth are not well adapted to a weightless environment. This is particularly true of the mouse which is intended to be used on a flat surface under the influence of gravity. The lightpen and the tracker ball both require pressure against a surface and so have an induced motion problem. The joystick may be better adapted from the point of view of induced motion since it requires that the user grip it to manipulate it. This raises the possibility that correction of the motion induced might be possible through the user's grip. However, there are obvious problems with this approach for fine-grained movements, but there is a great deal of experience with the use of joysticks in weightless environment from such tasks as remote manipulation.

A better approach may be solved by further development of innovative pointing devices specifically aimed at use in a weightless environment. One possibility is a freely movable hand-held "mouse" which induces 2-D motion on a screen. Of course, the full six degrees of freedom of motion with such a device also open up the possibility of control of three-dimensional simulations or real actions. Devices of this kind are available and investigations into their use and refinement should be encouraged.

Another innovative kind of pointing technology even better adapted for space is eye tracking. Eye tracking has the dual advantages of no significant induced motion and hands-free operation. It has the disadvantage of intrusive apparatus. It may be particularly appropriate for activity in a space suit where the eye-tracking apparatus can be incorporated into the helmet with no increment in discomfort or inconvenience. Further work is needed both to develop less intrusive forms of eye tracking and on the use of eye tracking control in extra-vehicular activity.

Earth-based direct manipulation interfaces generally operate within the context of fixed workstations. While there are many space station tasks for which this is perfectly appropriate, there are others where a more portable arrangement is required or preferable. FVA is the most common, but other examples include inventory, inspection, and communication tasks. Work on in-helmet displays is needed for EVA to complement the work on eye-tracking. Other work on hand-held or otherwise portable display and pointing devices is needed for the on-board tasks requiring mobile interactive devices.

Natural Language Interaction Via Keyboard

Typed natural language input and output is not a modality in its own right, but a variation on character-oriented interaction. However, it is sufficiently different from typical command language interaction that it is worth considering separately.

A low-level, but nevertheless significant, artifact of the redundancy of human language is that natural language will usually require many more keystrokes than a command language designed for a specific interaction task. This means that the remarks above about the undesirability of the significant amounts of typing involved in command language interaction apply with greater strength to typed natural language interaction. Also for rapid interaction or interaction with an expert user, the amount of typing involved typically makes natural language interfaces unacceptably slow.

Natural language interaction, however, has the important advantage over command language interaction that it allows the user to express things in a way that is natural for him, rather than having to learn an artificial (and frequently arcane) command language. It is thus more suitable for casual users and could help to meet the goal of making a wide variety of space station systems accessible to many different users of varying skill levels.

This argument in favor of natural language interaction presupposes that the interfaces can handle any form of expression that a user cares to come up with and is relevant to the underlying application. At the current state-of-the-art, this is an invalid assumption. In practice, natural language interfaces fall well short of full coverage on syntactic, semantic, and pragmatic grounds, even for the restricted domain of discourse implied by a specific underlying application. This leads to the habitability problem (Watt, 1968) in which many of the advantages of naturalness and lack of learning disappear because the user has to learn what is still essentially a subset of English (or whatever natural language is being used) artificially restricted by the limitations of the natural language processing system. This problem can sometimes even make the language more difficult to learn than a simple command language because the limitations are less easy for the user to identify and remember. On the other hand, these problems can be minimized by appropriate human engineering for interfaces to appropriately limited applications. However, this is very time-consuming and expensive at the time the interface is developed since it involves detailed observations of many users interacting with the system and repeated extensions of the natural language coverage until all the commonly occurring syntax, semantics, and pragmatics are handled.

Perhaps the most important reason for not using natural language interaction is that most interaction can be handled more easily by direct manipulation or other graphically-oriented means. Moreover, as the section titled "Graphically-Oriented Interaction" points out, graphical interaction is likely to be more suitable for the space station environment than character-oriented interaction in general. Whenever the user is trying to select between a limited number of

alternatives or is trying to manipulate objects or access information that can be presented to him in an intuitive spatially-distributed manner, then natural language interaction (or any other form of keyboard interaction) is likely to prove inferior to graphical interaction. There are, however, some circumstances in which natural language or command language interaction is preferable to graphical interaction, including:

- When there is a large range of options to choose between, especially when the options can be composed in a combinatorially explosive kind of way;
- When there is no convenient way to distribute the information in a two-dimensional space;
- When a suitable spatial distribution exists, but the resulting space of information is so large that only a small fraction of it can be presented to the user at any one time;
- When the user is looking for information that is distributed across several spatially-distinct items, so that retrieval of the information by direct manipulation would require iterative examination of each of the relevant interface components.

These conditions are not true for most interactive situations, but come up frequently enough for natural language to be considered as a secondary mode of interaction for many applications to supplement a largely direct manipulation interface. To be effective in this role the natural language interaction has to be suitably integrated with the direct manipulation interaction. Some work has been done in this area on how to use visual context to help interpret pronouns and other anaphoric and deictic references by the user and also to allow intermixing of pointing and natural language input (Bolt, 1980; Hayes, 1987a). However, integrated natural language and graphical interfaces could provide significant benefits given an appropriate research effort.

Speech Interaction

Although a combination of typed natural language and graphical interaction offers some attractive advantages, natural language interaction through speech offers many more. While the habitability problems mentioned in the section titled "Natural Language Interaction Via Keyboard" remain, spoken input is much more rapid and natural than typing the same words. Moreover, the voice and ears offer channels of communication quite separate from the hands and eyes. Speech input leaves the hands free and speech output leaves the eyes free for other tasks (either computer interaction or interaction with the physical world).

In terms of suitability for speech interaction, the space station environment has one specific advantage and one specific disadvantage. The advantage is the absence of any need for speaker-independent speech recognition. At the present state-of-the-art in speech processing,

considerably better results can be obtained if the speech recognition system has been trained in advance on the specific characteristics of a speaker's voice (through recordings of the speaker saying a predetermined set of words several times). Given the relatively small number of people that will be on-board the space station at any given time, their relatively long training period, and their relatively long stay, such system training is unlikely to be a problem. The specific disadvantage of the space station environment is the relatively high level of ambient noise that can be expected inside it, at least if the experience of the Shuttle is a guide. Ambient noise is problematic for speech recognition. At the current state-of-the-art, resolving this problem would probably require the use of a close-speaking microphone of some kind. This itself has the disadvantage of being intrusive and inconvenient to take off and put back on.

The current state-of-the-art in speech processing is still fairly limited. In addition to the speaker-dependent and ambient noise limitations mentioned above, the better commercially available systems tend to be able to handle only small vocabularies (less than a thousand words is typical) and pauses between each word or group of words that the system recognizes as a lexical units (so-called connected speech recognition, as opposed to continuous speech recognition in which no pauses are needed). However, this is a field where rapid advances are occurring and new commercial developments plus a very active academic research program are pushing back all of these limitations. In fact, speaker-independent, large (10,000 word plus) vocabulary, continuous speech recognition in noisy environments is likely to be available within the lifetime of the space station, and systems in which a subset of these restrictions have been relaxed are likely in the early part of the space station's lifetime.

Given these prospects for advancement and the inherent advantages of speech interaction, it seems natural for NASA both to plan on a significant role for voice in space station human-computer interfaces and to keep track of or actively support research on speech processing. Nevertheless, even if the underlying speech technology advances as projected above, other problems remain that will require solution before speech can make its full contribution to human-computer interaction on the space station.

First, speech interaction on its own is quite unsuitable for some kinds of interaction, particularly analogue/continuous commands--it would be very difficult to control a remote manipulation device through a series of "left a bit", "down a bit" kinds of commands. Moreover, even in situations where speech could be used, such as the specification of discrete commands in an inventory tracking system, it may not always be the preferred mode of interaction. For instance, if the arguments to a particular command all have relatively complex verbal descriptions, but there are only four of them, it is probably simpler, more mnemonic, and more reliable to let the user input the argument by pointing at a menu or set of icons representing them. Both of these situations indicate the need for techniques for integrating speech interaction with other modalities including pointing and 3-D manipulation. Speech can then be seen as a complementary channel for

issuing discrete commands during continuous/analogue manipulations while both hands are occupied, such as releasing catches during a remote manipulation task. It can also be seen as a supplementary channel for issuing whatever commands or portion of commands are convenient during a discrete command interaction, and as a stand-alone interaction medium for discrete commands whenever hands-free operation is necessary or convenient. Many of the same research issues arise in integrating speech with other modalities as were described in the section titled "Natural Language Interaction Via Keyboard" for the integration of typed natural language and graphical interaction. These issues include resolution of deictic phrases ("this one", "that") and other pronouns, use of the user's visual context in interpreting what he says, and methods of combining input from pointing and speech to form a single command. Although interesting explorations have already been undertaken in this area (Bolt, 1980; Hayes, 1986), these issues all require further research.

In addition to problems of integration with other input modalities, speech interaction raises some interesting problems of its own related to managing the dialogue between human and computer. The first problem concerns when the computer should listen, i.e. when it should try to interpret the speech that its users are producing. The users will speak to other people (or sometimes to themselves) as well as to the machine and attempts by the machine to interpret speech not directed at it is only likely to cause trouble. Techniques that have been explored here include physical switches (typically foot switches on Earth) or switches based on key phrases (such as "listen to me" and "stop listening") that have to be uttered to start and stop the machine trying to interpret speech. These devices are clumsy and detract from the feeling of naturalness that spoken interaction should provide, but will probably be necessary until speech systems become sophisticated enough to make positive determinations that spoken input is not being directed at them. The prospect of such an ability is well beyond the horizon of current research.

Another dialogue issue with special implications for speech is that of ensuring reliable communication. An interactive speech interface must ensure that it understands the user accurately; that the user is confident of this; that the user becomes aware when the system has failed to understand correctly; and that the user is able to correct such errors when they arise. Humans have developed sophisticated conventions (Sacks et al., 1974; Schegloff et al., 1977) for ensuring that communication is indeed robust in this way. Unfortunately, many of these conventions rely on a level of understanding and intelligence that is unrealistic for machines. However, to have smooth conversations, ways must be found to perform the above functions that are both suitable for the limited intelligence of current machines and fit reasonably well with human conventions. A limited amount of work has been done in this area e.g., (Hayes and Reddy, 1983), but much more is needed.

Finally, there is the same problem of habitability that arises for typed natural language interfaces. For speech, however, the problem can be even worse since the user is less well able to be deliberate and

precise in his choice of words and phrasings while speaking than while typing. Moreover, when speech is used as a stand-alone human-computer interaction modality, there is no possibility of reminding the user through a display about the limitations of the domain of discourse or the phrasings that can be used. Work is needed here to find better ways of developing a reasonably habitable subset of a natural language for a restricted domain, to develop ways for the system to encourage the user to stay within the bounds of the restricted language through appropriate output of its own, to devise methods for partial understanding when a user strays outside the bounds of the restricted language, and to develop interaction methods for steering the user back on track when he does stray as he inevitably will.

Novel I/O Modalities

The interaction modalities discussed so far are conventional in the sense that they have already been widely used (this is least true of speech) in earthbound interfaces and other space systems. However, the numerous challenges posed for human-computer interaction by the space station and the recent emergence of some novel and innovative interaction modalities suggest that it is worthwhile also to consider some of these less-developed modalities for use in the space station.

An innovative input modality of potentially considerable utility on the space station is the use of gesture. The conventional use of a mouse or other pointing device in conjunction with a display screen is a limited form of gesture, but it is possible to sense and interpret a much broader range of human gesture by machine. Large scale gestures involving whole limbs are not practical for the space station because of the constraints of a weightless environment, but smaller-scale gestures are quite suitable. The least problematic form of gesture from the point of view of the induced motion problem is eye motion. As already discussed in the section titled "Graphically-Oriented Interaction", eye tracking can be used as a substitute for pointing via a mouse or other conventional pointing device. It is particularly well suited for use with in-helmet displays.

A more radical departure from conventional technology is the interpretation of hand and finger gestures. Technology is emerging that will allow a machine to recognize a full range of small manual gestures made in restricted spatial context. There is a large range of gestures that have associated conventional meanings (such as yes, no, get rid of it, move it from place to place, etc.). This suggests that interfaces that accepted such gestures as input could be very easy and intuitive to learn and natural to use. It might even be possible to resolve any motion problems induced by gesturing through the use of balanced symmetrical gestures which employ two equal and opposite motions.

We have discussed two ways in which gesture can be used in innovative ways for computer input. There may well be others. In general, there is a need for imaginative exploration of the whole range of ways in which human movement compatible with a weightless, noisy

environment can most easily be sensed by machine.

Another potentially promising area for innovation in interaction techniques involves output by means other than fixed screens and simple audio feedback. In-helmet displays hold significant promise in this direction. Although such displays are most natural in circumstances in which the user has to wear a helmet anyway, such as EVA, they can also improve human-computer interaction in other circumstances. Current investigations, including some at NASA-Ames, have shown the utility of in-helmet displays for presenting a complex 3-D world view to the user. This work involves the use of direct-eye projection, rather than an actual display screen inside the helmet. It provides the illusion of a 3-D world by sensing the direction in which the user's head is pointing and adjusting the projection accordingly. This is a good example of the kind of innovative work in novel interaction modalities that needs to be undertaken to exploit fully the potential for human-computer interaction on the space station.

Other kinds of novel output modalities on which further research could bring useful results include force or tactile feedback on joystick-type direct manipulation or analogue tasks and acceptably unobtrusive speech output. Force and tactile feedback has been used regularly in flying and remote manipulation tasks, but has been little explored for use in human-computer interaction for more abstract tasks, such as manipulating a set of computer files. Force or tactile feedback through a joystick on such problems could enhance the directness of the "feel" of direct manipulation interfaces and also be useful as an indicator of urgency, importance, or difficulty. Speech output has also been used before, but a recurring difficulty is getting the speech output to fit naturally into the flow of an interaction. Speech output is by its nature transitory and must be given at just the right point in the interaction and be repeatable by the user if desired. Moreover, the speech output should not occur so frequently that it becomes distracting to the user. Just as in the case with input modalities, much work is needed in the form of imaginative explorations over a large range of untried and speculative output modalities.

Finally in this section, we turn to the idea of expert interfaces, i.e. interfaces that require considerable expertise and training to operate, but offer high rates of very efficient interaction in return. The high degree of training that will be undergone by many space station personnel provides good opportunity for use of innovative expert interfaces, involving coordinated use of multiple limbs, eyes, etc. in multiple modalities for high efficiency interaction. Flying is best explored example of such an activity, and many of the techniques developed with flying have been successfully transferred to docking and other such maneuvers in space. Another source of ideas for expert interfaces can come from musical performance (Buxton, 1986). Players of such instruments as the organ learn after a long period of training to use all four limbs in a coordinated fashion to produce an enormously high rate of command input to the instrument. For interaction tasks that are important enough to justify the large training periods involved and could benefit from a high data transfer rate, interfaces

which draw on the experience of flying and musical interfaces are well worth investigation.

INTELLIGENT INTERFACES

The need to plan for change in interfaces comes not only from the possibility for advances in interface modalities and the techniques used with them, but also from the increasing possibility of the development of intelligent interfaces. Intelligent interfaces are still a research area, rather than a set of proven interface techniques, but the potential benefits of truly intelligent interfaces in terms of ease of use make them an area worthy of investigation for future space station interfaces. Intelligent interfaces also fit very well with the increasing development of intelligent, autonomous application systems for space use. If an application exhibits intelligent task behavior, then it should also behave intelligently in its interaction with its user.

An initial fundamental distinction should be made in considering the potential of intelligent interfaces is the distinction between conversational or agent-like systems and tool or machine-like systems. Almost all current interfaces are of the tool/machine-like kind. Users of such systems accomplish a task by controlling a (hopefully) responsive, but essentially unintelligent system. Direct manipulation interfaces (see section titled "Graphically-Oriented Interaction") are the archetype of this kind of interface since they encourage the user to feel that he is directly controlling the world that the underlying system deals with. However, command language interfaces can also be thought of as tool/machine-like since they respond in predictable ways to a fixed set of commands. The user is left feeling firmly in control.

Conversational/agent interfaces, on the other hand, are intended to give the user an entirely different feeling. Users of conversational/agent systems are intended to feel that they are negotiating with a subservient, but intelligent, system. They accomplish their tasks through negotiation with and through the agency of the system, rather than through direct actions of their own. Conversational systems thus have much greater possibilities for intelligent interaction than machine-like systems. Conversational systems also do not fit well with the direct manipulation or command language styles of interface, but fit much better with natural language or speech interfaces which naturally lend themselves to a dialogue style. Interfaces to intelligent, autonomous application systems can also make good use of a conversational style of interaction.

The user of a conversational equipment reservation system might, for instance, request (in natural language) the reservation of a certain piece of equipment and then be engaged by the system in a dialogue concerning the period of the reservation and if the equipment was unavailable the possibility of substitute equipment or substitute times. The user of a tool/machine-like interface to the same underlying functionality would, on the other hand, expect to be forced

to specify the reservation times through constraints on the interaction enforced by the interface. If equipment was unavailable at the desired time, he would also expect to have to initiate a search himself through alternative times and substitute equipment. It is clear that the culture within NASA is very much oriented to tool/machine-like interfaces and moreover to interfaces in which the degree of control exercised by the user is very high. There are historical reasons for this related to the importance placed from early on in the space program (Loftus, 1986) on having as much human control as possible available so that there would be the maximum chance of fixing any problems that arose. As systems increase in complexity, the tool/machine-like interfaces have tended to reduce the amount of complexity (and therefore fine control) available to the user without, however, crossing over the line that separates tools from agents. At the current state of the art, this approach is entirely as it should be. There are no successful operational interfaces anywhere that could fairly be described as true conversational/agent systems². However, the promise of intelligent conversational systems remains. If this promise is successfully realized, then it offers an attractive way of achieving the goal of having a large variety of non-mission-critical space station system easily available to a broad class of users.

The key to the development of conversational/agent interfaces lies in the development of detailed models of the task and the user. To produce intelligent agent behavior, it is necessary to use Artificial Intelligence techniques to model what tasks the user can accomplish through the interface, how he can achieve his goals, and what his current goals and state of knowledge are. Previous work that has tried to do this includes (Huff and Lesser, 1982; Mark, 1981; Card et al., 1983).

This detailed level of modelling is necessary for intelligent agent-like behavior because, without it, the interface can only respond to the user's individual actions and the very local context. Using our equipment reservation example, knowledge of what purpose the user might be trying to achieve through use of a particular piece of equipment could allow the system to suggest a suitable alternative. Without that knowledge, the system can only respond on the availability of a particular piece of equipment.

This kind of modelling becomes much harder when the user is pursuing a goal that involves several system actions. An agent system then has to determine the nature of the higher level goal from observation of the individual actions. An electronic mail system, for instance, might observe that the user is trying to write a message out to a file and then use the contents of the file as the body of a message to another system user. If it recognized from this that the user was simply trying to forward the message to the other user, it could suggest an abbreviated method of doing so. Since individual system actions can often fit into many plans and since system users often interleave plans to achieve several goals, the detection of such larger scale goals out of lower level actions is a very hard task. A system that has such an ability can, however, assist the user in a variety of ways including suggesting simpler ways of doing things (as in the example above),

warning about pitfalls that it can foresee could lead to the user's current plan not achieving his overall goal, offering to take over and complete the plan it believes the user to be following, or offering to perform the next action or actions in the plan whenever it becomes clear what they are.

The kinds of task and user modelling abilities mentioned above could be used in conjunction with any kind of interface, not just one that uses natural language. However, agent-like interfaces fit particularly well with natural language for two reasons. First, natural language is a natural medium for the kinds of negotiation that arise when a system is trying to respond to the goals it believes its user to have rather than direct commands. Second, the goal and task models themselves can be very useful in natural language and speech understanding. The biggest single problem in natural language processing is handling ambiguity of various kinds (syntactic, semantic, referential, etc.) and if one version of the ambiguity makes sense in the context of the other user model and the other does not, then the one that does not fit can be eliminated.

The whole area of conversational modelling is still in its infancy. Much work remains to be done to produce usable systems. However, progress in this field is necessary for truly intelligent interfaces, whether or not they are based on natural language. Given the potential benefits of intelligent interfaces to the space station, it is an area of research well worth pursuing.

The same kind of techniques that go into pure conversational systems can also be used in conjunction with more conventional interaction techniques to produce a hybrid kind of interface that incorporates both conversational/agent and tool/machine-like components. The basic flavor of such an interface is essentially tool/machine-like. The conversational component serves as medium through which the system and user can exchange comments about what is going on in the central tool/machine-like component. The user can also use the conversational component to instruct the system indirectly to perform actions or present information that he could perform or request directly (though perhaps more tediously) through the tool/machine-like component.

A system of this kind has several advantages. First, pure conversational systems are unsuitable for any task that can be performed effectively through direct manipulation techniques, and particularly for tasks that involve continuous/analogue interaction. Adding a conversational/agent component to a tool/machine-like direct manipulation interface for performing such tasks allows the basic task to be performed in the most efficient manner, but also allows components of that task that could benefit from a conversational approach to do so. Examples of conversational interaction in such a situation include: the user requesting information that would require multiple actions to retrieve through the direct manipulation interface; the user asking questions about how to use the direct manipulation interface component; the system volunteering information about more efficient ways to use the direct manipulation component; the user requesting the system to achieve a higher level goal that would require extensive interaction with the direct manipulation component.

A second advantage of this kind of hybrid system is that the conversational component does not have to be used at all if the user does not so desire. This kind of arrangement may be the best way to introduce conversational systems into a culture like NASA's that has good reason to be cautious about such systems. The unproven nature of conversational/agent systems suggests that they be introduced in a way that gives their user alternative methods of accomplishing all their tasks.

This kind of hybrid agent/machine-like interface requires the same technological underpinnings as pure conversational systems and hence the same research program. However, it also requires additional work on how to integrate the two components in a smooth way. Some work (Negronponte, 1981. Bolt, 1980; Hayes, 1987b) has already been done in this area, but much more is required.

PLANNING FOR CHANGE IN INTERFACES

The previous two sections have discussed some of the potential developments in interface modalities and techniques that will generate the need for change in human-computer interfaces during the life of the space station. In this section, we turn to the issue of how to deal with such change.

User Interface Management Systems

The essence of the approach discussed here is based on hooking, i.e. designing software for future extension and modification. The kind of hooking envisaged is determined by the assumption that it is unnecessary and probably infeasible to rewrite the underlying application systems whenever interfaces change. This means that the application systems need to be hooked in such a way that new interface systems can be developed for them without changes to the applications themselves. This in turn means that applications and interfaces must be written in as separate a way as possible with communication between them as narrow and as tightly defined as possible.

There is already a substantial body of work in the human-computer interaction literature on this kind of separation between application and interface, e.g. (Tanner and Buxton, 1983; Hayes and Szekely, 1983; Hayes et al., 1985; Wasserman and Shewmake, 1984; Jacob, 1984; Yuntan and Hartson, 1984). The systems developed to achieve this kind of separation are known as user interface management systems (UIMSS). However, work to date is far from achieving a consensus on the best way to achieve the desired separation or indeed the degree of separation that is desirable, appropriate, or possible. This is unfortunate from the point of view of building the software for the space station IOC, since to achieve any useful degree of separation both interface and application must be built using a strict model of the kinds of communication that can occur between application and interface. Decisions made now on this kind of communication will affect the

possibilities for interface/application separation for the life of the space station. Since research work in this area is far from reaching a conclusion about what is the best model of communication, whatever model is adopted now is likely to be considerably less than optimal. However, adopting some model may be better than none at all, so the remainder of this section reviews current research and future directions in the area of UIMS work.

The basic model adopted by most work on user interface management systems is shown in Figure 1. The user communicates with the UIMS which in turn communicates with the application. Communication between the UIMS and the application is achieved through a carefully defined protocol which limits the kind of interaction that can occur. A typical repertoire of communication events might include:

- request from the UIMS to the application to perform a particular operation with a certain set of parameters
- notification by the application of completion of an operation
- update by the application of a variable indicating progress towards completion of an operation
- error message from the application
- request from the UIMS for a check on the semantic validity of a proposed parameter for an application operation
- reply from the application to such a request

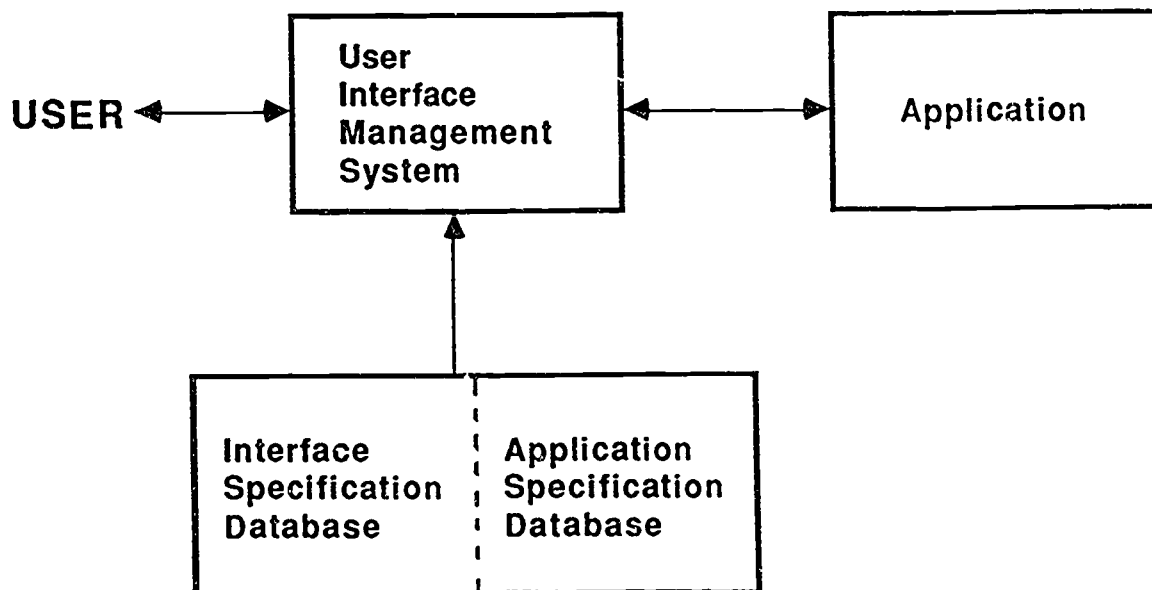


FIGURE 1 Model of communication in a UIMS

The precise content of the messages that flow between UIMS and application is defined by a declarative data base, the Application Specification Data Base of Figure 1, which specifies what actions and operations the application is capable of.

This model is not the one adopted by the most usual approach to interface standardization, that of providing a set of standard subroutines for high-level interface actions, such as getting the user to choose a value from a fixed set by presenting him with a pop-up menu. A typical interface subroutine for this task might take a set of choices as a parameter and return one of the choices. The subroutine would take care of the details of presenting the user with the menu and interpreting his mouse movements in making a choice from it. A disciplined use of a comprehensive package of such subroutines can thus provide a significant degree of low-level consistency across applications that use it. However, it cannot provide some of the other advantages of the kind of separation between interface and application described above, as we shall see.

The kind of separation between application and interface shown in Figure 1 can allow the interface to change without any alteration to the underlying application, whether or not the interface is provided by a UIMS. A UIMS goes further by defining the behavior of the interface itself through another declarative data base (possibly integrated with the application specification data base). This interface specification data base governs the details of the way the user is able to issue commands to the application. It would govern, for instance, whether commands were selected from menus, from an array of icons, through a command language line, etc., or whether a particular parameter to a specific command would be selected from a menu, from a row of "radio buttons", or typed into a field on a form, etc.. The UIMS provides a basic set of facilities to perform these various kinds of interaction, and the interface developer chooses the desired kind of interaction out of this cookbook by an appropriate interface specification. This arrangement has several advantages:

- **Consistency:** Since interfaces for different applications use the same basic set of UIMS-provided facilities, the interfaces will be consistent at the level of interaction details (how menus work, how icons are selected, etc.). Careful design of the UIMS interface specification formalism can also lead to consistency at a higher level. Consistency of this kind is very important in the space station, particularly for those less mission-critical interfaces where not all users may be fully expert. The transfer effects made possible through consistent interface behavior will greatly facilitate interaction with unfamiliar interfaces. Moreover, consistency avoids the negative transfer effects that can impair operation of even familiar interfaces.
- **Ease of interface development:** Specifying an interface through the interface specification formalism of a UIMS should be significantly less effort than programming one from scratch.

The UIMS formalism should provide high-level abstractions that allow the interface developer to specify the interface in terms that relate to the functionality of the interface as perceived by the user, rather than having to program it in a conventional manner at a level of detail more closely related to the implementation of the interface. This remains true even if the conventional implementation uses a high-level subroutine package of interface operations - using a subroutine package still places the emphasis on implementation, rather than abstract interface operations.

- Easier convergence on good interfaces: Despite all the advances in human-computer interaction that have occurred and continue to occur, the only known way to produce an excellent interface that fully meets the needs of its users is to build (or realistically simulate) the interface, let users interact with it, and modify it to resolve the problems that are observed. It is generally necessary to go around this loop many times before the interface performs satisfactorily, so anything that makes the loop easier and cheaper to follow is likely to improve the quality of the resulting interface by allowing more iterations. The UIMS model can speed up the modification part of the loop since interface modification can be done through modification of the declarative interface specification, rather than reprogramming in a conventional sense. This leads to a speed up in the loop as a whole.
- Ease of involvement of human factors experts: Since the UIMS model does not require programming to specify interface behavior, the interface specification can be done directly by people who are specialists in human-computer interaction, rather than by programmers. This allows better division of labor during interface/application development. Also, since programmers often think in terms of implementation ease and efficiency, rather than thinking about the interface from the user's point of view, better initial interfaces are likely to result if they are produced mainly by human factors specialists.

Of this set of advantages, only the first, consistency, and that at a relatively low level, is shared by the alternative approach of using a set of standardized interface subroutines. The other advantages all rely on a level of separation between interface and application that the subroutine approach does not provide.

Given this significant set of advantages for the UIMS approach, the natural question is why are all interfaces not produced through UIMSS. The answer is that current UIMS system approach the ideal described above only imperfectly. There are several specific problems.

The primary problem is that the constraints imposed by the need for an interface specification make it hard to provide ways of specifying interfaces that are carefully tailored to the needs of an individual application. Solutions to this problem (Szekeley, 1987) have tended to

introduce a procedural component into the interface specification formalism. The ability to program interaction allows the interface builder to tailor interface behavior to individual interface needs. The problem with this solution is that it tends to negate the benefits of the UIMS approach, such as consistency and ease of interface modification, that depend on the interface being specified declaratively. The way around this difficulty may be to include a procedural component in the interface specification formalism, but organize it at as high a level of abstraction as possible from the interface point of view. The procedural component could then be seen as a highly specialized programming language for interface specification. Such a language could conceivably maintain consistency by encouraging through its available constructs a particular style of interaction. Ease of use for rapid interface development and use by human-computer interaction specialists would be promoted by the high-level of the abstractions involved. A great deal more research would be needed to bring this idea to fruition, but the potential payoff could be great.

A second problem with current UIMS work is that the model of communication between application and interface is too limited. Many UIMS models allow only a subset of the list of message types listed above as flowing over the UIMS/application link. And even that list is insufficient for a sizable portion of applications, especially those involving graphical or analogue manipulation, which need a much closer coupling with their interfaces than that list of communication events allows. Again, the solutions that have been explored (Szekeley, 1987; Myers and Buxton, 1986) tend to change the model in the direction of tailoring the UIMS/application link to the needs of particular applications through use of a specialized programming language - a move away from the cleanest form of the UIMS model. A compromise here may be to develop several general UIMS/application communication protocols for large classes of applications with similar needs, while still leaving open the possibility of specialized communication protocols for particular applications.

A final problem with current UIMS work concerns the potential discussed earlier for interfaces employing multiple interaction modalities in effective coordination. The coordination of the different modalities increases the challenge for the UIMS model, and the use of a UIMS approach with multiple modalities has not been explored.

Work is needed to overcome all these problems if the UIMS approach is to be practical for the space station. Unfortunately, if the UIMS approach is to be used at all, a UIMS/application communication model must be adapted before the underlying applications are developed. Since meeting the needs of complex applications through a UIMS model is still a research problem with no clear solution, the only practical way a UIMS approach can be adopted for the space station IOC is to choose that (probably quite large) subset of simpler space station applications that can be adequately serviced by currently well-developed UIMS/application communication protocols. Research in extending the limits of applicability of these protocols could

nevertheless be useful for new systems developed after IOC. If these practical difficulties of adopting a UIMS approach appear too formidable for IOC, the fall-back position would be disciplined use of a comprehensive package of interface subroutines. This fall-back approach would provide the major advantage of a significant level of consistency across applications.

Interface Development Environments for Rapid Prototyping

A topic highly related to the UIMS approach to interfaces is that of interface development environments. Since the only known way to generate excellent interfaces is through an iterative process of creation, testing with users, and modification, a rapid prototyping facility for interfaces can materially improve the quality of interfaces produced by making it easier and faster to go around this loop. The rapid prototyping facilities most useful from this point of view allow interfaces to be seen and interacted with as they are developed, rather than forcing the interface developer to create the interface through working in a programming language or other formalism distinct from the interface itself. Examples of this approach include (Gould and Finzer, 1984; Myers and Buxton, 1986). They can be thought of as interface editors analogous to a what-you-see-is-what-you-get (wysiwyg) text editors. Such interface editors are a relatively new arrival on the human-computer interaction scene; their utility means they deserve a great deal more research attention.

Although rapid prototyping facilities can exist independently of the UIMS approach to interface design, they fit well with it. The cleanness of the based separation between application and interface in the UIMS model makes an interface development environment particularly useful in conjunction with a UIMS approach. A UIMS interface can be developed before the real application is available (or without incurring the expense of running the real application) by creating a dummy application that operates according to the same UIMS/application protocol as the real application. Coupled with a rapid prototyping facility, this capability allows rapid development of interface mock-ups to provide cheap and fast initial "sanity checks" on interfaces as they are developed.

Another intriguing possibility with wysiwyg interface development environments is their use (probably in restricted mode) by end users to reconfigure interfaces to their personal needs or preferences. So long as the interface modification facilities are made as easy to operate as the interfaces themselves, and so long as they do not interfere with the normal operation of the interfaces, this kind of facility could serve to improve significantly the level of personal satisfaction that space station users find with their interfaces.

Work in the area of wysiwyg interface development facilities has been almost entirely concentrated on graphical direct manipulation interfaces. This is natural in that it is the visual aspect of the interfaces that is most natural to specify in this manner. However,

additional work is needed both to develop techniques for this kind of interface further, and to extend the natural interface specification techniques to multi-mode interfaces as well.

CONCLUSIONS

This paper has focussed on change in space station interfaces - the reasons that it must be expected and ways to plan for it. We have identified several topic areas associated with these two aspects of change in space station interfaces in which further research effort would be beneficial. We conclude by listing several broad areas in which we particularly recommend the support of further work.

- investigation of speech recognition techniques and natural language processing techniques for use with spoken input, plus the integration of both of these modalities with direct manipulation interfaces;
- exploration of innovative I/O devices suitable for the space station environment;
- work on the user and task modelling needed to support conversational interfaces and the integration of such interfaces with machine-like direct manipulation interfaces;
- continued development of the UIMS concept, coupled with highly interactive interface development environments for all interface modalities.

NOTES

1. The complementary concept of scarring (designing hardware for future extension and modification) is also well established, but is not addressed in this paper.
2. Though see Mark (1981), Carbonell, et al., (1983), and Douglass and Hegner (1982), for examples of successful experimental agent systems.

REFERENCES

- Bolt, R. A.
1980 Put-that-there: voice and gesture at the graphics interface. Computer Graphics 14(3):262-270.

- Buxton, W.
1986 There's more to interaction than meets the eye: some issues in manual input. Pp. 319-337 in User Centered System Desigr. D. A. Norman and S. W. Draper, eds., Erlbaum, New Jersey:
- Carbonell, J. G., Boggs, W. M., Mauldin, M. L., and Anick, P. G.
1983 The XCALIBUR project: a natural language interface to expert systems. Proceedings of the Eighth International Joint Conference on Artificial Intelligence. Karlsruhe, August:
- Card, S. K., Moran, T. P., and Newell, A.
1983 The Psychology of Human-Computer Interaction. Hillsdale, N.J.: Erlbaum.
- Druglass, R. J. and Hegner S. J.
1982 An Expert Consultant for the UNIX Operating System: Bridging the Gap Between the User an Command Language Semantics. Los Alamos National Laboratory.
- Gould, L. and Finzer, W.
1984 Programming by rehearsal. Byte 9(6):187-210.
- Hansen, W. J.
1971 User engineering principles for interactive systems. Pp. 523-532 in Proceedings of the AFIPS, Fall Joint Computer Conference.
- Hayes, P. J.
1986 Steps towards Integrating Natural Language and Graphical Interaction for Knowledge-based Systems. Proceedings of the Seventh European Conference on Artificial Intelligence, Brighton, July, Pp. 456-465.
- 1987a Using a knowledge base to drive an expert system interface with a natural language component. In J. Hendler, ed., Expert Systems: The User Interface. New Jersey: Ablex.
- 1987b Intelligent interfaces to expert systems. In T. Bernold, ed., User Interfaces, Gateway or Bottleneck? North Holland.
- Hayes, P. J. and Reddy, D. R.
1983 Steps toward graceful interaction in spoken and written man-machine communication. International Journal of Man-Machine Studies 19(3):211-284.
- Hayes, P. J. and Szekely, P. A.
1983 Graceful interaction through the COUSIN command interface. International Journal of Man-Machine Studies 19(3):285-305.

- Hayes, P. J., Szekely, P. A., and Lerner, R. A.
 1985 Design Alternatives for User Interface Management Systems Based on Experience with Cousin. Proceedings of CHI85, San Francisco, April.
- Huff, K. E. and Lesser, V. R.
 1982 Knowledge-based command understanding: an example for the software development environment. Computer and Information Sciences. University of Amherst, Massachusetts.
- Hutchins, E. L., Hollan, J. D., and Norman, D. A.
 1986 Direct manipulation interfaces. Pp. 87-124 in D. A. Norman and S. W. Draper, eds., User Centered System Design. New Jersey: Erlbaum.
- Jacob, R. J. K.
 1984 An executable specification technique for describing human-computer interaction. In H. R. Hartson, ed., Advances in Human-Computer Interaction. New Jersey: Ablex.
- Loftus, J. P.
 1986 Space: Exploration-Exploitation and the Role of Man. Johnson Space Center: NASA
- Mark, W.
 1981 Representation and inference in the consul system. Pp. 375-381 in Proceedings of the Seventh International Joint Conference on Artificial Intelligence. August, Vancouver.
- Martin, J.
 1973 Design of Man-Computer Dialogues. New Jersey: Prentice-Hall.
- Myers, B. A. and Dixton, W.
 1986 Creating highly interactive and graphical user interfaces by demonstration. Pp. 249-258 in Computer Graphics: SIGGRAPH '86 Conference Proceedings, August, Dallas, Texas.
- Negronponte, N.
 1981 Media room. Proceedings of the Society for Information Display 22(2):109-113.
- Sacks, H., Schegloff, E. A., and Jefferson, G.
 1974 A simplest semantics for the organization of turn-taking for conversation. Language 50(4):696-735.
- Schegloff, E. A., Jefferson, G., and Sacks, H.
 1977 The preference for self-correction in the organization of repair in conversation. Language 53(2):361-382.

- Shneiderman, B.
1981 Direct manipulation: A step beyond programming languages. Computer 16(8):57-69.
- Smith, D. C., Irby, C., Kimball, R., Verplank, W., and Harslem, E.
1982 Designing the star user interface. Byte 7(4):242-282.
- Szekeley, P. A.
1987 Separating User Interface and Application. Ph.D. Th., Carnegie-Mellon University Computer Science Department.
- Tanner, P. and Buxton W.
1983 Some Issues in Future User Interface Management System (UIMS) Development. IFIP Working Group 5.2 Workshop on User Interface Management, Seeheim, West Germany, November.
- Wasserman, A. I. and Shewmake, D. T.
1984 The role of prototypes in the user software Engineering (USE) methodology. In H. R. Hartson, ed., Advances in Human-Computer Interaction. New Jersey: Ablex
- Watt, W. C.
1968 Habitability. American Documentation 19:(3):338-351.
- Williams, G.
1984 The Apple Macintosh computer. Byte 9(2):30-54.
- Yunten, T. and Hartson, H. R.
1984 Supervisory methodology and notation (SUPERMAN) for human-computer system development. In H. R. Hartson, ed., Advances in Human-Computer Interaction. New Jersey: Ablex.

COGNITIVE FACTORS IN THE DESIGN AND DEVELOPMENT OF SOFTWARE IN THE SPACE STATION

Peter G. Polson

Achievement of the operational and productivity goals for the Space Station will require extensive use of a wide variety of computer-based systems ranging from application programs that run on general purpose work stations to specialized embedded computer systems that monitor, operate, and trouble shoot critical subsystems, e.g., environmental and power control systems (Anderson and Chambers, 1985; Johnson et al., 1985). However, improperly designed user interfaces for these systems will compromise these goals.

The objectives of this chapter are to characterize major problems involved in the design of human-computer interfaces for systems on the Space Station and show how systematic application of empirical and theoretical results and methodologies from cognitive psychology and cognitive science can lead to the development of interfaces that reduce training cost and enhance space station crew productivity. This chapter focuses on four issues: 1) transfer of user skills, 2) comprehension of complex visual displays 3) human-computer problem solving, 4) management of the development of usable systems.

PROBLEMS

Transfer of User Skills

Inconsistent user interfaces in which the same basic function is performed by several methods in different contexts reduces transfer and interferes with retention (Polson, 1987; Postman, 1971). The Space Station's numerous computer-based systems and applications programs will be developed by different organizations over a period of many years. Inconsistency will be the rule rather than the exception unless extraordinary measures are taken in the design of user-interfaces for these systems. Popular and powerful applications programs developed for personal computers could be realistic models for software developed for the Space Station.

The typical popular applications program for a personal computer has been developed by an independent organization; the program has a great deal of functionality which is the reason for its commercial success.

The user interface is unique to the application being embedded in the application's code. Effective use of the application requires specialized training and several weeks of experience. There is no consistency across different popular applications. For example, they can have very different methods for editing operations on a text string. Thus, editing an axis label on a graph, editing an operating system command, or modifying a line of text with an editor all require different sequences of user actions.

The Comprehension of Complex Visual Displays

Complex visual displays using graphics, color, and possibly motion will be used in the space station to present various kinds of information to crew members carrying out complex tasks. Poorly formatted, poorly organized, and difficult to comprehend displays will have negative impacts on the productivity. Such displays increase training costs, difficulty of complex tasks, and probability of serious operator errors.

There exists extensive knowledge of the processes involved in the perception of basic visual properties like color and form (Graham, 1965; Walraven, 1985), and there are numerous guidelines for display layouts and use of symbols and color (e.g. Smith and Moser, 1984; Kosslyn, 1985). However, there is no systematic knowledge of how people comprehend complex displays or use the information presented in such displays to perform complex tasks. There are no general principles for the development of effective complex displays.

Human-Computer Problem Solving

NASA has extremely ambitious plans for the use of artificial intelligence and robotics in the space station. The proposed application areas include information management, life support systems operations and monitoring, electrical power systems operations and monitoring, and guidance and navigation. Many of these tasks on the Space Station will be performed by systems with significant embedded intelligence in order to satisfy mission, technological, and economic constraints and to achieve productivity goals (Anderson and Chambers, 1985).

The use of artificial intelligence techniques can significantly increase the complexity of a system from the point of view of its human user. The crew member must now understand both the task performed by the system as well as the characteristics of the "intelligent" control program (Hayes, 1987). Waterman (1986) notes that expert systems are "brittle" when pushed beyond the very narrow domain of their real expertise can fail with little or no warning. Uncritical use of the current state-of-the-art in expert systems' technology could decrease productivity of the crew and endanger their safety. Achievement of NASA's plans for the applications of artificial intelligence in the

space station will require extensive basic research and rapid advances in the state-of-the-art.

SOLUTIONS

Four solutions are proposed for the problems outlined in the preceding sections: 1) Use of information processing models of tasks in the design process, 2) allocation of adequate resources to user-interface development, 3) use of user interface management systems, and 4) use of existing expertise in NASA.

Detailed Information-processing Models

The first, and most important, solution is that designs for applications programs, complex visual displays, and cooperative human-computer problem solving systems be based on detailed, information-processing models of cognitive processes involved in the performance of specific tasks. Information-processing models describe the knowledge, cognitive operations, and user actions required to perform a task. These models can also be used to generate predictions of usability parameters, e.g. training time, productivity, and mental work load, and they can be used to isolate design flaws in proposed versions of a computer-based system.

Information-processing models describe what transfers, the knowledge necessary to perform the task, and thus they can be used in the design of consistent user interfaces that facilitate transfer of user skills. Information-processing models can make important contributions to the development of effective complex visual displays. The models describe both the knowledge necessary to successfully complete a task, what is to be displayed, and the processes involved in extracting that knowledge from displays, how it is to be displayed.

Information-processing models are an important component in the successful development of effective human-computer problem solving systems. There is general agreement that successful human-computer problem solving systems will incorporate models of the task and the user (Hayes, 1987). Current theoretical methodologies in cognitive psychology and cognitive science can be used to develop both kinds of models.

Management of the Design Process

The second solution involves successful management of the development process for computer-based systems. The typical development process for complex computer-based systems in the military, NASA, and the civilian sector does not allocate enough resources to usability considerations. The primary focus of the process is on developing a system with specified functionality. Functionality is necessary but not sufficient for usability. Usability, training time and

productivity, is typically evaluated late in the design cycle when it is far too late to make changes that improve usability. The design of highly productive complex computer-based systems requires solving simultaneously two interrelated sets of design problems involving functionality and usability.

What is proposed in this chapter is that usability and functionality considerations receive equal weight during all phases of the design cycle. The preliminary version of the system is evaluated for usability. If the system fails to meet usability goals, the design is revised. The revised design is then evaluated. This iterative process continues until the design meets both usability and functionality goals (Gould and Lewis, 1985; Hayes, 1987).

User Interface Management Systems

The third solution involves the use of appropriate technologies. Many of the problems involving transfer of user skills and consistency across applications can be solved using user interface management systems. The nature of these systems is discussed in Hayes (1987) and Hayes, Szekely, and Lerner (1985). They will not be discussed further here.

Existing Expertise in NASA

The fourth solution involves making effective use of the expertise already within NASA. What is being proposed here is similar to other modeling efforts currently underway in NASA dealing with problems of anthropometrics and habitability. OPSIM (Globus and Jacoby, 1986) is a computer model that simulates crew actions and interactions carrying out specific tasks under constraints imposed by different interior configurations, crew size and skills and other environmental factors. These simulated task scenarios are used to rapidly explore a large number of variables involving the environment and crew composition iteratively developing a more optimal design. Detailed models of the cognitive operations and physical actions required to carry out various types of tasks involving interaction between man and machine can be used in a similar fashion to optimize designs for user interfaces.

Alternative Solutions

Guidelines and Handbooks

Human factors guidelines (Smith and Mosier, 1986) and handbooks summarize information ranging from design goals and methodology to specific data on perceptual and motor processes. Guidelines and handbooks contain parametric information about basic perceptual and motor processes and information on limitations of classes of interaction techniques. However, they are of limited use in

characterizing higher-level cognitive processes, e.g. comprehension, learning, and problem solving. Guidelines propose reasonable design goals for cognitive aspects of a system, but they contain little or no advice on how to achieve such goals. Examples of cognitive guidelines include "minimize working memory load" and "minimize the amount of information the user has to memorize".

Usability parameters characterize the use of a system to perform a task, e.g. training time, productivity, and user satisfaction. Developing a system that optimizes usability parameters requires understanding of the task and the cognitive processes involved in performing the task. Most features incorporated into user interfaces are not good or bad per se. Usability is determined by interactions of the specific features of a design with the structure of a task. Guidelines do not contain necessary information about task structure, the knowledge required to perform a task, or the dynamics of the cognitive processing required to perform the task. Our knowledge of cognitive processes is in the form of detailed information processing models of the performance of complex tasks.

Many writers (e.g. Gould and Lewis, 1985; Hayes, 1987) argue that successful interface design is an iterative process. This view is strongly championed in this chapter. It is not possible to derive an optimal interface from first principles. Accumulated experience, information in guidelines and handbooks, and careful theoretical analyses can lead to the development of a reasonable initial trial design. However, this design has to be evaluated, modified, and evaluated again. In other words, guidelines and handbooks are not enough.

Empirically Based Modeling Strategies

Gould and Lewis (1985) and Carroll and Campbell (in press) seriously question the theoretically driven design and evaluation processes championed in this chapter. They argue that there are serious limitations of current modeling techniques, e.g. the limitations on our knowledge of comprehension of complex visual displays. They champion empirically-based modelling and evaluations methodologies. Many successful, complex systems, e.g. today's generation of highly automated aircraft, evolved from a combination of increasing technical capabilities, e.g. highly reliable microprocessors, and extensive operational experience (Chambers and Nagel, 1985).

However, relying on empirical methods to evaluate trial designs has serious limitations. They include difficulties in extrapolating results, doing experiments to evaluate complex systems, and evaluating transfer of training. For example, in a very complicated system, it may not be feasible to do empirical studies to evaluate a large number of tasks or to evaluate transfer between many tasks. If the current version of a trial design has unacceptable usability parameters, a designer has the very difficult task of deciding what attributes of the current design should be changed in order to improve performance. A theoretical model provides an explicit decomposition of the complex

underlying processes. This additional detail describing the underlying processes can be very valuable in making well motivated changes leading to the next iteration of the design process.

OUTLINE OF REMAINDER OF CHAPTER

The remainder of this chapter is organized into five sections. The first provides a general characterization of the kinds of theoretical models of cognitive processes that we argue should be the basis for the design of highly usable computing systems. The next section describes a detailed analysis of the process involved in the transfer of user skills and presents summaries of empirical results supporting these theoretical analyses. This section also provides a description of current theoretical models of human-computer interaction. Transfer is a well understood problem. The objective of this long section is to provide an illustration of a successful solution. The next section describes some of the difficult problems involved in the design of effective complex visual displays. The fourth section discusses the problems involved in the development of effective cooperative man-machine systems. The final section makes recommendations for further research.

MODELS OF COGNITIVE PROCESSES

The information processing framework (Newell and Simon, 1972; Gardner, 1985) provides the basis for the development of detailed process models of tasks performed on the Space Station. These theoretical analyses can be used as the basis for the design of human-computer interfaces that have minimal training costs and for the task and user models incorporated into human-computer problem solving systems.

The Information Processing Framework

An information processing model incorporates representations of the task, the knowledge required to perform the task, and the processes that operate on the representation to perform the task (Gardner, 1985). Such models are often formalized as computer simulation programs. The framework characterizes the general architecture of the human information processing system which in turn constrains the nature of the representations and the processes that operate on them, e.g., limited immediate memory. Newell and Simon (1972) and Anderson (1976, 1983) have proposed that the human information processing system can be described as a production system. The following section describes production system models of human-computer interaction.

Models of Human Computer Interaction

The GOMS model (Card et al., 1983) and Cognitive Complexity Theory (CCT) (Kieras and Polson, 1985) both characterize the knowledge necessary to make effective, routine use of software tools like an operating system, a text editor, or a data-base manager. The GOMS formalism describes the content and structure of the knowledge underlying these skills. CCT represents this knowledge as production rules which permits one to quantify amount. CCT incorporates all of the assumptions of the GOMS model. The production rule formalism enables one to derive quantitative predictions of training time, transfer of user skills, and performance. The next two sections describe each framework.

The GOMS Model

The GOMS model represents a user's knowledge of how to carry out routine skills in terms of goals, operations, methods, and selection rules.

Goals represent a user's intention to perform a task, a subtask, or a single cognitive or a physical operation. Goals are organized into structures of interrelated goals that sequence cognitive operations and user actions.

Operations characterize elementary physical actions (e.g., pressing a function key or typing a string of characters), and cognitive operations not analyzed by the theory (e.g., perceptual operations, retrieving an item from memory, or reading a parameter and storing it in working memory).

A user's knowledge is organized into methods which are subroutines. Methods generate sequences of operations that accomplish specific goals or subgoals. The goal structure of a method characterizes its internal organization and control structure.

Selection rules specify the conditions under which it is appropriate to execute a method to effectively accomplish a specific goal in a given context. They are compiled pieces of problem solving knowledge. They function by asserting the goal to execute a given method in the appropriate context.

Content and Structure of a User's Knowledge

The GOMS model assumes that execution of a task involves decomposition of the task into a series of subtasks. A skilled user has effective methods for each type of subtask. Accomplishing a task involves executing the series of specialized methods that perform each subtask. There are several kinds of methods. High-level methods decompose the initial task into a sequence of subtasks. Intermediate-level methods describe the sequence of functions necessary to complete a subtask. Low-level methods generate the actual sequence of user actions necessary to perform a function.

A user's knowledge is a mixture of task-specific information, the high-level methods, and system-specific knowledge, the low-level methods. The knowledge captured in the GOMS representation describes both general knowledge of how the task is to be decomposed as well as specific information on how to execute functions required to complete the task on a given system.

Cognitive Complexity Theory

Kieras and Polson (1985) propose that the knowledge represented in a GOMS model be formalized as a production system. Selection of production systems as a vehicle for formalizing this knowledge was theoretically motivated. Newell and Simon (1972) argue that the architecture of the human information processing system can be characterized as a production system. Since then, production system models have been developed for various cognitive processes (problem solving: Simon, 1975; Karat, 1983; text comprehension, Kieras, 1982; cognitive skills: Anderson, 1982).

An Overview of Production System Models

A production system represents the knowledge necessary to perform a task as a collection of rules. A rule is a condition-action pair of the form

IF (condition) THEN (action)

where the condition and action are both complex. The condition represents a pattern of information in working memory that specifies when a physical action or cognitive operation represented in the action should be executed. The condition includes a description of an explicit pattern of goals and subgoals, the state of the environment, (e.g., prompts and other information on a CRT display), and other needed information in working memory.

Production Rules and the GOMS Model

A production system model is derived by first performing a GOMS analysis and then writing a program implementing the methods and control structures described in the GOMS model. Although GOMS models are better structural and qualitative description of the knowledge necessary to perform tasks, expressing the knowledge and processes in the production system formalism permits the derivation of well motivated, quantitative predictions for training time, transfer, and execution time for various tasks.

Kieras and Bovair (1986), Polson and Kieras (1985) and Polson et al. (1986) among others have successfully tested assumptions underlying these predictions. These authors have shown that the amount of time

required to learn a task is a linear function of the number of new rules that must be acquired in order to successfully execute the task and that execution time is the sum of the execution times for the rules that fire in order to complete the task. They have shown that transfer of training can be characterized in the terms of shared rules.

TRANSFER OF USER SKILLS

In a following section, research on transfer of user skills in human-computer interaction will be reviewed. This research shows that it is possible to give a very precise theoretical characterization to large transfer effects, reductions in training time on the order of three or four to one. These results strongly support the hypothesis that large transfer effects are due to explicit relationships between different tasks performed on the same system or related tasks performed on different systems. Existing models of the acquisition and transfer of cognitive skills enable us to provide precise theoretical descriptions of these transfer processes. These same models can in turn be used to design consistent user interfaces for a wide range of tasks and systems that will promote similar large reductions in training time and saving in training costs.

A Theoretical Model of Positive Transfer

The dominant theoretical approach for explaining specific transfer effects is due to Thorndike and Woodward (1901) and Thorndike (1914). Thorndike assumed that transfer between two tasks is mediated by common elements. Common elements acquired in a first task that successfully generalize to a second do not have to be relearned during the acquisition of the second task. If a large number amount of the knowledge required to successfully perform the second task transferred, there can be a dramatic reduction in training time.

Kieras and Bovair (1986) and Polson and Kieras (1985) proposed that a common elements theory of transfer could account for positive transfer effects during the acquisition of operating procedures. The common elements are the rules. Tasks can share methods and sequences of user actions and cognitive operations. These shared components are represented by common rules. It is assumed that these shared rules are always incorporated into the representation of a new task at little or no cost in training time. Thus, for a new task in the middle of a training sequence, the number of new unique rules may be a small fraction of the total set of rules necessary to execute this task.

Examples of Successful Transfer

This section briefly describes results from the human-computer interaction literature demonstrating the magnitudes of the transfer

effects and showing how OCT (Kieras and Polson, 1985) can explain these results.

Polson et al. (1986) found very large transfer effects, on the order of four to one reductions in training time, for learning to perform a simple utility task on a menu-based, stand-alone, word processor. Their theoretical analysis showed that a significant portion of the knowledge, when quantified in terms of number of rules, required to perform these tasks were in consistent with low-level methods for making menu transitions, entering parameters, and the like.

Singley and Anderson (1985) found large transfer effects between different text editors, e.g., transfer from a line to a screen editor. Polson, Bovair, and Kieras (1987) found effects of similar magnitude for transfer between two different screen editors. Their theoretical analysis showed that editors share common top level methods that decompose the task of editing a manuscript into a series of subtasks involving individual changes in the manuscript. Furthermore, even very different editors share low-level methods, e.g., cursor positioning. Text editing is a task where transfer is mediated by knowledge of the general structure of the task as well as shared methods.

The Xerox STAR is a workstation that was explicitly designed to maximize the transfer of methods both within a given application as well as across different applications (Smith et al. 1983). All commands have a common format. The user first selects an object to be manipulated using specialized selection methods for different kinds of text or graphic objects. The operation is selected by pressing one of four command keys on the keyboard. For example, hitting the delete key causes the selected object to be deleted.

Ziegler et al. (1986) carried out transfer experiments with the STAR workstation. They studied transfer between text and graphics editors. They showed that common methods acquired in one context were successfully transferred to the other leading to very large transfer effects. Further, they were able to provide a quantitative analysis of the magnitude of these transfer effects using a production system model like those of Polson et al. (1987).

An Example of the Impact of Low Level Inconsistencies

Karat et al. (1986) examined transfer between three highly similar word processing systems that were intended by their designers to facilitate the transfer of user skills from one system to another. The first system was developed as a menu-based, stand alone word processor. A major goal in the design of the follow-on systems was to facilitate transfer from the dedicated, stand-alone, word processor to word processors hosted on a general purpose personal computer and a departmental computing system.

Karat et al. evaluated the magnitude of transfer effects from the dedicated version of the system to the other two system environments. The transfer effects were disappointingly small. Karat et al. found users' difficulties transferring their skill were due almost entirely to subtle differences in low level-methods. For example, many problems

were caused by the fact that the dedicated version of the system has specialized, labeled function keys. On the general purpose personal computer and the departmental computer system versions, the user had to learn and retain the locations of the corresponding functions on an unlabeled, generic keyboard. Inconsistencies in key assignments for activating known functions disrupted performance when users attempted to transfer their skills from one version of the system to another.

Implications for the Design of Systems in the Space Station

The research reviewed in preceding sections shows that common methods are transferred across tasks and application leading to large reductions in training time, on the order of 100% to 300%. However, the Karat et. al. results show that these transfer effects are fragile and can be reduced by minor but arbitrary differences in low-level methods let alone more extensive inconsistencies. For example, the method for centering text is identical on both the dedicated and personal computer versions of the systems except that the centering function is activated on the dedicated version by Control-Shift C and by Control-Shift X on the personal computer version. This small inconsistency disrupted the performance of skilled users of the dedicated version forcing them to stop and refer to documentation to find the correct function key. This inconsistency was caused by the fact that Control-Shift C already used by many applications programs to abort and return to the top level of operating system.

The potential for serious inconsistencies in common methods across different systems and application in the Space Station is much greater than the example of the three word processing system studied by Karat et. al. They were all developed by a single manufacturer with the explicit goal of permitting transfer of skills developed on the dedicated version of the system.

COMPREHENSION OF COMPLEX VISUAL DISPLAYS

Rapid developments in hardware and software technology permit the generation and presentation of very complex displays combining text, color, motion, and complex visual representations. There is limited understanding of how to effectively utilize these new capabilities. There is extensive knowledge of the basic visual processes underlying color and form perception (Graham, 1965; Walraven, 1985). Detailed models of the comprehension of complex visual displays do not exist. There is some systematic work on the effective graphical presentation of quantitative information (e.g., Kosslyn, 1985; Tufte, 1983). The widely acclaimed book The Visual Display of Quantitative Information by Tufte is a collection of design guidelines.

Today, development of effective complex displays relies almost entirely on empirically-based, iterative design methods (Gould and Lewis, 1985). A good illustration of how effective these methods can be is shown in an experiment reported by Burns et al. (1986). These

investigators were concerned with the problem of display format optimization. They designed a set of alternative displays to be used in orbital maneuvering tasks onboard the Space Shuttle. The new displays grouped information by function and include more meaningful abbreviations and labels. Burns et al. (1986) had both non-experts and Space Shuttle crew members retrieve specified items of information from the current, operational displays and the reformatted experimental displays.

Reformatted displays improved both speed and accuracy for the non-expert subjects. The changes in format had no effects on Space Shuttle crew member performance, and the reformatted displays improved their accuracy. These results are surprising. Extensive training and experience should have enabled the crew members to develop specialized skills to deal with even non-optimal displays. Any changes in display format should have disrupted these skills leading to reductions in performance for highly trained crew members. One possible conclusion is that the current displays are so far from optimal that even brief experience with the reformatted displays enabled trained crew members to perform at a level equal to their performance with actual displays.

The Burns et al. (1986) experiment shows that application of our current knowledge of visual perception and guidelines for formatting and labeling can lead to significant improvements of performance in an empirically-based iterative design process. However, the situation in the Space Station is more complex. The display technology for the systems onboard the Space Shuttle used small, alpha-numeric CRTs. Displays onboard the Space Station will make extensive use of graphics and color. In other words, increase capabilities provided by new display technology will enable developers to generate truly incomprehensible displays.

Furthermore, there are important transfer and consistency issues. Conflicting uses of symbols, color and motion cues, and inconsistent formats across applications will have the same impact on users as inconsistent methods for entering text, increased training time and probabilities of user errors. Dealing with issues involving more complex displays, consistency, and the use of displays as interfaces to systems with significant embedded intelligence are more complex design problems. The design problems will have to be solved using the combination of empirically-based evaluation methods combined with detailed models of the task and a theory of the comprehension of visual displays.

Consider the design problems involved in developing the displays for systems with significant embedded intelligence like the Space Station's environmental controls and power systems. Effective displays should be based on 1) an understanding of the knowledge required to successfully perform critical tasks, e.g., trouble shoot a malfunction, 2) a characterization of the cognitive processes involved in extracting the necessary information from the display, 3) and a description of how the information is utilized to complete the task. In other words, what is required is a complete theory of the comprehension of complex visual displays.

Ellis and his colleagues (Ellis et al., 1985; Kim, Won Soo et al., 1985) have proposed a methodology for the development of effective specialized displays for spatial tasks involving control of objects in three dimensional space with a full six degrees of freedom, e.g. the JPL Telerobot demonstrator, and Space Station Proximity Operations Displays. Ellis and his colleagues propose a design methodology that creates a very tight link between the characteristics of the task, a theoretical understanding of the perceptual processes, and empirical demonstrations that the displays actually facilitate performance of the task. This design strategy can be generalized in all various types of displays and tasks.

HUMAN-COMPUTER PROBLEM SOLVING

NASA has articulated a very ambitious design philosophy for expert systems to be used on the Space Station calling for the development of cooperative human-computer problem solving systems. Many issues concerning the design of such systems can be understood from experience with highly automated commercial aircraft (Chambers and Nagel, 1985), automatic test equipment (Richardson et. al., 1985), and automated control systems for nuclear power plants. Some of the issues are: 1) vigilance of the human operator, 2) safe transition from automatic to manual modes of operation, 3) maintenance of skills necessary to perform tasks manually, 4) successful completion of a task after the automatic system has failed, 5) allocations of functions between man and machine, 6) and the development of truly symbiotic human-computer problem solving systems. Although the basic issue have been identified, there are no well worked out general solutions nor are there any operational examples of symbiotic human-computer problem solving systems.

Autonomous vs. Cooperative Systems

Hayes (1987) distinguishes between conversational/agent and machine/tool-like systems. In a conversational/agent system, the user interacts with an intelligent agent to accomplish a task. Robots that carry out complex tasks under human supervision and systems with sophisticated natural language interfaces are examples. Machine/tool-like systems are directly controlled by their users although they can be highly automated carrying out a whole sequence of low level steps without direct intervention. Examples include auto-pilots, automatic test equipment (ATE) and application programs like text editors and spreadsheets.

There also is a second important dimension, autonomy. Some systems, once initialized by their users, carry out their task completely autonomously or only make use of the human user as a low level sensor and manipulator. Examples include auto-pilots, ATE systems, and most expert systems. Auto-pilots and ATE systems are not normally considered intelligent. However, they carry out extremely complex

tasks autonomously. They may not be classified as intelligent systems in that they carry out their tasks using well understood algorithms. Many expert systems imply the human user as a low-level sensor and manipulator. The task is carried out autonomously. The user can ask for explanations of the final results or why the system requested a given piece of data in the process of completing the task (e.g., Shortcliffe, 1976).

Limitations of Current Expert Systems

Intelligent systems can actually complicate the task of human user, e.g., telerobots and applications with natural language interfaces. Bejczy (1986) shows that intelligent agents can impose additional difficulties for users because they have to understand both the control program and the task. For example, no natural language interface is capable of responding correctly to unrestricted input. Such interfaces understand a limited subset of natural language and may have no or limited capabilities for reasoning about the task. Thus, even if the user's request is parsed correctly, resulting commands may be an incomplete and/or incorrect sequence of operations necessary to complete the task.

Consider the problem of effective handoff from automatic to manual operation in a troubleshooting task, e.g., finding a serious fault in the power distribution system. Current expert systems do not make the transition from automatic to manual operation gracefully. Waterman (1986) observes that expert systems have narrow domains of expertise and they have no capability to reason about their limitations. Because they can't reason about their limits, such systems are little use in assisting a human problem solver once they have failed to find the cause of a serious fault. Thus, the system can fail catastrophically leaving its user with a task of manually diagnosing a serious fault.

Building a system capable of reasoning about its limits and providing the user with a useful explanations regarding failure is beyond the current state-of-the art. However, it's exactly this kind of capability that is required in a truly cooperative system. In summary, current expert systems are not cooperative problem solving systems. In the process of performing their task, humans serve in a very low level subservient role and when systems fail, they fail catastrophically providing their users with little or no information for the reason of the failure and no assistance in continued efforts to solve the problem.

Being able to reason about its own limitation is difficult because of constraints embedded in the fundamental properties of current knowledge representation schemes (Jackson, 1986). The rules in current expert systems contain a complex mixture of control knowledge and domain specific and general problem solving knowledge. Such systems have no explicit model of domain principles or any specific knowledge of their strategies. Exactly this kind of knowledge is required to produce coherent explanations (Clancy, 1983). This type of knowledge is also required to reason about limitations.

Cooperative Human-computer Problem Solvers

NASA's goals are far more ambitious than the development of autonomous intelligent problem solvers with explanation capabilities. It is repeatedly proposed in various NASA documents to develop cooperative or symbiotic human-computer problem solvers (Johnson et al. 1985; Anderson and Chambers, 1985).

Discussions about the possibility of developing such systems have a surprising uniformity. The authors observe that powerful problem solvers can be developed if systems exploit the complimentary strengths of human and machine permitting one to compensate for the weaknesses of the other. The next issue is function allocation. The discussion of function allocation begins with a general assessment of the strengths and weaknesses of human and computers as problem solvers. This assessment is in the form of a characterizations human and machine components listing the strengths and weaknesses of each. Typical listings are in Johnson et al., 1985, pp. 27-28; Richardson et al., 1985, pp. 47-49; Anderson and Chambers, 1985. What is striking about these lists is their consistency. The following is taken from Richardson et al. (1985, pp. 47-49).

The strengths of the human component of the system are:

1. Processing of sensory data.
2. Pattern recognition.
3. Skilled physical manipulation but limited physical strength.
4. Limited metacognitive skills, e.g. ability to reason about limits of knowledge and skill.
5. Slow but powerful general learning mechanisms.
6. A large, content-addressable permanent memory.

The weaknesses of the human problem solver are:

1. Limited working memory.
2. Limited capacity to integrate a large number of separate facts.
3. Tendency to perseverate on favorite strategies and malfunctions; set effects and functional fixity.
4. Limited induction capabilities.
5. Lack of consistency; limitations on the ability to effectively use new information.
6. Emotional and motivational problems.
7. Limitations on the availability of individuals with the necessary abilities and skills.
8. Limited endurance.

The current generation of expert systems and highly autonomous automatic systems, e.g. ATE's make use of human sensory processing, pattern-recognition, and manipulative skills. Most authors recognize this and point out that their objective in developing cooperative problem solving systems is to exploit human's cognitive capabilities as well as these level skills. Continuing to quote Richardson, et al., the strength of the computer component of the system are:

1. Large processing capacity.
2. Large working memory.
3. Capabilities of making consistent mechanical inferences taking into account all relevant facts.
4. Processing and utilizing large amounts of actuarial information.
5. Capabilities to store and retrieve training and reference material.
6. Availability of system is limited only by reliability of basic computer technology.
7. No motivational or other related problems.

The weaknesses of the machine component of the system are

1. No or very limited capacity to adapt to novel situations.
2. No or very limited learning abilities.
3. No or very limited metacognitive abilities, e.g., understanding of own limitations.
4. Very difficult to program particularly the current generation of expert systems.

Examples of Cooperative Systems

The best examples of cooperative systems are intelligent training systems (ITS) (Sleeman and Brown, 1983; Polson and Richardson, 1987). The main components of an ITS are: 1) the expert module or task model, 2) the student module or user model, and 3) the tutor module or explanation subsystem. A cooperative, intelligent problem solving aid has to have real expertise about the task, an accurate model of the other intelligent agent that it is interacting with (the human user), and the capability of conducting sophisticated dialogues with the user. Richardson et al. (1985) argue that the machine component should attempt to compensate for known limitations and failure modes that are characteristics of all forms of human problem solving: They are working memory failures, set and functional fixity, inference failures, and attentional limitations.

One important role for a cooperative intelligent system would be to reduce information overload by selectively displaying information relevant to the highest priority subcomponent of a task. Chambers and Nagel (1985) describe the cockpit of a Boeing 747 with its several hundred instruments, indicators, and warning lights as an example of where skilled pilots can be simply overwhelmed by the amount of available information. Plans for highly automated aircraft of the 1990s incorporate selective displays on color CRTs of a small subset of the total information about the state of the aircraft that is relevant to the current task. The ability to display relevant information would prevent information overload and augment human working memory by providing an external representation relevant information about the system's state.

Other proposals for the role of the computer in a cooperative system focus on its computational capabilities. Memory limitations prevent

human users from adequately integrating information about the current state of the system and archival information concerning likelihoods of component failures. Thus, the machine takes on the role of filter, memory aid, and inference engine compensating for known general weaknesses in the human information processing system.

Possible Scenarios - Serious Problems

These proposals are consistent with the large body of data about the strength and weaknesses of human diagnostic reasoning and problem solving. However, implementing these proposals into a functioning system can cause serious difficulties. Consider a situation involving the power distribution system of the Space Station where several interacting failures have occurred. The system makes a series of incorrect inferences about the cause of the faults and displays to the human partner information irrelevant to successful solution of the problem. Such misinformation could effectively block successful solution by the human user. It's essentially a set manipulation. The misinformation would be especially damaging if the system were normally successful.

Other problems could result if the system makes incorrect inferences from its model of the human user. Assume the system has concluded, correctly, that it is incapable of independently diagnosing the faults in the power distribution system. Using its advanced explanation capabilities, it explains to its human partner its understanding of the current state of the power distribution system and various partial results obtained in attempting to diagnose failures. In the process, system presents a series of complex displays showing the current state of the power distribution system. The expert human user recognizes a complex pattern of interrelated events and informs the computer of the correct solution to the problem. The system responds by attempting to evaluate the human partner's input using information contained in its user model. This model has a very detailed description of the limits of the human information processing system, and the system incorrectly concludes that the human partner is incapable of making the correct diagnosis on the basis of such complex input and the solution is rejected.

Conclusions

Many readers may think that the scenario presented in the preceding section is overdrawn. Of course, NASA would never tolerate the fielding of a system that was capable of effectively overruling a Space Station crew member. However, a system in which human users can override the machine partner compromises the goal of developing truly cooperative human-computer problem solving systems. Information overload, working memory failures, and failures to integrate historical data in making diagnoses are highly probable failure modes of human users. The incorrect inference made by the machine described in the

preceding scenario is not unreasonable and would probably be correct in most situations. Experience with intelligence tutoring systems (Polson and Richardson, in press) shows that such cooperative systems are exceedingly difficult to construct.

RECOMMENDATIONS FOR FURTHER RESEARCH AND CONCLUSIONS

This section contains information on recommendations for further research and concludes that the difficulties in developing truly productive computer-based systems are primarily management problems.

Information Processing Models

Recommendation 1. Support the development of the software tools required to rapidly develop information processing models of tasks performed on the Space Station.

This chapter has recommended that information processing models of cognitive processes be the basis for the design of applications programs, complex visual displays and cooperative human-computer problem solving systems. A theoretical technology should be applied on a large scale to solve interface design problems on the Space Station. Unfortunately, the development of information processing models is currently an art and not a robust design technology. Furthermore, these models can be extremely complex simulating basic psychological process in detail (Anderson, 1983). What is required are engineering models (Newell and Card, 1986; Kieras & Polson, 1985).

Development of an effective modeling facility is an engineering problem, albeit a difficult one. There are no advances required in the theoretical state of the art in cognitive psychology. Models of various cognitive processes have to be integrated into a single simulation facility, e.g., models of perceptual, cognitive, and motor processes. Higher level languages should be developed that automate the generation of the simulation code and the detail derivation of models. A simulation development system will be required for designers to rapidly develop models of adequate precision for use in a timely fashion in the design process.

The Comprehension of Complex Displays

Recommendation 2. Support an aggressive research program on the processes involved in the comprehension of complex, symbolic displays.

Many tasks on the Space Station will require that crew members interact with complicated displays. Examples include monitoring and trouble shooting of complex subsystems, manipulation and presentation of scientific data, and interacting with expert systems to carry out trouble shooting and maintenance tasks. Rapid advances in computer and

display technology will enable designers to develop complex displays making using of symbolic, color, and motion cues. Effective displays that facilitate performance on these complex tasks can have large positive effects on crew productivity. The complexity of the tasks and the freedom given to the designer by the display technology require that successful designs be based on explicit models of how information in such displays is used to perform these tasks.

Development of models of the comprehension of complex displays requires important contributions to cognitive theory. Current research in cognition and perception provides a solid foundation on which to build such models. It is possible that models of comprehension of complex displays can be based on the extensive body of theoretical results obtained on the processes involved in text comprehension (e.g., van Dijk and Kintsch, 1983). Excellent work on related problems is already going on within NASA; research programs in this area could be modeled in the work of Ellis and his colleagues briefly described in a preceding section.

Human-Computer Problem Solving

Recommendation 3. Design and support an aggressive research program leading to the eventual development of cooperative, human-computer problem solving systems.

Although the many analyses characterizing cooperative human-computer problem solving are correct, development of a useful cooperative system requires solutions to unsolved problems in expert system design, artificial intelligence, and cognitive science. A well structured research program would generate many intermediate results, components of the eventual cooperative system, that are useful in themselves on the Space Station. These include robust, high performance expert systems, advanced explanation subsystems, and various problem solving tools to assist the crew in management of the Space Station systems.

Consider utilities of an inspectable expert system and of an inference engine tool. By an inspectable expert system, we mean a system that displays intermediate states of its diagnostic processes during trouble shooting. The expert systems tool presents to the trained user intermediate results of the trouble shooting process using of complex, symbolic displays. Properly designed, such information gives the human expert the information necessary to confirm a diagnosis or take over effectively if the expert system fails. Most current automatic test equipment simply reports success or failure, e.g., a red light or a green light. An inspectable expert system would be a dramatic improvement over diagnostic systems with such limited feedback.

Another useful subsystem would be an inference engine, a tool that combines information about system state with actuarial data on the likelihoods of different failure modes. This system would be designed to enable a skilled human user to do what if calculations and serve as

a memory aid reminding the crew member of infrequently occurring faults that are likely to be overlooked.

Inspectable expert systems are within the state-of-the-art and would serve as a very useful test bed for research on comprehension of complex symbolic displays and on the design of such displays. An interactive inference engine could be seen as a primitive prototype of a cooperative problem solving system. Both tools can be very useful in an operational environment and both are important intermediate steps in the eventual development of high performance cooperative systems.

There are important areas of research in cognitive science that will have to be better developed before it will be possible to build successful cooperative human-computer problem solving systems. These include models of human diagnostic reasoning, cooperative problem solving, and models of the processes involved in generating and comprehending useful explanations. A cooperative system must incorporate an extremely sophisticated model of its human partner which in turn requires a detailed understanding of how humans carry out the specific task performed by the system as well as the general characteristics of the human information processing system and its failure modes. User models are related to the problem of developing student models in intelligent training systems. Although progress is being made in the area of student modeling, there is still numerous important unsolved problems (Polson and Richardson, 1987).

In summary, the design and development of cooperative, human-computer problem solving is the most difficult of the technological goals related to cognitive science associated with the Space Station. This goal will only be achieved by a long term, well managed research program.

In Reality, It's a Management Problem

It is widely recognized that the ambitious productivity goals for the Space Station can only be achieved with extensive use of automated systems that have effective user interfaces. However, there is a broad gap between good intentions and actual development practice. It is widely recognized today that complex systems developed for civilian, NASA, and military use are far from the current state-of-the-art in human factors presenting serious problems for their users. Often, design errors are so obvious that applications of simple common sense could lead to the development of more usable interfaces.

In the final analysis, development of usable systems is a management problem. Consistent application of the current state-of-the-art in human factors and knowledge of cognitive processes during all phases of the development process would have dramatic and positive effects on the productivity of the Space Station crew.

REFERENCES

- Anderson, J. R.
1976 Language, Memory, and Thought. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- 1982 Acquisition of cognitive skill. Psychological Review 89:369-406.
- 1983 The Architecture of Cognition. Cambridge, MS: Harvard University Press.
- Anderson, J. R., and Chambers, A. B.
1985 Human Centered Space Station Design.
- Bejczy, A. K.
1986 Human factors in space teleoperation. Transactions of the 2nd International Symposium on Next Generation Transportation Vehicles. Analfi, Italy, June 20-24.
- Burns, M. J., Warren, D. L., and Rudisill, M.
1986 Formatting space-related displays to optimize expert and nonexpert user performance. Pp. 274-280 in H. Mantei, and P. Orbeton, eds., Proceedings CHI'86 Human Factors in Computer Systems. New York: Association for Computing Machinery.
- Card, S. K., Moran, T. P., and Newell, A.
1983 The Psychology of Human-Computer Interaction. Hillsdale, N.J.: Erlbaum.
- Carroll, J. M.
1987 Interfacing Thought. Cambridge, MS: In press. Bradford Books/MIT Press.
- Carroll, J. M., and Campbell, R. L.
1987 Softening up hard science: a reply to Newell and Card. Human-Computer Interaction. In press.
- Chambers, A. B., and Nagel, D. C.
1985 Pilots of the future: Human or computer? Computer November:74-87.
- Clancy, W. J.
1983 The epistemology of a rule-based expert system: A framework for explanation. Artificial Intelligence 20:215-251.

- Ellis, S. R., Kim, W. S., Tyler, M., and McGreevy, M. W.
1985 Visual enhancements for perspective displays: perspective parameters. Proceeding of the International Conference on Systems, Man, and Cybernetics. IEEE Catalog No. 85CH2253-3. November, 815-818.
- Engle, S. E. and Granda, R. E.
1975 Guidelines for Man/Display Interfaces. Technical Report TR 00.2720. Poughkeepsie, NY: IBM.
- Garnder, H.
1985 The Mind's New Science: A History of the Cognitive Revolution. New York: Basic Books.
- Globus, A. and Jacoby, R.
1986 Space Station Operational Simulation (OPSIM). Moffett Field, CA: NASA-Ames Research Center.
- Gould, J. D. and Lewis, C.
1985 Designing for usability: key principles and what designers think. Communications of the ACM 28:300-311.
- Graham, C. H.
1965 Vision and Visual Perception. New York: John Wiley and Sons.
- Hayes, P.
1987 Changes in human-computer interfaces on the space station: why it needs to happen and how to plan for it. Humans in Automated and Robotic Space Systems. National Academy of Sciences. Washington, D. C.
- Hayes, P. J., Szekely, P. A., and Lerner, R. A.
1985 Design alternatives for user interface management systems based on experience with cousin. Pp. 167-175 in L. Borman, and B. Curtis, eds., Proceedings of the CHI 1985 Conference on Human Factors in Computing. New York: Association for Computing Machinery.
- Jackson, P.
1986 Introduction to Expert Systems. Workingham, England: Addison-Westley Publishing Company.
- Johnson, R. D., Bershader, D. and Leifer, L.
1985 Autonomy and the Human Element in Space: Final Report of the 1983 NASA/ASEE Summer Faculty Work Shop. Moffett Field: NASA-Ames Research Center.
- Karat, J.
1983 A model of problem solving with incomplete constraint knowledge. Cognitive Psychology 14:538-559.

- Karat, J. Boyes, L. Weisgerber, S., and Schafer, C.
1986 Transfer between word processing systems. Pp. 67-71 in M. Mantei and P. Orbeton, eds., Proceedings CHI'86 Human Factors in Computer Systems. New York: Association for Computing Machinery.
- Kieras, D. E.
1982 A model of reader strategy for abstracting main ideas from simple technical prose. Text 2:47-82.
- Kieras, D. E. and Bovair, S.
1986 The acquisition of procedures from text: A production-system analysis of transfer of training. Journal of Memory and Language 25:507-524.
- Kieras, D. E., and Polson, P. G.
1985 An approach to the formal analysis of user complexity. International Journal of Man-Machine Studies 22:365-394.
- Kim, W. S., Ellis, S. R., Tyler, M., and Stark, L.
1985 Visual enhancements for telerobotics: perspective parameters. Pp. 807-812 in Proceedings of the International Conference on Systems, Man, and Cybernetics. IEEE Catalog No. 85CH2253-3, November.
- Kosslyn, S. M.
1985 Graphics and human information processing: A review of five books. Journal of the American Statistical Association 80:499-512.
- Newell, A., and Card, S. K.
1986 The prospects for psychological science in human computer interaction. Human-Computer Interaction 1:209-242.
- Newell, A., and Simon, H. A.
1972 Human Problem Solving. Englewood Cliffs, New Jersey: Prentice-Hall.
- Polson, M. C., and Richardson, J.
1987 Foundations of Intelligent Tutoring Systems. In press. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Polson, P. G.
1987 A quantitative theory of human-computer interaction. In J. M. Carroll, ed., Interfacing Thought. In press. Cambridge, MS: Bradford Books/MIT Press.

- Polson, P. G., Bovair, S., and Kieras, D. E.
 1987 Transfer between text editors. Pp. 27-32 in P. Tanner, and J. M. Carroll, eds., Proceedings CHI'87 Human Factors in Computer Systems. New York: Association for Computing Machinery.
- Polson, P. G., and Kieras, D. E.
 1985 A quantitative model of the learning and performance of text editing knowledge. Pp. 207-212 in L. Borman, and B. Curtis, eds. in, Proceedings of the CHI 1985 Conference on Human Factors in Computing. New York: Association for Computing Machinery
- Polson, P. G., Muncher, E., and Engelbeck, G.
 1986 A test of a common elements theory of transfer. Pp. 89-83 in M. Mantei and P. Orbeton, eds., Proceedings CHI'86 Human Factors in Computer Systems. New York: Association for Computing Machinery.
- Postman, L.
 1971 Transfer, interference, and forgetting. In J. W. King and L. A. Riggs, eds., Woodworth and Scholberg's Experimental Psychology. New York: Holt, Rinehart, and Winston.
- Richardson, J. J. Keller, R. A., Maxion, R. A, Polson, P. G., and DeJong, K. A.
 1985 Artificial Intelligence in Maintenance: Synthesis of Technical Issues. (AFHRL-TR-85-7) Brooks Air Force Base, TX: Air Force Human Resources Laboratory.
- Shortcliffe, E. H.
 1976 Computer Based Medical Consultations: MYCIN. New York: American Elsevier.
- Simon, H. A.
 1975 Functional equivalence of problem solving skills. Cognitive Psychology 7:268-286.
- Singley, K., and Anderson, J.
 1985 Transfer of text-editing skills. International Journal of Man-Machine Studies 22:403-423.
- Sleeman, D., and Brown, J. S.
 1983 Intelligent Tutoring Systems. London: Academic Press.
- Smith, D. C., et al.
 1983 Designing the STAR user interface. In P. Degano and E. Sandewall, eds., Interactive Computer Systems. Amsterdam: North-Holland.

- Smith, S. L. and Moser, J. N.
1984 Guidelines for Designing User Interface Software.
EDS-TR-86-278, MTR-10090. Bedford, MS: The Mitre Corp.
- Thorndike, E. L.
1914 The Psychology of Learning. New York: Teachers College.
- Thorndike, E. L. and Woodward, R. S.
1901 The influence of improvement in one mental function upon the efficiency of other functions. Psychological Review 8:247-261.
- Tufte, E. R.
1983 The Visual Display of Quantitative Information. Cheshire, CN: Graphics Press.
- van Kijk, T. A., and Kintsch, W.
1983 Strategies of Discourse Comprehension. New York: Academic Press.
- Walraven, J.
1985 The colours are not on the display: a survey of non-veridical perceptions that may turn up on a colour display. Pp. 35-42 in Displays.
- Waterman, D. A.
1986 A Guide to Expert Systems. Reading, MS: Addison-Wesley Publishing Company.
- Ziegler, J. E., Vossen, P. H., and Hoppe, H. U.
1986 Assessment of Learning Behavior in Direct Manipulation Dialogues: Cognitive Complexity Analysis. ESPRIT Project 385 HUFIT, Working Paper B.3.3, Stuttgart, West Germany: Fraunhofer-Institut für Arbeitswirtschaft und Organisation.

DISCUSSION: DESIGNING FOR THE FACE OF THE FUTURE:
RESEARCH ISSUES IN HUMAN-COMPUTER INTERACTION

Judith Reitman Olson

The hardest part of generating a research agenda now for issues in human-computer interaction for the Space Station is not in finding important issues and unanswered questions that are in need of careful research. It is selecting those research issues and the approaches to them that will answer the questions we have in the year 2000. In the year 2000, we will have devices that we can only dream of today; the Space Station environment will have a mission, size, and complexity that today we can only begin to sketch out. Our job, therefore, is not to recommend a research program that will answer specific questions that we know will arise in the design of the future Space Station. Rather, it is to prepare for that future with a research plan that lays the foundation, a sound theoretical base, that will make specific results both easy to predict and simple to confirm empirically. Additionally, the research has to produce a development environment, a flexible hardware platform and programming environment, that allows rapid prototyping for empirical testing and easy final implementation. These bases will serve us well when we have to make specific designs for the year 2000.

INTERFACES OF THE SYSTEMS OF THE FUTURE

It is important to begin by noting those things that are likely to be different in the Space Station environment than they are in the environments we focus our research on now. The most obvious differences, well discussed by both Polson and Hayes, are that the Space Station environment is weightless (with concomitant difficulties in forceful action and counteraction), perhaps noisy (with difficulties for the implementation of speech recognition and sound production), and complex (with a small number of people doing many, varied tasks with the help of computers, some of which they will be expert in, some of which they will not).

In addition, the tasks performed in the Space Station differ in other, more fundamental ways from the tasks we use today in our laboratory research on human-computer interaction. By far the largest amount of current research focuses on the behavior of people doing operational tasks: wordprocessing, spreadsheet formulation and

analysis, database search in support of constructing a report. Our current research focuses on office tasks.

The Space Station, in contrast, is likely to have very little need for operational tasks; standard everyday tasks are more likely to be accomplished by ground personnel. Space Station personnel are more likely to be involved in:

- the monitoring and control of onboard systems (e.g., life support, experiment/manufacturing control),
- the occasional use of planning and decision systems (e.g., expert systems for medical diagnosis or for planning for changes in the mission), and
- The nearly constant use of communication systems (i.e., for both mission related information and for personal contact with friends and family), for both synchronous conversation and asynchronous messages.

ISSUES

There are important research issues that are common among these systems and the operational systems that we focus on today, but there are other, additional issues that are unique, requiring particular emphasis. The common issues, important to all future human-computer interaction, include:

1. How to design a system that is easy to learn and easy to use. One core feature of such a system is "consistency". Polson's paper makes the case for consistency--a detailed argument for the importance of specifically modeling the user's goals and the methods necessary to accomplish the goals with a particular system. This is a very important research approach that promises to give the right level of answers to questions about consistency that will arise in future designs.
2. A second core feature in making a system easy to learn and use involves a straightforward "mapping" between the way the user thinks about the task objects and actions and the way the system requires the user to specify them. For example, the mapping between the objects of wordprocessing, such as letters, words, and sentences, correspond much more closely to the objects in a visual editor than they do to the strings and line objects of a line editor. Moran (1983) has made a beginning in delineating this type of analysis; more theoretical work and empirical verification is necessary.

3. How to make decisions about what modes of input/output (and their combinations) are appropriate for a given environment and task. Hayes' paper discusses a number of considerations that must be taken into account when deciding among speech/visual/keyboard input and output modalities, as well as the use of appropriate combinations of these modalities.
4. What characteristics of the human information processor are primary determinants of the range of acceptable interface designs. One way of evaluating a design of an interface is to analyze it on the basis of the major processing that a user engages in in order to understand the output and generate the next input. For example, we can analyze an interface for its perceptual clarity (e.g. adherence to Gestalt principles of grouping for meaning), its load on working memory (e.g., how many sub-goals or variables must be retained for short periods in order for users to accomplish their goals), and its requirements for recall from long-term memory (e.g., how many specific rules must be learned and how similar they are to each other). This approach, the cognitive science of human-computer interaction, by its generality across all application interfaces, promises to provide a theoretical thread through a number of empirical investigations. With a body of empirical tests of its predictions, this approach can both provide a robust base for future design situations and grow in sophistication and precision as a base for understanding complex cognition, even outside the domain of human-computer interaction.

Progress on these topics will make substantial contributions to our understanding of how to design human-computer interfaces for the Space Station in the year 2000, just as they will for those interfaces in offices and on the factory floor.

As discussed above, however, the systems on the Space Station are less likely to include operational systems, like those used in research on the above "common" topics, and more likely to include planning and decision, monitoring and control, and communication systems. Additional, important research issues arise in considering these latter three types of systems:

1. What characteristics of an interface appropriately alert users to abnormal situations in systems that must be monitored. What advice, information, or immediate training can be given users of a monitoring system that will guide them to behave in a creative but appropriate manner.
2. How are voice, video, keyboard, pointing devices, etc. to be used singly and in combination in each of these three types of systems? Certainly voice and video have begun to be explored in synchronous communication systems (e.g., picturephone and slow-scan video teleconferencing). How can these modalities be

used to best advantage to support the need for long-term contact with friends and family when individuals are separated for a long time? How are privacy issues accommodated in such systems, both for personal communication and operational communication?

3. If users have to consult an expert system or if some intelligence is incorporated into a system, how is information conveyed to the user about whether the system is to be believed? Since current intelligent systems are "fragile," that is, easily put in situations for which their advice is not appropriate, we need to convey to the user information about the system's boundaries of capabilities. Or, better yet, we need to build intelligent facilities that allow the user to query or access the stored knowledge in ways that can make the advice fit new situations more flexibly.
4. Since the systems that Space Station users must deal with will be varied and the users will have varying expertise in either the task at hand or the particular system to be used, it is important to have the system provide requisite context or training. Training need not be a formal module that one accesses explicitly, as software training modules are designed today. The systems could be initially designed to be transparent (i.e, with objects and actions that fit the way the user thinks about the task), not requiring training. Or, they could be built to include a "do what I mean" facility or embedded "help" or "training" facilities, accessible either when the user requests it or when the system detects that the user is confused or doing things inefficiently.
5. Most of the current theoretical bases for the design of human-computer interfaces consider tasks that are well-known to the user. The GOMS analysis of Card et al. (1983), for example, is for skilled cognition. Kieras and Polson's (1985) production system formalism similarly considers only skilled performance of cognitive tasks. However, in the Space Station environment, users will be doing few routine tasks. They will be doing tasks that involve novel situations, situations that invoke creative problem solving, not routine cognitive skill. Space Station personnel, for example, may try to alter a system that their monitor has shown is malfunctioning; they may use the advice of a medical expert system to attend to a colleague who has an undiagnosed illness; they may use communication channels to acquire additional expertise from the ground crews to solve onboard problems or plan new missions. In order to understand how these interfaces should be designed, more emphasis should be made in research in the area of human problem solving. The focus should be, for example, on how to build systems that, minimally, do not interfere with the information the person needs to keep track of during complex problem solving. Ideally,

we want to be able to build systems that augment a person's abilities to explore and evaluate new actions in novel situations.

6. Furthermore, as Hayes' paper points out, most of our current research on human-computer interaction focuses on the use of a system as a "tool" not as an "agent." Our understanding of cooperative human behavior is woefully thin. Theoretical bases need to be established so that we can build systems that cooperate well with the human problem solver, so that systems can augment the intelligent human to produce an even greater level of understanding and action.

APPROACHES FOR THE UNDEFINED FUTURE

As stated at the beginning of this discussion, the most difficult aspect of the task of listing research issues that the Space Station of year 2000 will benefit from concerns predicting the Space Station environment and the technology that will be available at the time. We just don't know what the alternative design elements will look like. The best we can do at this time, therefore, is to recommend a research agenda whose results promise to be useful no matter what the environment and technology will be. At the core of these recommendations is research that centers on the capabilities of the human information processor, both as an individual and in a cooperative environment. The human will not have changed substantially by the year 2000.

Consequently, our understanding of human-computer interaction will benefit from research that accumulates results from a common theoretical core that:

1. delineates in detail the functioning of the human information processor, with particular emphasis on the interaction among cognitive resources and those resources involved in attention (for monitoring systems), problem solving (for expert systems and decision support systems), and communication,
2. within the domain of expert systems, explores the information a user needs and determines how it should be presented so that the user can assess the believability of the advice given, and
3. determines ways to help casual users of a variety of systems to use them without a great deal of "start up" effort, either through transparent design; effective, easy training; or embedded intelligent aids.

A salient aspect of this type of research is that it is based on cognitive models, not on design principles. Cognitive models allow the examination of the interaction of features of the task or interface, which principles cannot do. These cognitive models characterize

details of what the task requires and details of the human information processor. By running these models, the designer or researcher can determine in detail areas of difficulty in the interaction (e.g., where the working memory is overloaded with subgoals and parameters to be retained). Certain changes to the interface design could be tested by running these models without having to invest in the expense of a full-fledged usability study. The number of researchers approaching issues in human-computer interaction with cognitive models is currently very small; their numbers should be encouraged to grow.

Furthermore, research should have as one of its goals the transfer of the knowledge developed in the laboratory to the design and development process. This calls for development of:

1. analytic tools for assessing consistency in a particular design.
2. analytic tools for assessing the amount of effort required in mapping the users' natural way of thinking about the task (i.e., an object/action language) into that required by the system, and
3. guidelines that will assist the designer in decisions about which modality or combinations of modalities are appropriate for a particular task and situation.

And, if systems are to be built for an evolving future, they must be built with scars and hooks, as Hayes notes. Software should be designed so that it has places that will allow easy growth in capabilities or input/output devices. Furthermore, research is needed to develop:

1. a method and language that allows the system designer to incorporate good human factors into the target system (e.g., a "toolkit" with components that have been designed with consideration for research on their human factors), and
2. a method that allows system developers to rapidly implement trial interfaces, so that they can be tested with real end-users, and then turned quickly into production code.

It is clear from the papers in this session that funds devoted only to simple empirical studies of users' behavior with new, increasingly complex technology will not be sufficient for answering the questions of the year 2000 and beyond. In contrast, research that focuses on:

1. the abilities of the human information processor with concomitant widespread, specific, robust cognitive modeling, and

2. additions to the development life cycle to make the production of good software rapid

can produce research that can make the human-computer interfaces on the Space Station of the highest possible quality for their time.

REFERENCES

- Card, S. K., Moran, T. P., and Newell, A.
1983 The Psychology of Human - Computer Interaction, Hillsdale, N. J.: Erlbaum.
- Kieras, D. E., and Polson, P. G.
1985 An approach to the formal analysis of user complexity. International Journal of Man - Machine Studies 22:365-394.
- Moran, T.
1983 Getting into a system: external-internal task mapping analysis. Pp. 45-49 in Proceedings of the 1983 CHI Conference on Human Factors in Computing. New York: ACM.

SYNOPSIS OF GENERAL AUDIENCE DISCUSSION

Two general points were raised from the floor.

1. When discussing natural language interfaces for human-computer interaction, one should make a clear separation between those requiring auditory input and those accepting natural language. Although these two features are highly correlated, they need not be. One could consider a speech input that would restrict language to a subset, such as single word commands or even special codes. Similarly, there could be natural language input that was entered via keyboard. Although there is an additional memory load imposed on the user if speech input accepted only a subset of natural language, there may be some applications that could effectively use this mode.
2. Allen Newell wished to emphasize the importance of having specific, detailed cognitive models as the basis for designing human-computer interfaces. The current researchers who are using this approach is very small, and though growing exponentially, the growth rate is very "leisurely." The approach has the advantage of not only specifying details of the processing mechanisms of cognition and their interaction, but also of specifying the details of the task the user is engaged in. Having the details of the task can provide benefits beyond redesign of the interface. They could serve as the basis from which the task itself could be redesigned, affording productivity enhancements from a straightforward efficiency analysis. Newell recommended a strong incentive be established for researchers to conduct their work in the context of cumulative, model-based theories of cognition, and let the design principles fall from them.

SESSION IV
COMPUTER AIDED MONITORING AND DECISION MAKING

ROBUSTNESS AND TRANSPARENCY IN INTELLIGENT SYSTEMS

Randall Davis

INTRODUCTION

Developing and building a space station will confront problems of significant complexity in an extraordinarily demanding environment. The station's size and complexity will make necessary the extensive use of automation for monitoring and control of critical subsystems, such as life support. The station complexity, along with the novelty of space as an environment, means that all contingencies cannot be anticipated. Yet the hostility of the environment means the consequences of failure can be substantial.

In such situations, robustness and transparency become essential properties of the systems we develop. A system is robust to the degree that it has the ability to deal with unanticipated events. A system is transparent to the degree that its operation can be made comprehensible to an observer.

This paper is concerned with these two properties--robustness and transparency--from a number of perspectives. We claim that they are crucial to the space station undertaking (and indeed to any situation with similar levels of complexity and similar consequences of failure). We argue that they are fundamental properties of models and system designs based on those models. As a result, robustness and transparency cannot easily be grafted on afterward; they must be considered at the outset and designed in. We explore how this might happen, i.e., how these two properties translate into constraints on system design and describe a number of research efforts that may lead to better understanding of how such design might be accomplished.

It is useful at this point to establish some simple vocabulary. By "system" or "device" we mean the hardware whose behavior we wish to understand and control. The power distribution system, for example, would include all the cables, batteries, fuel cells, solar arrays, switches, etc., that supply power to the station. By "model" we mean a description of that hardware that will allow us to analyze, interpret, diagnose, and guide its behavior. The model may be implicit in a program designed to monitor the hardware or it may exist in the mind of the human doing the same job. When expressed explicitly, it is typically written in terms of schematics, performance curves, engineering drawings, etc. The model also may be implicit in a program

designed to monitor the hardware or it may exist in the mind of the human doing the same job. In any case it provides the basic framework used to understand the device.

While we speak broadly of systems and models, our concern here is for the most part with systems of physical devices and the associated engineering models of them; much of what we say is likely to carry over to software as well. Models of human behavior and social systems are largely beyond what we attempt to do here.

Unanticipated Events: Motivation

Because much of what we discuss is motivated by the difficulties of dealing with unanticipated events, it is worth taking a moment to consider what they are and why they are important. By unanticipated events we mean any occurrence requiring a response that has not been previously planned for, analyzed, and the appropriate response determined.

One compelling example might occur if the life support system monitors present a collection of readings that indicate a malfunction but do not match any known pattern of misbehavior. The readings need to be analyzed and an appropriate response initiated, yet this cannot be done "by the book;" it requires that we reason through what could have happened to produce such readings.

The importance of such events arises from their inevitability, due to both the complexity of the space station and the novelty of the environment. Unanticipated events and interactions are a fact of life for complex, large scale systems because the number of different kinds of things that can go wrong is so vast, and our ability to do exhaustive formal analyses of fault events has rather modest limits. Space is a sufficiently novel environment that we have no comprehensive catalog of standard fault models that can be checked ahead of time.

Unanticipated Events: Example

During STS-2, the second space shuttle mission, an interesting sequence of events lead at one point to the recognition that a fuel cell was failing and later to the realization that in its degraded state it could conceivably explode. This sequence of events helps to illustrate both the inevitability of unanticipated events and the kinds of knowledge and reasoning needed to deal with them.

Some brief background will help make the events comprehensible. The basic function of the 3 fuel cells (Figure 1) is to produce electricity by combining hydrogen and oxygen in a carefully controlled reaction using potassium hydroxide as a catalyst. The combustion product is water, removed from the cell by the water removal system (Figure 2): damp hydrogen enters the condenser at the right, pulled along by the flow produced by the motor and pump at left. The motor is also turning a separator that pushes condensed water droplets toward the walls of the chamber where they accumulate due to surface tension (recall this

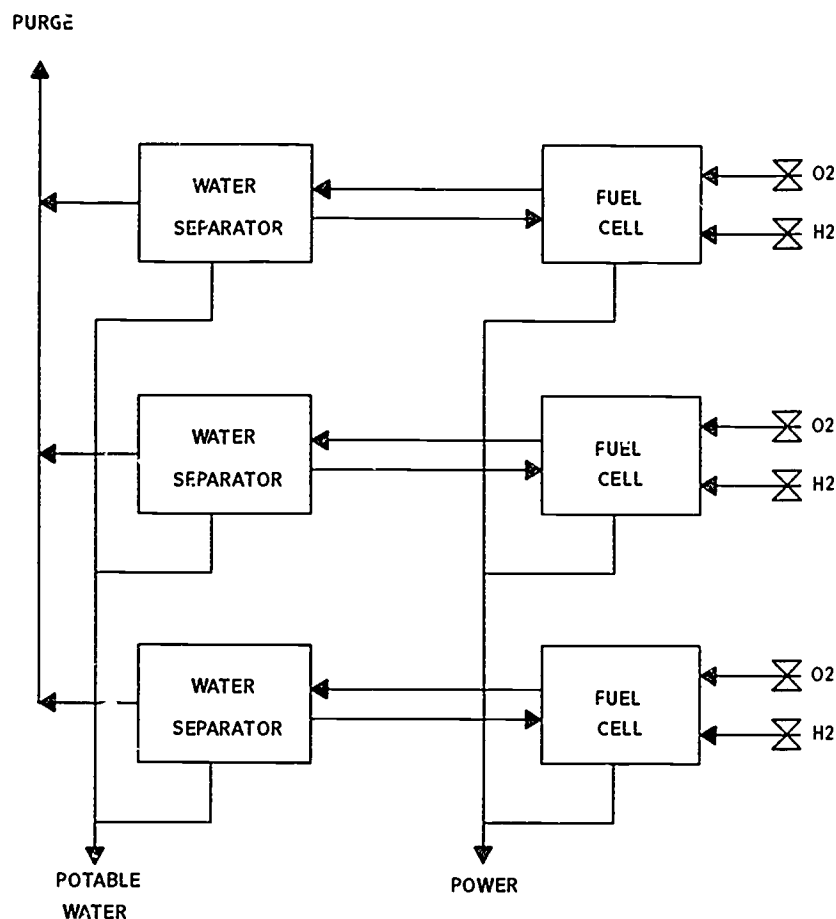


FIGURE 1 The fuel cell and water separation system.

is a Og environment). The now drier hydrogen returns to the fuel cell, while the annulus of water continually being formed at the separator is picked up and guided to the water storage area. A meter at the outlet monitors water pH, checking for contamination (e.g., potassium hydroxide from the fuel cell), since the water is intended for consumption.

In very much abbreviated form, the sequence of events leading to early mission termination of STS-2 proceeded as follows (Eichoefer, 1985):

Pre-Launch:

During pre-launch activities, the fuel cell pH meters register high.
Interpretation: Familiar, unexplained anomaly.

Pre-Launch:

At various times oxygen and hydrogen flow meters read high; at one point oxygen flow goes off-scale.
Interpretation: Sensors malfunctioning.




FIGURE 2 Details of the water separation unit. Source: Gerald Eichhoefer (July 1985).

- + 3:00 Fuel cell 1 (FC1) begins to shed load; the other two assume more load.
Interpretation: Cell may be failing.
Controllers consider purging FC1. Degraded performance suggests possible flooding; pH high also suggests flooding; purging will remove water. Purging FC1 rejected--purged KOH might solidify, blocking purge line that is common to all 3 cells.
- + 3:25 Crew asked to test pH manually. If sensor is correct, potable water may be getting contaminated by KOH.
- + 4:25 Crew too busy with other duties to perform test.
- + 4:40 FC1 off loads significantly
Interpretation: Clear failure.
- + 4:51 FC1 isolated from remainder of electrical system and shut down.

- + 5:48 Mission evaluation room recognizes new failure mode for the cell in the current situation. Once it is shut down pressure slowly drops, but can drop at different rates on each side. If pressure differential becomes large enough, gas hubbles from one side can cross to the other, possibly combining explosively.
- + 7:52 FCl restarted with reactant valves closed; reactants consumed and voltage in cell drops to 0.

Post-mission analysis of the fuel cell and water separator revealed that the pH meter had been working correctly and that a small particle blocked the nozzle in the water separator of cell 1, preventing water removal to the storage area. The water backed up first in the separator and later in the cell, flooding the cell (hence the high pH), leading to performance degradation, consequent load shedding, and eventual failure.

Lessons From The Example

This example is useful for a number of reasons. It illustrates, first, robustness and transparency in the face of unanticipated events. The reasoning was robust in the sense that the blockage had not previously been anticipated, yet engineers were able to reason through how the device worked, and were able to recognize and predict a novel sequence of potentially serious consequences. The reasoning was transparent in the sense that the story above is comprehensible. Even given the very small amount of information in Figures 1 and 2 and the short description above, the description of the events "makes sense."

Second, it suggests the difficulty of a prior identification and analysis of all failure modes and all the ways those failures may combine. Even with all the careful design, testing, and previously experience with fuel cell technology, a new mode of cell failure was encountered.

Third, it illustrates the kind of knowledge and reasoning that was required to understand, diagnose, and repair the problem. The knowledge involved information about structure (interconnection of parts) and behavior (the function of a component labeled "motor" or "pump"), supplied by the diagrams in Figures 1 and 2. Knowledge of basic chemistry and physics was also involved, used to understand the behavior potassium hydroxide in solution and the notion of surface tension. Importantly, the reasoning relies on causal models, descriptions of devices and processes that capture our ordinary notion of what it means for one event to cause another (e.g., the motor causes the pump to turn which causes the hydrogen and water to move through the condenser, etc.).

The reasoning involved was of several varieties., The fourth event above, for instance, illustrates reasoning about behavior to predict consequences: if the cell is flooded, potassium hydroxide can get in the water, meaning it can get to the water separator and then into the

water storage. Another form of reasoning involved working from observed symptoms to diagnoses and then to repair actions: If FCI is shedding load, it's an indication of degraded performance, which suggests flooding. Flooding in turn suggests purging as a repair. Simple knowledge of connectivity and chemistry ruled out that action in the event above at + 3:00: it might have blocked the common purge line.

Finally, it offers a simple way summarizing much of what this paper is about: while all of the reasoning above was done by people using their models of the devices in question, we suggest giving computers exactly the same sort of knowledge and reasoning abilities. They could, as a result, perform as far more effective assistants.

We believe this can be done by supplying them with something like the diagrams of Figures 1 and 2, with knowledge about structure, behavior, an understanding of causality, chemistry, physics, electronics, and more. In short, we need to give them the same understanding of "how things work" that we use in everyday engineering reasoning.

The aspiration, of course, is easy, execution is considerably more difficult; this is clearly no small undertaking. In the remainder of this paper, we examine some of the research issues that arise in attempting to make this happen.

- How can we provide descriptions usable by a machine that are equally as rich as those in Figures 1 and 2? Consider, for example, how much knowledge is captured by the simple labels motor, pump, and condenser.
- How can we provide the kinds of reasoning abilities displayed above?
- How can we provide the ability to judiciously select the correct model for a given problem? Consider how our view shifted from one grounded in physics, to one oriented towards chemistry, to one grounded in electronics, as the need arose.
- How can we provide the ability to simplify a complex model, selecting out just the "relevant" details? Consider what a drastic, yet useful, simplification Figures 1 and 2 are of the actual devices. (Consider too what a misleading statement it was, above, to say "Even given the very small amount of information in Figures 1 and 2 ..., the description of the events makes sense." It makes sense precisely because the right level of detail was chosen. How might we get a machine to do that?)
- For that matter, how do human engineers do all these things?

Unanticipated Events As A Focus

Unanticipated events like the blockage of the water separator are an appropriate focus for this paper because this symposium aims to identify research issues for future attention rather than incremental improvement to current practice. Some useful techniques already exist for simulation, fault insertion, and creation of error recovery procedures for foreseeable events. Additional work is in progress on techniques for error avoidance and in designing systems that are error tolerant. There is also a well-established approach to producing robustness through man-machine combinations: divide the work so that the more routine tasks fall to the machine and rely on the human for resourceful responses to atypical events. All of these are appropriate, important, and will continue to contribute to system design.

But new research issues arise in part by asking what relevant things we don't know how to do very well, or at all. From that perspective, unanticipated events present a set of interesting and important challenges, providing an appropriate focus for this paper.

They also lead to increased concern about transparency. Other rationales already exist for transparency, including giving users an understanding of the system's reasoning so they know when to rely on the conclusions, and the importance of keeping the system accessible to human comprehension and possible intervention. Dealing with unanticipated events adds additional motivation, most visible in the question of system override: to determine whether a system's response is based on inappropriate assumptions (e.g., an inappropriate model), we need first to know what those assumptions are. Transparency helps make this possible.

Agenda

Our discussion now proceeds in three basic steps. First, to help make clear the difficulties involved in robustness, we explore briefly some non-solutions to the problem. Second, we identify two broad categories of attack that are likely to offer some leverage on the problem: developing models and reasoning methods powerful enough to handle unanticipated events, and developing techniques for coping with situations where only imperfect models are available. Finally, we describe a number of specific research topics that will help to develop the models, methods and techniques needed to produce robustness and transparency.

SOME NON-SOLUTIONS TO THE PROBLEM

Before proposing a new attack on a problem, it's worth asking whether the problem can be tackled with known techniques. We consider three plausible approaches and explore why each of them fails to provide the degree of robustness we believe is necessary.


One traditional approach is the use of man-machine combinations, relying on the human to handle non-routine situations. This is, of course, useful and can be quite effective over a range of problems. In the fuel cell problem of STS-2, for instance, routine monitoring was handled automatically, while exceptions were analyzed by human experts.

It is also clear, however, that systems currently being designed and used are sufficiently complex that this will no longer be sufficient, unless we can make our automated assistants smarter. Some nuclear power and chemical processing plants, for instance, are complex enough that non-routine events lead to massive overload on human information handling abilities. So many alarms were triggered during the Three Mile Island accident, for instance, that not only was it effectively impossible to interpret them, even detection became problematic as multiple alarms masked one another. Somewhat more immediately relevant, during shuttle mission STS-9 an alarm was triggered more than 250,000 over 3 days, due to an unanticipated thermal sensitivity in a Spacelab remote acquisition unit, along with an oversight in user software.

It is likely that similar and perhaps higher levels of complexity will be involved in the space station. As a result, we need to do more than rely on the human half of the team to handle all exceptions. We need to upgrade the ability of our machines to interpret, diagnose, and respond to unanticipated events, enabling man-machine combinations to remain effective in the face of complex systems and novel environments.

A second route of attack on the problem might appear to be the creation of more reliable software through improved software engineering, program verification, or automatic programming¹. Unfortunately all of these solve a problem different from the one at hand here. The issue is illustrated in the figure below: techniques for production of reliable software all assist in ensuring that a program matches its specifications. Unanticipated events, however, will by definition not show up in the specifications. The problem here is not so much one of debugging code, it is the creation and debugging of the model and specifications.

Finally, given its wide popularity, we might ask what expert system technology² might be able to contribute to the difficulties we face. Here too the answer is that they have little to offer. The fundamental limitation in these systems arises from the character of the knowledge they use. Traditional expert systems gain their power by collecting empirical associations, if-then rules that capture the inferences human experts have learned through experience. We refer to them as empirical



associations to indicate the character of the knowledge they capture—associations, typically between symptoms and diseases, gathered as a result of human experience.

Importantly, those associations are typically heuristic rather than causal; i.e., they capture what experts have observed to happen without necessarily being able to explain why it should be so. A medical diagnosis system, for example, might have a rule of the form "a college student complaining of fatigue, fever, and sore throat is likely to have mononucleosis." The rule offers useful guidance even if the experts cannot provide a detailed causal (i.e., physiological) explanation for why the conclusion follows. Indeed the power of the technology comes in part from the assistance it provides in accumulating large numbers of fragmentary rules of thumb for tasks for which no well-defined causal theory exists.

One important consequence of this kind of knowledge, however, is a kind of brittleness. Current generation expert systems are idiots savant, providing impressive performance on narrowly defined tasks and performing well when the problem is exactly suited to the program's expertise. But performance can degrade quite sharply with even small variations in problem character. In general the difficulty arises from a lack of underlying theory: since the rules indicate only what conclusions follow and not why, the program has no means of dealing with cases that "almost" match the rule, or cases that appear to be "minor" exceptions. Indeed, they have no notion of what "almost" or "minor" could mean.

"FIGURING IT OUT"

Having reviewed some existing technology that does not appear capable of providing the degree of robustness needed, we turn now to considering what kinds of ideas and technologies would help solve the problem.

The basic thrust of our argument is quite simple. As size and complexity of systems increase, we see a decrease in the opportunity to do an exhaustive a priori analysis and pre-specify appropriate responses. The space station will likely be complex enough to preclude such analysis; the novelty of the environment increases the chance of unanticipated challenges.

To deal with such situations we need a new approach to building intelligent systems, one based on a simple premise: when you can't say in advance what will happen, the ability to "figure out" how to respond becomes much more important. Where knowledge-based systems, for instance, "know" what to do because they have been given a large body of task-specific heuristics, we require intelligent systems capable of figuring out what to do.

This ability should play a supporting role and is clearly not a replacement for existing approaches. Where we can anticipate and analyze of course we should, and where we can construct effective fault tolerant systems we should. But as system complexity grows and the number and seriousness of unanticipated events increases, we need the

flexibility and breadth of robust problem solving systems to deal with them.

The key question, of course, is how to construct systems with this property. In the remainder of this paper we suggest several ways of looking for answers to that question.

MODELS AND ENGINEERING PROBLEM SOLVING

Faced with an unanticipated event in a complex system, a powerful way to figure out what to do is by reasoning from an understanding of the system, a model of "how it works." A behavioral model, for instance, can be of considerable help in dealing with complex software like an operating system. In dealing with a complex physical device, a model of structure and function (schematics and block diagrams), along with an understanding of causality can be essential in understanding, interpreting and debugging behavior³.

How might we proceed, for example, when faced, with a set of sensor readings from the fuel cells that indicate malfunction but do not match any known pattern of misbehavior? The most robust solution appears to be grounded in knowing how it works, i.e., creating and using models that capture structure, behavior, and causality at an appropriate level of detail. We need to know what the component pieces are, how they each work, how they are interconnected, and so forth.

We argue that, in the most general terms, the creation, selection, and use of appropriate models is the most powerful approach to the problem⁴. It is in many ways the essence of engineering problem solving. Since, as we discuss in more detail below, models are abstractions, the process of model creation and selection is essentially one of deciding which abstraction to apply. Faced with a complex system to be analyzed, an engineer can bring to bear a powerful collection of approximations and abstractions.

As a relatively simple example in electrical engineering, for instance, an engineer may decide to view a circuit as digital or analog, linear or non-linear. But even to approach the problem as one of circuit theory means we have made the more basic assumption that we can model the circuit as if signals propagated instantaneously, and hence ignore electrodynamic effects. Models and their underlying abstractions are thus ubiquitous in this kind of problem solving.

We believe that an important source of power in the problem solving of a good engineer is the ability to create, select, use, and understand the limits of applicability of such models. Consequently, we believe that a powerful approach to building robust problem solving programs is to identify and capture the knowledge on which that modeling ability is based. Similarly, a powerful approach to building transparent problem solving problems is to make that knowledge explicit in our programs. One general thrust of the research we suggest is thus broadly concerned with advancing our understanding of model creation, selection, and use, and demonstrating that understanding by creating programs capable of doing such things.

A second general thrust is made feasible by the fact that the space station is an engineered artifact, a device intended to accomplish a specific purpose whose design is under our control. As a result, we can also ask, how can we design in such a fashion that dealing with unanticipated events is easier? That is, given the inevitability of encountering such events and the difficulty of reasoning about them in complex systems, how should we design so that the reasoning and analysis task becomes easier? We speculate, for instance, about what "design for comprehensibility" might mean.

Other approaches we discuss that share the same basic mind set include understanding (and hence capturing in programs) "common sense" physical reasoning, and exploring the origins of robust problem solving in people, whose graceful degradation in performance is so markedly different from the behavior of automated systems.

We refer to this set of approaches as "making the best situation" because they have in common the assumption that it is in fact possible to model the system and approach the problem by asking how we can facilitate model creation and use.

But what about the alternative? how can we get robust behavior in situations where no effective model yet exists, in situations where the only available models are incomplete or insufficiently detailed for the task at hand? We term that set of alternatives "making the best of the situation," to suggest that, lacking a model to reason from, we have to fall back on some less powerful methods. In this we speculate very briefly about research in using multiple, overlapping but incomplete models.

MODELS AND PROGRAMS

Since much of our discussion is focused on models--creating them, using them, and determining their limitations--it is worth taking a moment to review briefly some of their fundamental properties. Since we will for the most part be concerned with embodying those models in computer programs, it is similarly worth reviewing briefly the relation between models and programs, understanding the role the computer plays in all this.

The Role of the Computer

Let's start with the role of the computer. Given the size and complexity of the space station, extensive use will have to be made of software to automate tasks like monitoring and control. Any such program inevitably embodies a model of the task at hand. Even a program as simple as one that monitors CO₂ and displays a warning when the level exceeds a threshold has, implicit in it, a much simplified model of the sensing device, the environment (e.g., that CO₂ is uniformly dispersed), what levels of CO₂ are safe, etc. Since models and computer programs are often so closely intertwined, it

is important to understand what the model can contribute and what the computer can contribute.

The computer brings to the task a number of useful properties. It offers, for example, a vast increase in information processing power. This power, in turn, makes possible the construction and use of models that are orders of magnitude larger than any we could create by hand. The power is useful even with simple models, where it makes possible determining less obvious consequences, as in cases when straightforward search in chess can determine the long-term consequences of a move.

The computer has also facilitated the construction of many different kinds of models, including those that are non-numeric. As a result of work in computer science generally and AI in particular, we now routinely build and compute with models that are symbolic, qualitative, and incomplete. Symbolic models embody non-numeric inferences (e.g., "if the current shuttle pilot is Joe, on screen 1 display elapsed time, Houston time, and fuel levels"). Qualitative models⁵ describe and reason about behavior using the language of derivatives commonly employed by engineers (e.g., "if the voltage at node N3 increases then rate of discharge of capacitor C4 will decrease"). Most current expert systems are based on models that are incomplete, in the sense that they cover a number of specific cases (e.g., "if valve V3 is open and the tank temperature is high, then close valve V4"), but may leave unspecified what action to take in other cases (e.g., what to do if V3 is closed).

Work in AI and cognitive science has facilitated understanding and capturing other types of models as well, including mental models⁶, the vastly simplified, occasionally inaccurate but effective representations of mechanism and causality that people use in dealing with the world. My mental model of how the telephone system works, for instance, is quite a bit different from reality, but quite useful.

The computer also brings to the table a strong degree of "mental hygiene." Models expressed in English and left to human interpretation produce a notoriously wide variety of conflicting results. The remarkably literal-minded character of computer-based models enforces a degree of precision that we might not otherwise achieve in areas outside of those handled with formal mathematical analysis.

Expressing a model in a program also makes it far easier to test it by example, since determining its predictions is a matter of running the program rather than working out the consequences by hand. This in turn facilitates finding ambiguities, oversights, and limitations, and thus aids in extending the model.

All of these are useful and important properties. But for our purposes even more important is what the computer doesn't bring to the task, what embodying the model in a program does not do. It does not by itself provide either robustness or transparency. Simply put, robustness and transparency are properties of models and systems, not properties of programs that may be monitoring or controlling those systems.

This single observation has two important consequences. First, it means that we cannot create robust or transparent systems simply by developing software. It will not do, for instance, to design a highly

complex system and then develop an equally complex piece of software that attempts to monitor, interpret, and perhaps control it. Layers of complexity will only make it more difficult to deal with novel situations.

Perhaps the simplest demonstration of the futility of this approach comes in dealing with events that may be outside the range of applicability of the program. The more complex the underlying system, the more complex the program needed to interpret it, i.e., the more complex the model of that system needs to be. And the more complex the model is, the more difficult it becomes to determine whether it is based on assumptions that do not hold for the current situation, and hence the current events are outside its range of applicability.

Second, if robustness and transparency are properties of models and systems, not properties of programs, it follows that they cannot be grafted on, they must be designed in. That is, we need to understand how to design in such a fashion that the resulting systems have those properties, and how to create models that have those properties. One of the research strategies we suggest in this paper is to turn this question around, and ask how the desire for systems with these two properties can be translated into constraints on system design. That is, is it possible to design in such a way that the resulting systems are easy to model robustly and transparently.

Robustness and Transparency in Models

We have argued that robustness and transparency are properties of systems and models rather than of programs and that a primary route to resourceful systems is the creation of models with these properties. But that isn't easy. To see why not, we examine the kinds of things that commonly get in the way.

Three common sources of failures of robustness are incompleteness, information overload, and incorrect level of detail. Models may be incomplete because information that should have been included was omitted. A particularly relevant example arose in the Solar Max repair during Mission 41-C. The initial attempt to attach to the satellite failed because additional, undocumented hardware had been added to the satellite near the attachment point, preventing the mating of the satellite and the attachment device. The lesson here is the obvious one: you can't reliably figure out what to do if your picture of the device in question is incomplete.

A second source of failure of robustness--information overload--occurs when information processing ability available is overwhelmed by the amount of data or the size of model. The data rate may be so high that it cannot be interpreted fast enough. The model itself may be so large that it outstrips the processing power available. The issue here is the same for man or machine: in either case the available processing power may be insufficient to use the model. The lesson here is the need to ensure that the models we build are computable with the power available.

Information overload is frequently a result of the third common source of failure: selecting the wrong level of detail, in particular choosing too low a level. Attempting to model the behavior of a digital circuit using quantum mechanics might be an interesting challenge, but would surely drown in detail. If, on the other hand, too high a level is chosen, the model omits relevant phenomena. For example, some circuit designs that are correct when viewed at the digital level may in fact not work due to effects that are obvious only when viewed at the analog level.

All of this leads us to a fundamental difficulty in designing and using models. Robustness depends in large measure on completeness of the model. Yet all models are abstractions, simplifications of the thing being modeled, so no model can ever be entirely complete. Nor in fact would we want it to be. Much of the power of a model arises from its assumption that some things are "unimportant details," causing them to be omitted. There is power in this because it allows us to ignore some phenomena and concentrate on others; it is this license to omit some things that reduces the information processing requirements of using the model to within tolerable levels.

But there is as a result a fundamental tension between completeness (and attendant robustness) and complexity. If we make no simplifying assumptions we drown in detail; yet any simplifying assumption we make may turn out to be incorrect, rendering our model incomplete in some important way. This in turn raises interesting questions, further explored below, including how we select an appropriate model, i.e., an appropriate set of simplifying assumptions, and how we might recover in the event that we select one that is inappropriate.

RESEARCH TOPICS

In this section we discuss in broad terms a number of research topics relevant to the overall goal of building systems that are both robust and transparent. For the most part, we proceed from the assumption that getting machines to assist in significant ways with reasoning about situations like the STS-2 fuel cell problem will require that they have appropriate models. We then ask how those models can be created and indeed how we can design the device from the outset in such a way that the model creation process is made simpler.

Model Selection and Creation

Selecting and creating models is perhaps the most fundamental issue in solving engineering problems and an important determinant of the robustness of the solution. It is a skill that is in some ways well known: it's what good engineers have learned to do through years of experience. The goal here is to understand that skill and experience well enough that it can be embodied in a program, allowing automated assistance in selecting and creating appropriate models.

In almost any design or analysis problem, the most basic question is how to "think about" the object in question, i.e., how to model it. Given the acknowledgment that all models are abstractions, it is futile (and as we have suggested, inappropriate) to seek perfect completeness and robustness. That in turn means that the modeling decision concerns what to pay attention to, i.e., what properties of the object are relevant to the task at hand and which can safely be ignored. Hence the goal is to find a model with two properties. First it should be complete enough that it handles the important phenomena. Second it should be abstract enough that it is computable and capable of producing a description at a useful level of detail (i.e., even if it were possible, it would be of little use to produce a picosecond, microvolt-level analysis of a circuit whose digital behavior is of interest). But naming the goal is easy; the research challenge is in finding a more precise understanding of what it means to "consider the task" and to determine when a model is "complete enough", "abstract enough", and at an appropriate level of detail.

One possible route to understanding the nature and character of models is to define the kinds of abstractions commonly used in creating them. This might be done by determining what kinds of abstractions are commonly (and often implicitly) employed by engineers. What are the rest of the terms like digital, analog, linear, etc.? Is there just an unstructured collection of such terms or is there, as we would guess, some sort of organizing principle that can be used to establish an ordering on them? If so, we might be able to say more concretely what it means to proceed from a more abstract to a more precise model and might be able to develop programs capable of such behavior. It is unlikely that there is a simple, strict hierarchy that will allow us to move in a single, unambiguous direction. Much more likely we will find a tangled graph of models; part of the task is to sort out the different kinds of interconnections likely to be encountered.

A second possible route to understanding the nature of models arises from the simple observation that models ignore details. Perhaps then different kinds of models can be generated by selecting different combinations of details to ignore. The task here is to characterize different "kinds" of details; the ideal set of them would not only generate known models but might suggest additional models as well.

By either of these routes—studying the kinds of abstractions used or the kinds of details ignored—we might be able to produce an array of different kinds of models. That brings us to the problem of model selection, determining which to use in a particular situation. Some assistance may be provided by knowing how the array of models is organized, i.e., what it means to be a "different kind of model." The difficulty arises in determining what the important phenomena are in the problem at hand and selecting a variety of model capable of dealing with it. How is it that a human engineer knows which approximations are plausible and which are likely to lead to error?

It is unlikely that we will ever be able to guarantee that the knowledge used for model selection is flawless or that the models given to the program are flawless. We thus need to confront the problem of detecting and dealing with models that are inappropriately chosen for

the task at hand or that are incomplete in some relevant detail. Human engineers at times make the wrong selection or use a faulty model, yet are capable of detecting this and dealing with it. How might we get machines to do the same?

Finally, note that progress on model selection will have an important impact on the somewhat loaded issue of system override. If, as we have argued, unanticipated events are inevitable, simply having a detailed model is not enough: events may occur that are outside the range of applicability of the model. This can be a particularly difficult problem because it concerns deciding "how to think about" the problem.

We argue that override is fundamentally a decision that a particular model is inappropriate. Consider the example of a program monitoring and controlling life support. We might be tempted to override its decisions if they seem sufficiently different from our own, but why should they differ? The most basic answer seems to be that the model the program is using to interpret sensor readings is inappropriate, i.e., based on assumptions that are not valid in the current situation.

The only objective way to discover this is by determining why that model was chosen, what approximations it embodies, and what the limitations are on those approximations. Since much of this information was used to make the model selection to begin with, leverage on the override problem can come from understanding model selection and, importantly, from making explicit both the model itself and the assumptions underlying it. This would give us reasonably objective grounds for the override decision, since the model and its underlying assumptions will be available, and can be examined and compared to the current situation. It also reminds us how important it is that such information be made explicit, rather than left implicit in the program code or the mind of the program author.

Model Specification Needs To Be Less Trouble Than It Is Worth

We have repeatedly stressed the importance of models as a basis for robust reasoning about complex systems. But specifying those models is not an easy task, for several reasons. At the simplest level the issue is volume: there is an enormous amount of information to be captured. Existing design capture systems don't deal well with the problem because they don't make the information collection process easy enough, nor do they offer sufficient payoff once the information is entered to provide a motivation for doing it. They are in general more trouble than they're worth.

For design changes in particular, it is today often easier simply to try out the change and then (maybe) go back and update the specification database. In the case of Solar Max, for instance, perhaps no one knew about the additional hardware because it had been added at the last minute and never documented. The problem of documenting code is similar: it's often easier to try it out, then document. Often the documentation never gets done because it simply isn't viewed as critical to the undertaking.

The problem is both organizational and technical. Organizational issues arise because design documentation is typically of least use to the original designer, who is most familiar with the object. There should be a value structure within the organization that makes clear the importance of supplying complete design specifications and emphasizes that, as in Solar Max, the consequences of even minor omissions can be serious.

But there is a more radical position on this issue that is surely worth exploring. It ought to be impossible to create or modify a design without doing it via a design capture system. Put slightly differently, there should be a design capture system so useful that no one would think of proceeding without it. The thought is utopian but not so far afield as it might seem. Existing VLSI design tools, for example, providing sufficiently powerful functionality that no major design would be done without them. Even their basic functions--schematic capture and edit, design rule checking, simulation--provide sufficient payback to make them worth the trouble.

Existing tools also illustrate important limitations: they capture the final result, but not the rationales, not the design process. An effective system would be one that was useful from the earliest "sketch on the back of an envelope" stage, and that captured (and aided) every step and decision along the way. The result would be a record that included not only the final design, but its intended functionality, all rationales for the design choices, etc.

The technical problems in creating such a system include standard concerns about a good interface, such as ease of use and portability; paper is still hard to beat. But the issues go considerably deeper than that. Engineers find communication with each other possible in part because of a large shared vocabulary and base of experience. Communication with a design capture system should be based on similar knowledge; the identification and representation of that knowledge is a sizable research task.

The relevant vocabulary includes concepts about structure (shape, connectivity, etc.) and behavior (what the device should do). Both present interesting challenges. While connectivity is relatively straightforward, a compact and appropriate vocabulary for shape is not obvious. Behavior can sometimes be captured by equations or short segments of code, but descriptions in that form soon grow unwieldy and opaque. We need to develop a vocabulary for behavior capable of dealing with considerably more complex devices.

There is also the problem of unspoken assumptions. If design capture systems simply transcribe what is expressed literally, forcing every fact to be made explicit, the description task will always be overwhelming. We need to understand and accumulate the knowledge and design conventions of engineers so that the system can make the relevant inferences about what was intended, even if not expressed.

Designing for: Testability, Diagnosability, Analyzability, Comprehensibility, Transparency,...

We have argued that the complexity of the station and the novelty of the environment preclude an exhaustive a priori analysis of contingencies and require instead an ability to figure out what to do in the face of unanticipated events. We have suggested that this in turn is best facilitated by "knowing how things work," i.e., having a model of structure and behavior.

The complexity of the systems we design clearly has an impact on both how easy it will be to create such models and how easy it will be to reason with them once they exist. Since we are in fact designing the station (rather than trying to model a naturally occurring system), it is worth asking what can be done at the design state to facilitate model creation and model use.

Design for Testability Design for testability is one relatively well known approach in this category. It acknowledges that newly manufactured devices have to be exhaustively tested to verify their correct operation before they are placed in service and suggests that we design in ways that facilitate this task. Substantial effort has been devoted to this in circuit design, with some success. Given the likely need for equipment maintenance and the difficulty of a house (station?) call by service technicians, it will be useful to design the station in such a way that basic diagnostic tests can easily be run on devices that may be malfunctioning. Where well known concepts like ensuring that signals are observable and controllable are likely to carry over easily, part of the research task here lies in extending techniques developed for simple digital circuits to deal with much larger subsystems.

Design for Diagnosability Designs for diagnosability is a less well understood task. Where testing involves methodically trying out all of the designed behaviors of the device, diagnosis is a process of reasoning from the observed symptoms of malfunction to identify the possibly faulty components. Diagnostic power is measured in part by discrimination ability: more powerful diagnostic reasoning techniques implicate fewer components. But some problems are inherently ambiguous—a device may be designed in such a way that the observed symptoms must correctly implicate a large number of different components. Design for diagnosability would involve designing in a way that avoids this situation. Put more positively, it would mean designing in ways that seek to minimize the number of components implicated by a malfunction.

One very simple observation along this line can be made by considering the topology of the device: the only subcomponents that can be responsible for an observed symptom are those that are "causally connected" to it. In an electronic circuit, for example, the most obvious causal connections are provided by wires. More generally,

there must be some sequence of physical interactions by which the error propagates from its source to the point where it is observed. The fewer such interactions, the fewer candidate subcomponents. Simply put, this argues for "sparse (modular) designs," i.e., those with relatively few interconnections.

Designs with uni-directional components (i.e., those that operate in a single direction and have distinct inputs and outputs, like logic gates and unlike resistors), also have smaller candidate sets. In devices with uni-directional components there is a single direction of causality, giving us a notion of "upstream" and "downstream" of the symptom. Only components that are upstream can be responsible for the symptom.

Diagnosis also involves probing, i.e., taking additional measurements inside the device, as well as generating and running tests designed to distinguish among possible candidate subcomponents. We might also examine design styles that facilitate both of these tasks.

Designing for Analyzability, Comprehensibility, Transparency Given our emphasis on being able to figure out what to do, perhaps the most fundamental thing to do early on is what might be called design for analyzability or comprehensibility. If we have to think about how the device works and reason through the possibly subtle effects of an unanticipated event, then let's at least make that easy to do. This may be little more than the traditional admonition to "keep it simple," here given the additional motivation of on-the-spot analysis and response.

Simplicity in design will aid in making that easy; it may present additional virtues as well. Simplicity often produces transparency, an important component in people's willingness to accept automated assistance with critical tasks. Simplicity will help achieve NASA's design goal of allowing crews to intervene at low levels in any station subsystem.

Finally, simplicity may also produce robustness by assisting in determining when a model is inappropriate. We argued above that the override decision is part of the model selection process and could be facilitated by making explicit the simplifying assumptions underlying each model. Those assumptions might not always be specified completely, at times it may be necessary to determine what they are. This is likely to be easier to determine if the model itself can be analyzed easily.

Robustness Requires Common Sense

Current expert systems are brittle in part because they lack common sense knowledge, that large collection of simple facts about the world that is shared across a culture. At the simplest it may include facts such as physical objects have mass and take up space, that two things cannot occupy the same space at the same time, or that objects that are unsupported will fall. In the absence of such an underpinning of world

knowledge, the system must interpret its rules with complete literal mindedness and can do little in situations in which the rules "almost" apply.

Consider for example a rule in a medical diagnosis expert system specifying in part that "the patient is between 17 and 21 years old." Does the rule apply if the patient is 16 years 11 months old? How about 16 years 5.9 months? Our common sense knowledge of the world tells us that the human body doesn't change discontinuously, so the rule is probably still relevant. Compare this with a rule that says "If the postmark date is after April 15, then the tax return is late." Here we know, again from common sense knowledge, that there is in fact a discontinuity. Each of these chunks of common sense is simple enough and easily added to a system; the problem is finding and representing the vast collection of them necessary to support the kind of reasoning people do with so little effort.

For engineering problem solving of the sort relevant to our concerns here there is another layer of what we might call engineering common sense that includes such facts as, liquids are incompressible, all objects are affected by gravitational fields, but not all objects are affected by electromagnetic fields, electromagnetic fields can be shielded, and so forth. Engineers also know large numbers of simple facts about functionality, such as what a valve does, and why a door is like a valve.

The research task here is the identification, accumulation, organization, and interconnection of the vast numbers of simple facts that make up common sense (Lenat et al., 1986) and engineering common sense. Only with this body of knowledge will we be able to create systems that are more flexible and less literal minded.

What is the Source of Human Robustness?

Since robustness in problem solving is a common trait of experienced engineers, we ought to take the obvious step of examining that behavior and attempting to understand its origins. What is it that human experts do, what is it that they know, that allows them to recognize and deal with inadequate models? Why is it that human behavior seems to degrade gracefully as problems become more difficult, rather than precipitously, as is the case with our current programs? Part of the answer may lay in the number of and variety of models they can use, along with their body of common sense knowledge.

Multiple Models

Thus far our approach has focused on creating robustness by reasoning from detailed models. But how can we get robust behavior in situations where no effective model yet exists? One quite plausible reason for this would be incomplete information: even assuming we know all the limits of the models we have, selection of an appropriate one might

depend on a fact about the system or environment that we simply don't have yet. In this section, we speculate on one possible approach to such problems.

One idea explored to some degree in the HEARSAY system (Erman, et al., 1980) for speech understanding involves the use of multiple knowledge sources, each dealing with a slightly different body of knowledge. Our imperfect knowledge about the task--interpreting an utterance as a sentence--means that none of the knowledge sources can be guaranteed to be correct. The basic insight here is to employ a group of cooperating experts, each with a different expertise, in the hope that their individual weaknesses are distinct (and hence will in some sense be mutually compensated) but their strengths will be mutually reinforcing.

A similar technique might be useful in engineering problem solving: lacking any one model believed to be appropriate, we might try using a collection of them that appear to be plausible and that have somewhat different conditions of applicability. Even given such a collection, of course, there remains the interesting and difficult problem of deciding how to combine their results when the outcomes are (as expected) not identical.

SUMMARY

We have argued that the complexity of the station and the novelty of space as an environment makes it impossible to predict and analyze all contingencies in advance. The hostility of the environment means the consequences of failure are substantial. In such situations, robustness and transparency become essential properties of the systems developed. Systems are robust to the extent that they can deal with events that have not been specifically anticipated and analyzed. They are transparent to the extent that they can make their reasoning comprehensible to an observer.

Given the inevitability of unanticipated events, robustness is best accomplished by "figuring out" what to do, rather than relying on a list of predetermined responses. But "figuring out," the sort of analysis and reasoning routinely done by engineers, can only be done if you "know how it works," i.e., have a model of the device. We thus believe that a key source of power in engineering reasoning is the collection of models engineers use, along with the approximations and abstractions that underlie the models. One major thrust of research then should be directed toward understanding the processes of model creation, selection, and simplification.

Given the serious consequences of working from incomplete information, a second major thrust should be devoted toward model and design capture. Existing systems for VLSI design are effective enough to make them essential tools, and hence effective in some aspects of design capture. We need to provide similar levels of tools for all varieties of design and need to understand how to capture design rationales as well as the final result of the design process.

Given the difficulty of the reasoning process even with complete information, we suggest turning the question around and asking what we can do at design time to make the reasoning task easier. We have speculated about what design for testability, diagnosability, and comprehensibility might mean, and suggest further exploration there as well.

Finally, it appears that additional leverage on the problem is available from examining human performance to determine the source of robustness in our own problem solving behavior, and from compiling the large body of common sense knowledge that seems to be a source of graceful degradation in human problem solving.

ACKNOWLEDGMENTS

Support for the preparation of this report came in part from a research grant from Digital Equipment Corporation, from the Defense Advanced Research Projects Agency of the Department of Defense, under Office of Naval Research contract N00014-84-K-0124, and from a research grant from the Wang Corporation. This paper benefitted significantly from comments on early drafts by Walter Hamscher, Brian Williams, Reid Simmons and Dan Weld.

NOTES

1. Rich and Waters, eds., Artificial Intelligence and Software Engineering, Morgan Kaufmann, 1986, is a recent survey of attempts to use AI approaches to this problem. It provides a historical overview and a wide-range view of the problem with extensive references. Also see the IEEE Transactions on Software Engineering.
2. Davis, Buchanan, Shortliffe, Production rules as a representation, Artificial Intelligence, February 1977, Pp. 15-45, provides an early overview of MYCIN, the first purely rule-based expert system. Waterman, A Guide to Expert Systems, Addison Wesley, 1986, is a recent text oriented toward commercial applications of the technology and provides a large set of examples and references.
3. Bobrow, ed., Qualitative Reasoning About Physical Systems, North-Holland, 1984, is the book version of the December 1984 issue of Artificial Intelligence, a special issue on that topic. Nine articles illustrate the variety of models and tasks attacked, including diagnosis, design verification, behavior prediction, etc.

4. Relatively little work addresses this topic directly. Patil, Szolovits, and Schwartz, Causal understanding of patient illness in medical diagnosis, Proc Seventh Intl Jt Conf on AI, Pp. 893-899, explores the combined use of three different kinds of models in diagnostic reasoning. Hobbs, Granularity, Proc Ninth Intl Jt Conf on AI, Pp. 432-435 speculates on ways of producing coarser grained models from fine grained ones.
5. See the deKleer, Williams, and Forbus articles in Bobrow, op. cit.
6. See, for example, Gentner and Stevens, Mental Models, Lawrence Erlbaum, 1983.
7. Breuer, A methodology for the design of testable large-scale integrated circuits, Report SD-MR-85-33, January 1985, Space Division, Air Force Systems Command, provides a wide-ranging overview of different testability techniques.

REFERENCES

- Eichoefer, Gerald
1985 MITRE Corp. Report. (Adapted from July 16 report)
- Lenat, and Prakash, and Shephard
1986 Using common sense to overcome brittleness. Pp. 65-85 in AI Magazine. Winter.
- Erman, and Hayes-Roth, and Lesser, and Reddy
1980 The hearsay-II speech understanding system: integrating knowledge to resolve uncertainty. Pp. 213-254 in Computing Surveys. June.

DECISION MAKING--AIDED AND UNAIDED

Baruch Fischhoff

INTRODUCTION

Decision making is part of most human activities, including the design, operation, and monitoring of space station missions. Decision making arises whenever people must choose between alternative courses of action. It includes both global decisions, such as choosing a station's basic configuration, and local decisions, such as choosing the best way to overcome a minor problem in executing an onboard experiment. Decision making becomes interesting and difficult when the choice is non-trivial, either because decision makers are unsure what outcomes the different courses of action will bring or because they are unsure what outcomes they want (e.g., what tradeoff to make between cost and reliability).

Much of science and engineering is devoted to facilitating such decision making, where possible even eliminating the need for it. A sign of good engineering management is that there be no uncertainty about the objectives of a project. A sign of advanced science is that there are proven solutions to many problems, showing how to choose actions whose outcomes are certain to achieve the chosen objectives. Where the science is less advanced, the hope is to routinize at least part of the decision-making process. For example, the techniques of cost-benefit analysis may make it possible to predict the economic consequences of a proposed mission with great confidence, even if those techniques cannot predict the mission's risks to lives and property or show how those risks should be weighed against its economic costs and benefits (Bentkover et al., 1985; Fischhoff et al., 1981). Or, current engineering knowledge may allow automation of at least those decisions where electronic sensors or human operators can be trusted to provide accurate initial conditions. Indeed, space travel would be impossible without extensive computer-controlled decision making for problems involving great computational complexity or time pressure (e.g., during launch).

An overriding goal of space science (and other applied sciences) is to expand both the range of problems having known solutions and the technological capability for deriving and activating those solutions without human intervention. In this pursuit, it is aided by concurrent efforts in other fields. Among them is cognitive science (broadly

defined), whose practitioners are attempting to diversify the kinds of problems that can be represented and solved by computer.¹

Yet, however far these developments progress, there will always be some decisions that are left entirely to human judgment and some elements of judgment in even the most automated decisions. For example, there is no formula for unambiguously determining which basic design configuration will prove best in all anticipated circumstances (much less unanticipated ones). Analogously, there is no proven way to select the best personnel for all possible tasks. When problems arise, during either planning or operation, judgment is typically needed to recognize that something is wrong and to diagnose what that something is. When alarms go off, judgment is needed to decide whether to trust them or the system that they mistrust. When no alarms go off, supervisory judgment is needed to decide whether things are, in fact, all right. However thorough training may be, each operator must continually worry about whether others have understood their (possibly ambiguous) situations correctly, and followed the appropriate instructions. When solutions are programmed, operators must wonder how good the programming is. When solutions are created, engineers must guess at how materials (and people) will perform in novel circumstances. Although these guesses can be aided and disciplined by scientific theories and engineering models, there is always some element of judgment in choosing and adapting those models, compounding the uncertainty due to gaps in the underlying science. Any change in one part of a system creates uncertainties regarding its effects on other system components. In all of these cases, wherever knowledge ends, judgment begins, even if it is the judgment of highly trained and motivated individuals (Fischhoff, 1987; McCormick, 1981; Perrow, 1984).

Understanding how good these judgments are is essential to knowing how much confidence to place in them and in the systems that depend on them. Understanding how those judgments are produced is essential to improving them, whether through training or judgmental aids. Such understanding is the goal of a loosely bounded interdisciplinary field known as behavioral decision theory. The "behavioral" is meant to distinguish it from the study of decision making in mainstream American economics, which rests on the metatheoretical assumption that people always optimize when they make decisions, in the sense of identifying the best possible course of action. Although plausible in some circumstances and essential for the invocation of economics' sophisticated mathematical tools, the assumption of optimization severely constrains the kinds of behavior that can be observed. It also leaves economics with the limited (if difficult) goal of discerning what desires people have succeeded in optimizing in their decisions. Behavioral decision theory is concerned with the conditions conducive to optimizing, the kinds of behavior that come in its stead, and the steps that can be taken to improve people's performance (Fischhoff et al., 1981; Kahneman et al., 1981; National Research Council, 1986; Schoemaker, 1983; von Winterfeldt and Edwards, 1986).

Research in this tradition draws on a variety of fields, including psychology, operations research, management science, philosophy, political science, and (some) economics. As it has relatively little

institutional structure, it might be best thought of as the conjunction of investigators with several shared assumptions. One is the concurrent pursuit of basic and applied knowledge, believing that they are mutually beneficial. A second is the willingness to take results from any field, if they seem useful. A third is interest in using the latest technology to advance and exploit the research. These are also the assumptions underlying this chapter, which attempts to identify the most promising and important research directions for aiding space station development. Because of the space station's role as a pioneer of advanced technology, such research, like the station itself, would have implications for a wide range of other applications.

The results of research in behavioral decision theory have shown a mixture of strengths and weaknesses in people's attempts to make decisions in complex and uncertain environments. These intuitive psychological processes pose constraints on the decision-making tasks that can be imposed on people and, hence, on the quality of the performance that can be expected from them. These processes also offer opportunities for decision aiding, by suggesting the kinds of help that people need and can accept. The following section provides a brief overview of this literature and points of access to it, couched in quite general terms. The next section considers some of the special features of decision-making in space station design and operation. The following three sections discuss the intellectual skills demanded by those features and the kinds of research and development needed to design and augment them. These properties are the needs: (a) to create an explicit model of the space station's operation, to be shared by those involved with it, as a basis for coordinating their distributed decision making, (b) to deal with imperfect systems, capable of responding in unpredictable ways, and (c) to manage novel situations. A concluding section discusses institutional issues in managing (and exploiting) such research, related efforts (or needs) in other domains, and the philosophy of science underlying this analysis.

SPACE STATION DECISIONS AND THEIR FACILITATION

Most prescriptive schemes for deliberative decision making (Behn and Vaupel, 1982; Raiffa 1968; von Winterfeldt and Edwards, 1986), showing how it should be done, call for performing something like the following four steps:

- a. Identify all possible courses of action (including, perhaps, inaction).
- b. Evaluate the attractiveness (or aversiveness) of the consequences that might arise if each course of action is adopted.
- c. Assess the likelihood of each consequence occurring (should each action be taken).

- d. Integrate all these considerations, using a defensible (i.e., rational) decision rule to select the best (i.e., optimal) action.

From this perspective, decisions are evaluated according to how well they take advantage of what was known at the time that they were made, vis-a-vis achieving the decision maker's objectives. They are not evaluated according to the desirability of the consequences that followed. Some decisions involve only undesirable options, while the uncertainty surrounding other decisions means that bad things will happen to some good choices.

The following is a partial list of decisions that might arise in the course of designing and operating a space station. Each offers a set of action alternatives. Each involves a set of consequences whose relative importance must be weighed. Each is surrounded by various uncertainties whose resolution would facilitate identifying the optimal course of action:

Deciding whether to override an automated system (or deciding what its current state actually is, given a set of indicators);

Deciding in advance how to respond to a potential emergency;

Deciding where to look for some vital information in a computerized database;

Deciding whether to proceed with an extravehicular operation when some noncritical, but desirable safety function is inoperative;

Deciding whether to replace a crew member having a transient medical problem (either when formulating general operational rules or when applying them at the time of a launch);

Deciding where to put critical pieces of equipment;

Deciding how to prioritize the projects of different clients, both in planning and in executing missions;

Deciding where to look first for the sources of apparent problems;

Deciding which ground crew actions deserve an extra double check;

Deciding whether the flight crew is up to an additional period in orbit;

Deciding what to do next in a novel manipulation task;

Deciding on the range of possible values for a parameter needed by a risk analysis of system reliability;

Deciding just how much safety will be increased by a design change, relying on a risk analysis to project its system-wide ramifications;

Deciding what to report to outsiders (e.g., journalists, politicians, providers of commercial payloads) about complex technical situations that they are ill-prepared to understand.

These decisions vary in many ways: who is making them, how much time is available to make them, what possibilities there are for recovering from mistakes, how great are the consequences of success and failure, what computational algorithms exist for deciding what to do, how bounded is the set of alternative actions, and where do the greatest uncertainties lie, in evaluating the importance of the consequences or in evaluating the possibilities for achieving them. What these decisions have in common is that some element of unaided human judgment is needed before an action is consummated, even if it is only the decision to allow an automated process to continue unmolested. Judgment is needed, in part, because there is some element of uniqueness in each decision, so that it cannot be resolved simply by the identification of a procedural rule (or set of rules) that has proven itself superior in past applications. The search for rules might be considered an exercise in problem solving. By contrast, decision making involves the intellectual integration of diverse considerations, applying a general purpose integrative rule intended to deal with novel situations and "get it right the first time." In "interesting" cases, decision making is complicated by uncertain facts (Wise, 1986), so that one cannot be assured of the outcome (and of which choice is superior), and of conflicting consequences, so that no choice is superior in all respects (and some tradeoffs must be made)².

As mentioned, the hope of behavioral decision theory is to discern basic psychological processes likely to recur wherever a particular kind of judgment is required. One hopes, for example, that people use their minds in somewhat similar ways when determining the probability that they know where a piece of information is located in a database and when determining the probability that they can tell when a anomalous meter reading represents a false alarm. If so, then similar treatments might facilitate performance in both settings³ (Fischhoff and MacGregor, 1986; Murphy and Winkler, 1984).

The need to make decisions in the face of incomplete knowledge is part of the human condition. It becomes a human factors problem (the topic of this volume) either when the decisions involve the design and operation of machines (broadly defined) or when machines are intended to aid decisions. Decisions about machines might be aided by collecting historical data regarding their performance, by having them provide diagnostic information about their current trustworthiness, by providing operators with training in how to evaluate trustworthiness

(and how to convert those evaluations into action), and by showing how to apply general organizational philosophies (e.g., safety first) to specific operating situations. Decision aiding by machines might be improved by enhancing the display of information that operators understand most poorly, by formatting these displays in ways compatible with users' natural ways of thinking, by clarifying the rationale for the machine's recommendations (e.g., its assumed tradeoffs, its decision rule, its treatment of uncertainty), and by describing the definitiveness of its recommendations. A better understanding of how people intuitively make decisions would facilitate attaining these objectives, as well as developing training procedures to help people make judgments and decisions wherever they arise. Just thinking about decision making as a general phenomenon might increase the motivation and opportunities for acquiring these skills.

DESCRIPTIONS OF DECISION MAKING

One way of reading the empirical literature on intuitive processes of judgment and decision making is as a litany of problems. At each of the four stages of decision making given above, investigators have identified seemingly robust and deleterious biases: When people generate action options, they often neglect alternatives that should be obvious and, moreover, are insensitive to the magnitude of their neglect. As a result, options that should command attention are out of mind when they are out of sight, leaving people with the impression that they have analyzed problems more thoroughly than is actually the case (Fischhoff et al., 1978; Pitz et al., 1980). Those options that are noted are often defined quite vaguely, making it difficult to evaluate them precisely, communicate them to others, follow them if they are adopted, or tell when circumstances have changed enough to justify rethinking the decision (Bentkover et al., 1985; Fischhoff et al., 1984; Furby and Fischhoff, 1987; Samet, 1975). Imprecision also makes it difficult to evaluate decisions in the light of subsequent experience, insofar as it is hard to reconstruct exactly what one was trying to do and why. That reconstruction is further complicated by hindsight bias, the tendency to exaggerate in hindsight what one knew in foresight (Fischhoff, 1975). The feeling that one knew all along what was going to happen can lead one to be unduly harsh on past decisions (if it was relatively obvious what was going to happen, then failure to select the best option must mean incompetence) and to be unduly optimistic about future decisions (by encouraging the feeling that things are generally well understood, even if they are not working out so well).

Even though evaluating the relative importance of potential consequences might seem to be the easiest of the four stages of decision making, a growing literature suggests that people are often uncertain about their own values. As a result, the values that they express can be unstable and unduly sensitive to seemingly irrelevant features of how evaluation questions are posed. For example, (a) the relative attractiveness of two gambles may depend on whether people are

asked how attractive each is or how much they would pay to play it (Grether and Plott, 1979; Slovic and Lichtenstein, 1983); (b) an insurance policy may become much less attractive when its "premium" is described as a "sure loss" (Hershey et al., 1982); (c) a risky venture may seem much more attractive when described in terms of the lives that will be saved by it, rather than in terms of the lives that will be lost (Kahneman and Tversky, 1979; Tversky and Kahneman, 1981). Thus, uncertainty about values can pose as serious a problem to effective decision making as can uncertainty about facts.

Although people are often willing to acknowledge uncertainty about what will happen, they are not always well equipped to deal with it, in the sense of assessing the likelihood of future events (in the third stage of decision making). A rough summary of the voluminous literature on this topic is that people are quite good at tracking repetitive aspects of their environment, but not as good at combining those observations with inferences about what they have not seen (Hasher and Zacks, 1984; Kahneman et al., 1982; Peterson and Beach, 1967). Thus, they might be able to tell how frequently they have seen or heard about deaths from a particular cause, but not be able to assess how representative their experience has been—leading them to overestimate risks to which they have been overexposed (Combs and Slovic, 1979; Tversky and Kahneman, 1973). They can tell what usually happens in a particular situation and recognize how a specific instance is special, yet have difficulty integrating these two (uncertain) facts—with the most common bias being to focus on the specific information and ignore experience (or "base rates") (Bar Hillel, 1980). They can tell how similar a specific instance is to a prototypical case, yet not how important similarity is for making predictions—usually relying on it too much (Bar Hillel, 1984; Kahneman and Tversky, 1972). They can tell how many times they have seen an effect follow a potential cause, yet not infer what that says about causality—often perceiving relations where none exist (Beyth-Marom, 1982; Einhorn and Hogarth, 1978; Shaklee and Tucker, 1980). They have a rough feeling for when they know more and when they know less, but not enough sensitivity to avoid a commonly observed tendency toward overconfidence (Fischhoff, 1982; Wallsten and Budescu, 1983).

According to decision theory, the final stage of decision making should involve implementation of an expectation rule, whereby an option is evaluated according to the attractiveness of its possible consequences, weighted by their probability of occurrence. Since it has become acceptable to question the descriptive validity of this rule, much research has looked at how well it predicts behavior (Dawes, 1979; Feather, 1982; Fischhoff et al., 1981; Kahneman et al., 1982; National Research Council, 1986; Schoemaker, 1983). A rough summary of this work would be that: (a) the expectation rule often predicts people's choices fairly well—if one knows how they evaluate the probability and attractiveness of consequences; (b) with enough ingenuity, one can usually find some set of beliefs (regarding the consequences) for which the rule would dictate choosing the option that was selected—meaning that it is hard to prove that the rule was not used; (c) expectation rules can often predict the outcome of

decision-making processes even when they do not at all reflect the thought processes involved--so that predicting behavior is not sufficient for understanding or aiding it; (d) those processes seem to rely on rules with quite different logics, many of which appear to be attempts to avoid making hard choices by finding some way to view the decision as an easy choice--for example, by disregarding consequences on which the otherwise-best option rates poorly (Janis and Mann, 1977; Montgomery, 1983; Payne, 1982; Simon, 1957).

The significance of these results from experimental studies depends upon how well they represent behavior outside the lab, how much insight they provide into improving decision making, and how adversely the problems that they reveal affect the optimality of decisions. As might be expected, there is no simple answer to any of these questions. Life poses a variety of decisions, some of which are sensitive to even modest imprecision in their formulation or in the estimation of their parameters, some of which yield an optimal choice with almost any sensible procedure, and some of which can tolerate occasional inaccuracies, but not recurrent problems, such as persistently exaggerating how much one knows (Henrion, 1980; Krzysztofowicz, 1983; McCormick, 1981; von Winterfeldt and Edwards, 1982). Placing decisions within a group or organizational context may ameliorate or exacerbate problems, depending on how carefully members scrutinize one another's decisions, how independent are the perspectives that they bring to that scrutiny, and whether that social context has an incentive structure that rewards effective decision making (as opposed to rewarding those who posture or routinely affirm common misconceptions) (Davis, 1982; Lanir, 1982; Myers and Lamm, 1976).

The robustness of laboratory results is an empirical question. Where evidence is available, it generally suggests that these judgmental problems are more than experimental artifacts, which can be removed by such "routine" measures as encouraging people to work harder, raising the stakes contingent on their performance, clarifying instructions, varying the subject matter of the tasks used in experiments, or using better educated subjects. There are many fewer studies than one would like regarding the judgmental performance of experts working in their own areas of expertise. What studies there are suggest some reason for concern, indicating that experts think like everyone else, unless they have had the conditions needed to acquire judgment as a learned skill (e.g., prompt, unambiguous feedback) (Fischhoff, 1982; Henrion and Fischhoff, 1986; Murphy and Winkler, 1984).

The evidentiary record is also incomplete with regard to the practical usefulness of this research. The identification of common problems points to places where human judgment should be supplanted or aided. The acceptance of decision aids (and aides) has, however, been somewhat limited (Brown, 1970; Fischhoff, 1980; Henrion and Morgan, 1985; von Winterfeldt and Edwards, 1986). One inherent obstacle is presenting users with advice derived by inferential processes different than their natural ones, leaving uncertainty about how far that advice is to be trusted and whose problem it really is solving. Developing (and testing) decision aids that took seriously the empirical results

of behavioral decision theory would be a useful research project. With regard to situations where decision aids are unavailable, there is some evidence that judgment can be improved by training procedures that recognize the strengths and weaknesses of people's intuitive thought processes (Kahneman et al., 1982; Nisbett et al., 1983). Here, too, further research is needed.

THE PSYCHOLOGICAL REALITY OF SPACE STATION DECISIONS

The recurrent demand for similar intellectual skills in diverse decisions means that any research into decision-making processes could, in principle, provide some benefit to the space station program. However, there are some conditions that are particularly important in the space station environment and, indeed, might rarely occur in less complex and technologically saturated ones. The challenges posed by such conditions would seem to be suitable and important foci for NASA-supported research. Three such conditions are described in the remainder of this section. Each subsequent section considers research issues pertinent to one of these conditions. In each case, significant progress appears possible, but would appear to demand the sort of sustained programmatic effort that NASA has historically been capable of mustering.

Condition 1: The need to create a widely shared model of the space station and its support systems. The technical knowledge needed to manage the space program is widely distributed over diverse locations on earth and in space, in different centers on earth, and across different people within each earth and space center. As a result, there are prodigious technical problems involved in ensuring compatibility, consistency, and concurrency among the computerized databases upon which these scattered individuals rely. Even if these problems of information transmission can be resolved, there is still no guarantee that the diverse individuals at the different nodes in the system will be aware of the information available to them, nor comprehend its meaning for their tasks, nor be alert to all changes that might affect their work. Even with a static database, there may be problems of understanding when the individuals have very different kinds of expertise, such that their contributions to the database cannot be readily understood (or evaluated) by one another.

The management of such systems requires the creation of some sort of system-wide model within which individuals can pool their knowledge and from which they can draw needed information. That model may be a loosely organized database, with perhaps a routing system for bringing certain information to the attention of certain people (attempting to strike a balance between telling them too much and too little). Or, it may be an explicit coordinated model, such as those used in design processes guided by procedures like probabilistic risk analysis (McCormick, 1981; U.S. Nuclear Regulatory Commission, 1983). These models assign new information into an integrated picture of the

physical system, possibly allowing computational predictions of system performance, which can be redone whenever the state of the system (or the theoretical understanding of its operation) changes. Shared models with such computational abilities can be used to simulate the system, for the sake of comparing the effects of design changes, training operators for emergencies, and troubleshooting (by seeing what changes in the system could have produced the observed aberrations). Such models are useful, if not essential, for achieving NASA's goal of allowing "crews to intervene at extremely low levels of every subsystem to repair failures and take advantage of discoveries" (NASA, 1986).

Less ambitious models include spreadsheets, status displays, even simple engineering drawings, pooling information from varied human and machine sources (although, ultimately, even machine-sourced information represents some humans' decisions regarding what information should and can be summarized, transmitted, and displayed). All such models are based around a somewhat artificial modeling "language" which is capable of representing certain aspects of complex systems. Using them effectively requires "fluency" in the modeling languages and an understanding of their limits. Thus, for example, decision analysis (Behn and Vaupel, 1982; Raiffa, 1968; von Winterfeldt and Edwards, 1986) can offer insight into most decision-making problems, if decision makers can describe their situations in terms of options, consequences, tradeoffs, and probabilities--and if they can recognize how the problem described in the model differs from their actual problem.

Probabilistic risk analyses can aid regulators and designers to understand the reliability of nuclear power plants by pooling the knowledge of diverse groups of engineers and operators--as long as everyone remembers that such models cannot capture phenomena such as the "intellectual common mode failure" that arises when operators misunderstand an emergency situation in the same way.

The creation, sharing, interpretation, and maintenance of such models are vital to those organizations that rely on them. The unique features of such models in the context of NASA's missions are their size and complexity, their diversity (in terms of the kinds of expertise that must be pooled), and their formality. That formality comes not only from the technical nature of much of the information but also from the need for efficient telecommunications among NASA's distributed centers. Formality complicates the cognitive task of communication, by eliminating the informal cues that people rely upon to understand one another and one another's work. It may, however, simplify the cognitive study of such communication by rendering a high portion of significant behavior readily observable. It may also simplify the cognitive engineering of more effective model building and sharing, insofar as better methods can be permanently and routinely incorporated in the appropriate protocols. Research that might produce such methods is discussed below.

Condition 2: The need to make decisions with imperfect systems.

Decisions involving uncertainty are gambles. Although it is an uncomfortable admission where human lives are at stake, many critical

decisions in space travel are gambles. The uncertainties in them come from the limits of scientific knowledge regarding exactly how various elements of a mission will perform, from the limits of engineering knowledge regarding how different system elements will interact, from the limits in the technical capacity for modeling complex systems, and from the unpredictability of human operators (who are capable of fouling and saving situations in novel ways). Indeed, despite NASA's deep commitment to planning and training, the nature of its mission demands that some level of uncertainty be maintained. It is expected to extend the limits of what people and machines can do. Performance at those limits cannot be tested fully in theoretical analyses and simulation exercises.

In order to gamble well, one needs both the best possible predictions regarding a system's performance and a clear appraisal of the limits of those predictions. Such an assessment of residual uncertainty is needed in order to guide the collection of additional information, in order to guide preparation for surprises, and, most important of all, to guide the decision as to whether a mission is safe enough to proceed (considering NASA's overall safety philosophy). Using information wisely requires an understanding of just how good it is.

Because gambling is so distasteful, there is constant activity to collect (and produce) additional knowledge, either to perfect the system or to clarify its imperfections. As a result, the state of knowledge and the state of the system will be in constant flux, even without the continual changes of state associated with its ongoing operations (e.g., testing, training, wear). Somehow, this new information must be collated and disseminated, so that those concerned with the system know what is happening and know how much one another knows. In this way, dealing with uncertainty is related to dealing with a shared model.

For operators, this residual uncertainty creates the constant possibility of having to override the system, in order to rescue it from some unanticipated circumstance or response. That override might involve anything from a mild course correction to a fundamental intervention signalling deep distrust of a system that seems on the verge of disaster. As the physical stakes riding on the decision increase, so do the social stakes (in the sense of the responsibility being taken for system operation and the implicit challenge to system designers). Thus, operators, as well as designers and managers, must be able to assess the system's trustworthiness and to translate that assessment into an appropriate decision.

The variety of individuals with knowledge that could, conceivably, prompt override decisions means that coping with uncertainty is an intellectual skill that needs to be cultivated and facilitated throughout the organization. It also means that the system's overall management philosophy must recognize and direct that skill. For example, a general instruction to "avoid all errors" implies that time and price are unimportant. Where this is not the case, personnel are left adrift, forced to make tradeoffs without explicit guidance. Such an official belief in the possibility of fault-free design may also

discourage the treatment of those faults that do remain. Many failsafe systems "work" only because the people in them have learned, by trial and error, to diagnose and respond to problems that are not supposed to happen. Because the existence of such unofficial intelligence has no place in the official design of the system, it may have to be hidden, may be unable to get needed resources (e.g., for record keeping or realistic exercises), and may be destroyed by any change in the system that invalidates operators' understanding of its intricacies. From this perspective, where perfection is impossible it may be advisable to abandon near-perfection as a goal as well, so as to ensure that there are enough problems for people to learn how to cope with them. Moreover, steps toward perfection should be very large before they could justify disrupting accustomed relationships. That is, technological instability can be a threat to system operation.

Condition 3: The need to make novel decisions, in non-routine situations. With nearly perfect systems, rare problems are always somewhat novel. Even when they have been anticipated and incorporated in contingency plans, there is always some uncertainty about whether the problems that arise can be identified with the comparable problems described in the plans. Where the plans can be retrieved, there is still some uncertainty about whether they will seem like the right thing to do once the contingency is confronted "in the flesh." The retrieval of plans is an exercise in pattern matching. However, it also involves a series of decisions regarding whether a contingency has arisen, which plan is meant to fit the current situation, and whether that plan is to be trusted.

Yet other decision problems will be entirely novel and unanticipated. Such situations might be considered the purest form of decision making, clearly calling for the integration of diverse pieces of information in an effort to identify the right course of action, often having to get it right the first time. Where time constraints are great, such decision making may involve just the raw exercise of intuitive thought processes. Raw intuition may also be the primary ingredient for more leisurely decisions, when there is no accepted structure for decision making. That may happen, for example, when problems fall at the intersection of several jurisdictions or when they require tradeoffs regarding which the organization lacks policy.

In such situations, decision making may be seen as involving several kinds of "research." These include understanding the interactions among subsystems previously thought to be relatively independent, discerning how the organization's underlying safety philosophy applies to a particular novel case, generating action options to evaluate, and ferreting shared misconceptions.

When there is an algorithmic procedure for deciding what to do, the novelty of a decision may lie in having to deal with a unique state of the physical system. Understanding that state requires more than the usual troubleshooting (i.e. diagnosing which of a known set of problems has produced the observed symptoms). Rather than that sort of (sophisticated) pattern matching, unique states require the equivalent

of on-line research. That research may involve short-term engineering analysis, using whatever aspects of the overall design model can be accessed within the time constraints. When formal models are inaccessible, then the analysis must be performed within the "mental models" of the decision makers and their aides. In either case, judgment is needed to choose the information-gathering procedures with the highest "yield," in terms of hypothesis testing.

In addition to the cognitive difficulties of making unique decisions, there may also be institutional difficulties to gaining support for unfamiliar actions based on interpretations of values and facts that are not explicitly part of organization's shared model. There not be the time needed for customary consensus-building efforts. There may not be clear recognition of the needed autonomy. There may be unusual exposure to being evaluated in the light of biased hindsight. There may be problems in coordinating the activities of those involved in implementing the decision. These difficulties affect the ability to anticipate the consequences of taking various actions, as well as decision makers' ability to take those actions that seem right to them.

RESEARCH NEEDS: CREATING A SHARED MODEL

The creation of explicit shared models demands several general intellectual skills. Each could be the source of problems and the object of research. Where procedures exist (or can be discovered) for enhancing those skills, there should be good opportunities to implement them widely (e.g., in the computer programs used for eliciting and presenting models). Something is known about the exercise of each of the skills.⁴ If the same skills recur in the creation of many kinds of models, then learning more about them could provide some generally useful knowledge. They are:

Skill 1: identifying and characterizing the key components of the system being modeled.

Skill 2: identifying and characterizing the interrelations between those components.

Skill 3: estimating quantitative model parameters.

Skill 4: evaluating the quality of the model.

In the case of a probabilistic risk analysis, exercise of the first skill would include determining which pieces of physical equipment (e.g., valves, controls, piping) are vital to system performance and describing them in sufficiently precise terms as to allow further analysis. The second skill includes determining which malfunctions in System X need to be considered when studying the performance of System Y, and what the functional form of their relationship is. The third skill might include determining the probable distribution of failure

rates for particular system components (e.g., valves, maintenance measures). The fourth skill involves actions such as determining the range of values to be used in sensitivity analyses, assessing the information yield of possible research activities, and determining how well the system is understood (as a prologue to deciding whether it is understood well enough for action to proceed).

Creating such engineering models can be seen as a special case of the general problem of eliciting information from experts. It differs from the perspective associated with what are usually called "expert systems." Here, the modeling language does not attempt to be a natural one. Rather, it is a flexible analytic language, capable of modeling a wide variety of situations and pooling the knowledge of diverse experts--if they can express themselves in the terms of the language. Thus, the core of the needed research programme is an examination of how people express their beliefs in the terms of abstract languages, and how they interpret the expressions of others' beliefs in the models that they share.

As with "expert systems," these models can help users understand (and communicate) the nature of their own expertise. Models force one to be explicit and allow one to simulate the effect of varying assumptions on model performance. However, if the language is awkward, or imprecise, or inconsistently interpreted, then users may not know what they are talking about. If the syntax is unintuitive, then users may not understand the implications of the relations that they have described. In such cases, expertise couched in terms of true natural languages, with their deep dependence on tacit knowledge, may not ensure expertise with the modeling language. There even may be a role for interpreters, helping experts express what they know in terms that the language can accept.

As a small example of the possibility of such difficulties, (Fischhoff et al., 1978) two groups of experienced garage mechanics were asked judge the completeness of tree-like graphic depictions of possible reasons why a car might not stop. One group judged a fairly complete tree, the second a tree from which major systems (e.g., battery, ignition) had been pruned. Even though the pruning removed systems judged to include approximately 50% of problems, the pruned tree was judged to be almost as complete as the full one. The (pruned) systems that were out of sight were effectively out of mind. Although these experts clearly knew about the missing systems, they had difficulty interpreting that knowledge in the terms of the model. Their expertise might have been better exploited by having them list specific instances of no-starts, rather than asking for direct estimates of completeness. A second set of examples lies in the research literatures documenting the difficulties that people have with testing hypotheses and discerning causal relations (Evans, 1982; Fischhoff and Beyth-Marom, 1983; Kahneman et al., 1982; Nisbett and Ross, 1980).

Understanding these properties of modeling languages is important to having realistic expectations from them. Improving people's fluency with them is critical to improving the quality of modeling and the ability of shared models to serve an organization's needs. From this

perspective, what is needed, in effect, is an understanding of engineering design as a cognitive and social process, focused on these explicit expressions of it.

Every modeling language (like every other language, presumably) is better at capturing some kinds of situations than others. For example, most engineering languages are ill-suited to describing the actions of humans within a technical system (Hollnagel et al., 1986; Rasmussen and Rouse, 1981); economic techniques, such as cost-benefit analysis, are ill-suited to treating goods that are not traded directly in an unrestrained market; military intelligence analyses have more of a place for quantitative, tactical information (e.g., about what the enemy has) than for qualitative, strategic information (e.g., about what the enemy really wants). Such situations leave users with the difficult task of integrating two qualitatively different kinds of information, differing in how readily they can be incorporated in the model. Research is needed into how to extend the range of modeling languages, and into how to help users deal systematically with those factors that are left out.

Once models have been created, they must be communicated, raising the question of who needs to know what. Some balance must be struck between telling too much and too little. One research approach to developing communication guidelines would come out of value-of-information analysis, asking what information effects the greatest difference in the expected value of the specific decisions that need to be made at different nodes (Raiffa, 1968). A complementary, cognitive approach would consider how broad and deep a picture people need to see in order to understand the interface between their own actions and those taken elsewhere. A third, more social approach would ask how people anticipate what others in the system know, so as to be able to interpret their actions (Gardenier, 1976; Metcalf, 1986).

After a model has been created, it must be updated, both as the system changes and as better information about it is received. Although the natural desire is always to be current, that can create problems of understanding and coordination. For example, with an evolving system, design changes that are introduced piecemeal may have system-wide ramifications that are never detected. Or, users may find it difficult to deal with a picture of the system that is never the same as when they last consulted it. Both of these kinds of problems might be ameliorated by relying instead on periodic model-wide updating, at the price of letting the model become increasingly out of date as the last revision becomes more distant in time. Presumably, these "cognitive concurrency" problems, and their recommended treatments, will vary with the nature of the system and the changes.

Better models (and better use of existing models) would directly produce some better decisions, in those situations where action follows directly from the analysis of the facts. In other cases, the facts do not speak for themselves, but must be considered in the light of organizational policies. In such cases, there may be some place for decision aiding. The shared model could attempt to identify relevant policies and extract their implications for particular decision

problems. To avoid the rejection that decision aids frequently have experienced, they would have to aid decisions without usurping decision-making responsibility. That calls, in part, for cognitive research (e.g., on how to display the assumptions and definitiveness of recommendations) and, in part, for social research (e.g., on how to justify aided decisions).

RESEARCH NEEDS: USING IMPERFECT SYSTEMS

The key to using imperfect systems is understanding their imperfections. In part, that is a question of factual knowledge about problems and their solutions. In part, that is a question of appraising the limits to one's understanding of the system. That understanding is essential to being ready for surprises.

As mentioned earlier, considerable research has examined people's ability to assess the limits of their own understanding (Wallsten and Budescu, 1983). Typically, it has shown weak positive correlations between how confident individuals are in their own knowledge and how extensive that knowledge is. Although individuals are more knowledgeable when they are more confident, the relationship is quite imperfect. The most common overall tendency is toward overconfidence. Similar results have been observed in various settings, including some involving experts making judgments in their areas of expertise (Henrion and Fischhoff, 1986; Hynes and Vanmarcke, 1976) and some involving people's assessment of their understanding of technical systems (Fischhoff and MacGregor, 1986).

Although it could express itself as overconfidence in the reliability of a system, overconfidence in one's own understanding could also express itself in undue readiness to override a system and assume personal control. This has, for example, been the experience with attempts to automate various kinds of clinical diagnosis (Dawes, 1979). It is, therefore, important to know how accurately the operators and designers of a system are able to assess the extent of their own understanding of its operations. If these assessments are inaccurate, then it becomes important to know what cognitive processes are involved in assessing confidence (e.g., what cues do operators attend to? how do they weigh conflicting cues?). These processes provide the points of leverage for improving their self-understanding (e.g., by training, restructuring information flows, formalizing the evaluation process).

One methodological obstacle to creating more realistic expectations is the difficulty of evaluating current expectations in operational settings. Some novel procedures are needed to extract expectations in a more or less online manner and then to compare them with actual system performance. It may be possible to meter performance in some way, or to create a "black box" that could be used to compare what operators thought was happening with what was really happening (following successful operations, as well as following unsuccessful ones).

Once the accuracy of expectations has been assessed, it must be communicated in ways that will appropriately shape operator (and designer) behavior. Research has shown that just telling people about a judgmental difficulty has little effect, without some instruction in how to think differently and in how to match abstract principles of thought and analysis to concrete problems (Fischhoff, 1982; Kahneman et al., 1982; Murphy and Winkler, 1984; Nisbett et al., 1983). Further research is needed in this aspect of helping people to use their minds better. It might include exploration of alternative statistics for characterizing either the system or observers' understanding of it. Information about system reliability could come in the form of various summary statistics, but also in the form of structural information that might provide insight into the nature of problems as well as their magnitude. For example, it might be helpful to know about unresolved tensions in the design team, about the kinds of individuals (if any) who represented the perspectives of operators during the design process, about the number (or recency) of changes in design philosophy, about the state of the science underlying the design, and about the kind of external peer review to which the design was subjected. Whether such cues contain valid information is an analytical question. Whether that information can be used is an empirical behavioral question.

Expectations are the product of applying general beliefs to specific situations, as they are revealed by a system's external indicators. Normally, designers do everything possible to improve a system's transparency, that is, the chances that its status and operation will be interpreted appropriately. Where transparency is less than complete, however, operators need to understand a system's imperfections. The degree to which a system facilitates that understanding might be termed its metatransparency. In principle, transparency and metatransparency might be quite independent. In practice, they might even vary inversely. For example, summary presentations of current system status could facilitate getting a general feeling for the system, but obscure the raw observations that provide cues to the reliability of that summary. More generally, any refinement to a system can disrupt those finer points of its behavior that provide vital cues to judgments of its reliability. Thus, designers might consider when operators would be better off with a system that is harder to read but has better understood quirks. To avoid such tradeoffs, they might be helped by research into how to introduce improvements without disrupting operators' local knowledge. This question is analogous to the questions of how to update models (discussed above) and how to avoid deskilling (discussed below).

One potential source of information regarding the limitations of a system is analysis of specific problems that it has had. Superior methods for incident analysis would be useful in this regard. One problem facing those methods is having mixed and conflicting purposes. Assigning blame, determining causality, and estimating the probability of future mishaps are missions that call for somewhat different and incompatible procedures. A second problem is the effect of hindsight bias, which can distort observers' interpretations of past events and

even the memories of direct participants (Pew et al., 1982). A third obstacle to effective event analysis is ambiguity in the definition of events. For example, if incidents are defined too narrowly, then the lessons learned may ensure that a particular event sequence will not recur, but give the feeling that a whole class of events has been treated. Here, too, research is needed into the cognitive processes contributing to these problems and the procedures for overcoming them.

If events are defined precisely, then they may be amenable to theoretical analysis of the optimal breadth (or level) of analysis. As the category of event being studied broadens, a wider set of evidence becomes available, at the price of being able to reach less precise conclusions and recommendations.⁵ There are other behavioral aspects of dealing with imperfect systems that might benefit from analytical work. One is evaluating the sensitivity of decision making to different kinds of imperfection in information (Henrion, 1980; Krzysztofowicz, 1983; McCormick, 1981; von Winterfeldt and Edwards, 1982). Another is seeing how uncertainty about different aspects of the system accumulate to an overall estimate of its reliability (e.g., do they cancel or amplify one another). Another is providing some insight into the asymptotic level of reliability possible with systems of different levels of complexity (Perrow, 1984).

The ultimate expression of a lack of confidence is the decision to override a system over which the operator exercises supervisory control. It would be useful to have a fuller description of the override decision. What cues set it off? What steps are taken to confirm suspicions? How wide a set of system components (or operator actions) is called into question? What is the residual core of solid beliefs about the system? What cues are interpreted as demonstrating the return of control? How does one override decision affect subsequent behavior? In addition to descriptions of such decisions, one would want evaluations of their validity. Such evaluations might be available in existing system performance statistics. Or, operators' concerns might direct further research about the system. What operators do in the essentially novel situations created by a decision to override is the topic of the following section.

RESEARCH NEEDS: MANAGING NON-ROUTINE SITUATIONS

Any system concerned with irregularities that pose serious threats to life and property must prepare for contingencies. One standard method for doing so is by contingency planning: possible problems are anticipated; the best solution to each is identified; those solutions are then incorporated in the training of operators. If successful, such exercises will lead to the decision regarding the appropriate response being made well before any contingency arises. Such deliberate decisions should benefit from the reduced time pressure, reduced (emotional) stress, and greater ability to recruit diverse experts (or even to conduct research) which comes with planning. In this view, operators will be relieved of the need to make decisions in non-routine situations, by making those situations familiar in the form

of hypothetical experiences (even if those have yet to be experienced in reality). The decisions will be made by the contingency planners, leaving the operators to decide that some contingency has arisen and to decide which one it is. Then, the correct plan is accessed and executed.

Contingency planning requires a number of intellectual skills, each of which could benefit from study directed at ways to augment it. At the planning stage, these skills include the ability to imagine contingencies at all, the ability to elaborate their details sufficiently, the ability to generate alternative responses for evaluation, the ability to evaluate those responses critically in the hypothetical mode, and the ability to communicate the resultant decisions to operators. At the execution stage, these skills include the ability for operators to diagnose their crisis situations in ways that allow them to access the correct plan. Failures at either of these stages may result in ineffective decisions or in operators wondering about the appropriateness of the decisions that they are required to implement.

These problems are analogous to those facing effective emergency training in simulators. One worries, for example, that those who develop simulator exercises, teach the textbook responses, and evaluate operators' performance share some deep misconceptions about the system's operation--so that some critical contingencies are never considered. One also worries that spotting contingencies in the simulator might be quite different from spotting them in reality, where the system may have a different operating history or different social setting, or where operators are not as primed to expect problems (which typically come at enormously high rates in simulators). Understanding how people perform the component tasks in contingency planning might help decrease the number of non-routine decisions that have to be made (by making contingency planning more effective) and help assess the need for making non-routine decisions (by assessing the limits of contingency planning).

Such understanding might also help reduce the threats posed by undue reliance on contingency planning. One such threat is taking too seriously designers' idealizations of the system. Such models often provide a convenient basis for generating problems and exercises. They may even be used to run automated simulators. However, it is in the nature of models that they capture but a piece of reality, often without a clear (and communicated) understanding of just what that piece excludes. In some cases, a model is actually made to do double duty, being used by designers to discover limitations of the system (leading to design changes) and by trainers as though it represented a stable, viable operating system.

More generally, one needs to worry about how routine system operations affect operators' ability to deal with non-routine situations. Inadvertently inculcating undue faith in a basic design that typically functions well would be one kind of interference, as would acting as though contingency planning had routinized the treatment of novel situations. Institutional threats might include failing to train for handling non-routine situations or failing to

reward those who naturally have the skills for doing so (assuming that such skills could be discerned). The previous section suggested the possibility that the continuous introduction of design improvements or the polishing of synthetic data displays might disrupt operators' ability to "read" the system's state and to diagnose novel situations.

A general theoretical perspective for such research would be to consider the particular informational ecology in which judgment is acquired as a learned skill. Whenever that ecology changes, then there is some need to refine or alter judgmental skills, and some threat of negative transfer. A variant on this threat is deskilling, whereby useful intellectual skills are allowed to wither or are neutralized by design features or changes. For example, as automation increases, operators will increasingly be faced with near-perfect systems, which fail so seldom that there is little opportunity to learn their idiosyncracies. The problems of getting operators "back in the loop" so that they can cope with non-routine decisions may require some reduction in automation and perfection. The result of deautomation might be an increased rate of errors overall, but a reduced rate of catastrophic ones (a result that would be hard to prove given the low rate of occurrence for catastrophes). Research on these issues would seem hard and important.

Whenever there is some significant chance that contingency planning will not do, some capability is needed for making decisions in real time, starting from a raw analysis of the situation (perhaps after going part of the way with an inappropriate contingency plan). Training (and rewarding) the relevant intellectual skills (i.e., basic decision-making abilities) would seem extremely important. Much more needs to be known about how it can be done. For example, operators need to be able to generate good options regarding what might be happening and what might be done about it. Studies of creativity, in vogue some years ago, ostensibly examined this question. However, they used rather simple tasks and rather simple criteria for evaluating options (typically, the more the better). One potential aid to testing those options that are generated would be on-line, real-time system simulators. These could help operators diagnose the situation that they see by simulating the situations that would arise from various possible initiating conditions. They could also allow simulating the effects of various interventions. Getting such systems to work suggests some interesting computing and interface design problems.

A somewhat different kind of aid would be base-rate information describing typical performance of the system (or ones like it) under particular conditions. That information might describe, for example, what kinds of manipulations (in general) give one the best chance of being able to recover if they do not seem to be working, what manipulations provide the most diagnostic information about their failings, what are the best sources of information about current system status. Such statistical information might prove a useful complement to causal information about the system's intended operation. Its collection would represent an institutional commitment to learning from experience systematically.

It is often assumed that the choice of actions follows directly from diagnosing of the situation and anticipating of the effects of possible interventions. However, all decisions are contingent on objectives. Most organizations have complex objectives, some admitted and some implicit. Decision making can be paralyzed if the implications of those general values cannot be extracted for particular situations. It can be disastrous if the interpretations are inappropriate. Here, too, a mixture of analytical and behavioral work may help to improve that application and anticipate misapplications.

CONCLUSIONS

Research Management

The topics described here were selected for their implications for the design and operation of equipment such as would be found in the space station and its support systems. They are, however, described in terms of the general psychological processes that they involve. As a result, they could be pursued both as part of the development work for specific NASA systems and as basic research issues examined in laboratory settings intended to represent low-fidelity simulations of the actual NASA environments. Similarly, NASA could contribute to concurrent research prompted by other systems that place similar intellectual demands on designers and operators. Such connections would help to ensure the transfer of technology from NASA to the general community concerned with automation.

Insofar as this research deals with problems relevant to other technologically saturated environments, it should be able to learn from developments there. One relevant trend is the increasing scrutiny that is being given to the quality of expert judgment in technical systems. Some of that interest comes from within, out of concern for improving the engineering design process. Other interest comes from outside, out of the efforts of critics who wish to raise the standard of accountability for technological problems. In the face of that criticism, expert judgment proves to be a particularly vulnerable target. Although there is frequently great faith within a profession in the quality of its judgments, there is not that much of a research base on which to base a defense (Feyerabend, 1975; Morgan et al., 1981; Nelkin, 1984). Such research would have considerable basic, applied, and even political interest.

A second relevant trend is the introduction of computers into industrial settings. The creation of equipment has always carried an implicit demand that it be comprehensible to its operators. However, it was relatively easy for designers to allow a system to speak for itself as long as operators came into direct contact with it. Computerization changes the game by requiring explicit summary and display of information (Hollnagel et al., 1986). That, in turn, requires some theory of the system and of the operator, in order to know what to show and how to shape the interface. That "theory" might be created in an ad hoc fashion by the system's designers. Or, there

might be some attempt to involve designers with some expertise in the behavior of operators, or even representatives of the operators themselves (even in places where they do not have the high status of, say, pilots). A prejudice of this article, and other pieces written from a human factors perspective, is that concern over operability should be raised from the very inception of a project's development. Only in that way is it possible to shape the entire design with operability as a primary concern, rather than as a tack-on, designed to rescue a design that has been driven by other concerns. As a result, raising these issues is particularly suited for a long-term development project, such as that concerning this working group and volume.

Philosophy

A fundamental assumption of this chapter is that much of life can be construed as involving decisions (i.e., the deliberate choice among alternatives, often with uncertain information and conflicting goals). A corollary assumption is that the basic cognitive (or intellectual) skills involved in decision making have wide importance--if they can be understood and facilitated.

These are hard issues to study. However, even if they cannot be resolved in short order, system performance might be improved simply by drawing attention to them. A task analysis of where such skills arise can increase sensitivity to them, grant legitimacy to operators' complaints regarding problems that they are experiencing, and encourage a folklore of design principles that might serve as the basis for subsequent research.

The decision-making perspective described here is strongly cognitive, in part, because the decision theory from which it is drawn offers a widely applicable perspective and a well-defined set of concepts. As a result, there is a relatively high chance of results rooted in this perspective being generally applicable. Moreover, there may be some value to a general habit of characterizing decision-making situations as such. Within this context, there is still place to ask about issues such as the effects of stress, tension, conflict, fatigue, or space sickness on these higher-order cognitive processes (Wheeler and Janis, 1980).

This perspective sees people as active in shaping their environment and their decision problems. It could be contrasted with an operation research-type perspective in which people are reduced to system components and behavioral research is reduced to estimating some performance parameters. Focusing on what people do, rather than on the discrepancy between their performance and some ideal, increases the chances of identifying interventions that will help them to use their minds more effectively.

ACKNOWLEDGEMENTS

Support for preparation of this report came from National Science Foundation Grant SES-8213452 to Perceptonics, Inc. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the National Science Foundation. The Foundation's support is gratefully acknowledged. The thoughtful comments of Lita Furby, Ken Hartmann, Azad Madni, Ola Svenson, and members of the Symposium working group were also greatly appreciated. Correspondence may be addressed to the author at Department of Engineering and Public Policy, Carnegie-Mellon University, Pittsburgh, PA 15213.

NOTES

1. The chapters in this volume by Buchanan, Davis, Howell, Mitchell, and Newell provide other points of access to this literature.
2. The relationship between problem solving and decision making bears more discussion than is possible here, see National Research Council, 1986 for additional information.
3. In this particular case, there seems to be such generality, unless experience provides the sort of feedback needed to acquire probability assessment as a learned skill.
4. Fischhoff (in press) is an attempt to provide access to this literature, expressed in the context of the judgmental component of risk analyses for hazardous technologies.
5. Furby and Fischhoff (1986) discuss related issues in a very different context.

REFERENCES

- Bar Hillel, M.
 1980 The base-rate fallacy in probability judgments. Acta Psychologica 44:211-233.
 1984 Representativeness and fallacies of probability judgment. Acta Psychologica 55:91-107.

- Behn, R. D., and Vaupel, J. W.
1982 Quick Analysis for Busy Decision Makers. New York: Basic Books.
- Bentkover, J. D., Covello, V. T. and Mumpower, J., eds.
1985 Benefits Assessment: The State of the Art. Dordrecht, The Netherlands: D. Reidel.
- Beyth-Marom, R.
1982 Perception of correlation reexamined. Memory and Cognition 10:511-519.
- Brown, R. V.
1970 Why executives do not use decision analysis. Harvard Business Review 48:78-89.
- Combs, B., and Slovic, P.
1979 Causes of death. Journalism Quarterly 56:837-843.
- Davis, J. H.
1982 Group Performance. Reading, MA: Addison-Wesley.
- Dawes, R. M.
1979 The robust beauty of improper linear models. American Psychologist 34:571-582.
- Einhorn, H. J., and Hogarth, R. M.
1978 Confidence in judgment: Persistence in the illusion of validity. Psychological Review 85:395-416.
- Evans, J. St. B. T.
1982 The Psychology of Deductive Reasoning. London: Routledge and Kegan Paul.
- Feather, N. T.
1982 Expectancy, Incentive, and Action. Hillsdale, NJ.
- Feyerabend, P.
1975 Against Method. London: Verso.
- Fischhoff, B.
1975 Hindsight Foresight. Journal of Experimental Psychology: Human Perception and Performance 1:288-299.
1980 Clinical decision analysis. Operations Research 28:28-43.
1987 Judgmental aspects of risk analysis. In Handbook of Risk Analysis. Prepared for the Office of Management and Budget by the National Science Foundation. (In press.)

- Fischhoff, B., and Beyth-Marom, R.
1983 Hypothesis evaluation from a Bayesian perspective. Psychological Review 90:239-260.
- Fischhoff, B., Lichtenstein, S., Slovic, P., Derby S., and Keeney, R.
1981 Acceptable Risk. New York: Cambridge University Press.
- Fischhoff, B., and MacGregor, D.
1986 Calibrating databases. Journal of the American Society of Information Sciences 37:222-233.
- Fischhoff, B., Slovic, P., and Lichtenstein, S.
1978 Fault trees: Sensitivity of estimated failure probabilities to problem representation. Journal of Experimental Psychology: Human Perception and Performance 4:330-344.
- Fischhoff, B., Watson, S., and Hope, C.
1984 Defining risk. Policy Sciences 17:123-139.
- Furby, L., and Fischhoff, B.
1987 Rape self-defense strategies: A review of their effectiveness. Victimology. (In press.)
- Gardenier, J.
1976 Toward a science of marine safety. Proceedings of a Conference on Marine Safety.
- Grether, D. M., and Plott, C.
1979 Economic theory of choice and the preference reversal phenomenon. American Economic Review 69:623-638.
- Hasher, L., and Zacks, M.
1984 Automatic processing of frequency information. American Psychologist 39:1172-1188.
- Henrion, M.
1980 Sensitivity of Decisions to Miscalibration. Unpublished Ph.D. dissertation. Carnegie-Mellon University.
- Henrion, M., and Fischhoff, B.
1986 Uncertainty assessment in the estimation of physical constants. American Journal of Physics 54:791-798.
- Henrion, M., and Morgan, M. G.
1985 A computer aid for risk and other policy analysis. Risk Analysis 5:195-207.
- Hershey, J. C., Kunreuther, H. C., and Schoemaker, P. J. H.
1982 Sources of bias in assessment procedures for utility functions. Management Science 28:936-954.

- Hollnagel, E., Mancini, G., and Woods, D., eds.
1986 Intelligent Decision Support in Process Environments.
Heidelberg: Springer Verlag.
- Hynes, M., and Vanmarcke, E.
1976 Reliability of embankment performance prediction. In
Proceedings of the ASCE Engineering Mechanics Division
Specialty Conference. Waterloo, Ontario, Canada:
Waterloo University Press.
- Janis, I. L., and Mann, L.
1977 Decision Making. New York: Free Press.
- Kahneman, D., Slovic, P., and Tversky, A., eds.
1982 Judgment under Uncertainty: Heuristics and Biases.
New York: Cambridge University Press.
- Kahneman, D., and Tversky, A.
1979 Prospect theory. Econometrica 47:263-292.
- Krzysztofowicz, R.
1983 Why should a forecaster and a decision maker use Bayes
Theorem. Water Resources Journal 19:327-336.
- Ianir, Z.
1982 Strategic Surprises. Ramat Aviv: Tel Aviv University
Press.
- McCormick, N. J.
1981 Reliability and Risk Analysis. New York: Academic Press.
- Metcalf, J., III.
1986 Decision making and the Grenada Rescue Operation. Pp.
277-297 in J. G. March and R. Weissinger-Baylon, eds.,
Ambiguity and Command. Marshfield, MA: Pitman.
- Montgomery, H.
1983 Decision rules and the search for a dominance structure.
In P. Humphreys, O. Svenson and A. Vari, eds., Analyzing
and Aiding Decision Processes. Amsterdam: North Holland.
- Morgan, M. G., Henrion, M., and Morris, S. C.
1981 Expert Judgments for Policy Analysis. Brookhaven National
Laboratory, Brookhaven, NY.
- Murphy, A. H., and Winkler, K. L.
1984 Probability of precipitation forecasts. Journal of the
American Statistical Association 79:489-500.

- Myers, D. G., and Lamm, H.
1976 The group polarization phenomenon. Psychological Bulletin 83:602-627.
- National Research Council.
1986 Report of the research briefing panel on decision making and problem solving. Washington, D. C.: National Research Council.
- NASA Briefing
1986 Briefing Material, Johnson Space Center. May, Houston.
- Nelkin, D., ed.
1984 Controversy: Politics of Technical Decisions. Beverly Hills, CA: Sage.
- Nisbett, R. E., Krantz, D. H., Jepson, C., and Kunda, Z.
1983 The use of statistical heuristics in everyday inductive reasoning. Psychological Review 90:339-393.
- Nisbett, R. E., and Ross, L.
1980 Human Inference: Strategies and Shortcomings of Social Judgment. Englewood Cliffs, NJ: Prentice Hall.
- Payne, J.
1982 Contingent decision behavior. Psychological Bulletin 92:382-401.
- Perrow, C.
1984 Normal Accidents. New York: Basic Books.
- Peterson, C., and Beach, L.
1967 Man as an intuitive statistician. Psychological Bulletin 63:29-46.
- Pew, R., Miller, D. C., and Feeher, C. E.
1982 Evaluation of Proposed Control Room Improvements through Analysis of Critical Operator Decisions. EPRI NP-1982. Palo Alto, CA: Electric Power Research Institute.
- Pitz, G. F., Sachs, N. J., and Hoerboth, J.
1980 Procedures for eliciting choices in the analysis of individual decisions. Organizational Behavior and Human Performance 25:396-408.
- Raiffa, H.
1968 Decision Analysis. Reading, MA: Addison-Wesley.
- Rasmussen, J., and Rouse, W. B., eds.
1981 Human Detection and Diagnosis of Detection Failure. New York: Plenum.

- Samet, M. G.
1975 Quantitative interpretation of two qualitative scales used to rate military intelligence. Human Factors 17:192-202.
- Schoemaker, P. J.
1983 The expected utility model: Its variants, purposes, evidence, and limitations. Journal of Economic Literature 20:528-563.
- Shaklee, H., and Tucker, D.
1980 A rule analysis of judgments of covariation events. Memory and Cognition 8:459-467.
- Simon, H.
1957 Models of Man: Social and Rational. New York: Wiley.
- Slovic, P., and Lichtenstein, S.
1983 Preference reversals: A broader perspective. American Economic Review 73:596-605.
- Tversky, A., and Kahneman, D.
1973 Availability: A heuristic for judging frequency and probability. Cognitive Psychology 4:207-232.
- Tversky, A., and Kahneman, D.
1981 The framing of decisions and the psychology of choice. Science 211:453-458.
- U. S. Nuclear Regulatory Commission.
1983 PRA Procedures Guide. (NUREG/CR-2300) Washington, D. C.: U. S. Nuclear Regulatory Commission.
- von Winterfeldt, D., and Edwards, W.
1982 Costs and payoffs in perceptual research. Psychological Bulletin 93:609-622.
1986 Decision Analysis and Behavioral Research. New York: Cambridge University Press.
- Wallsten, T., and Budescu, D.
1983 Encoding subjective probabilities: A psychological and psychometric review. Management Science 29:151-173.
- Wheeler, D. D., and Janis, I. L.
1980 A Practical Guide for Making Decisions. New York: Basic Books.

Wise, B. P.

1986 An Experimental Comparison of Uncertain Inference Systems. Unpublished Ph.D. dissertation. Carnegie-Mellon University.

DISCUSSION: ISSUES IN DESIGN FOR UNCERTAINTY

William C. Howell

Reviewing the presentations of Drs. Davis and Fischhoff, one would be hard pressed to find critical omissions in the slate of issues set forth regarding human participation in the space station's judgment/decision/problem-solving requirements. The problem facing the R&D team, like that facing the future operators of the system itself, is deciding which of the plethora of options to address first--and to what depth--in the absence of complete knowledge. Agenda will have to be set, priorities established among research objectives (all of which seem worthy), and decisions made on when understanding has reached a sufficient (albeit far from ideal) level to move on to either development or the next agenda item.

The present discussion, therefore, will focus on some of these programmatic considerations. It would, of course, be presumptuous for anyone to prejudge the relative merit of research programs yet to be proposed for a moving target such as the evolving space station concept. Nonetheless, current knowledge is sufficient to begin the process so long as it is with the clear understanding that frequent stock-taking and consequent reorientation will undoubtedly be required as research findings accumulate, design decisions are made, and the entire system takes shape. Research never proceeds in as orderly a fashion as we anticipate in our plans and proposals because Mother Nature doesn't read them. One never knows when she will choose to reveal some important secret that will divert the whole process!

And finally, the discussion of priorities should in no way be construed as a call for serial research. The philosophy endorsed here is consistent with a theme that runs through the entire symposium: parallel research efforts must be carried out at various levels of specificity on a representative sample of the total problem space if the program is to evolve--and continue to develop--in the most efficacious manner. The pressure to focus too narrowly on the most well-defined or immediate problems is all too prevalent in undertakings of this magnitude having the level of public visibility that the space station enjoys. Many of the problems sure to arise "downstream" are in areas where the present knowledge base is at best primitive. Attention must be given now to expanding those knowledge bases if we are to avoid costly delays in development and/or costly design mistakes as the total system takes shape.

Model Building

Both presentations emphasize the importance of developing a conceptual model or set of models of the space station. Together, Davis and Fischhoff sketch out the essential features of such modeling and the kinds of research questions that must be addressed in order to make it useful. I shall not repeat their observations, except to note one point of contrast and to explain why I believe model building deserves a top priority.

First the contrast. Davis makes a distinction between aspects of the total system about which there is and is not sufficient information to construct models. Where it is deemed feasible, chiefly in the physical domain, the trick is to make the models--and the systems they represent--"resourceful" and comprehensible. Where it is not, the issue becomes one of finding alternatives to modeling. Fischhoff, on the other hand, seems to have in mind a more comprehensive kind of modeling effort: one that encompasses a variety of domains and levels of understanding. Here the emphasis is on integrating what we know even incompletely, and providing a framework upon which to build new understanding.

Whichever concept one prefers, and I lean toward the latter, the research issues are largely the same. Both call for exploring new ways to capture and express properties of the system that will promote understanding across disciplines; both recognize that to do so requires a better grasp of certain cognitive functions than we now have. There are, in my view, at least four main reasons to emphasize a broad modeling effort (Meister, 1985).

First, the process of model building is the most expeditious way to organize our knowledge and ignorance, not only at the outset, but as the knowledge base grows and the system evolves. Assumptions, facts, parameter estimates, areas of uncertainty etc. can be clearly articulated; gaps that need to be filled, or estimates that need to be refined, can be identified. More than anything, a conceptual model can ensure that even the most pragmatic research has a better chance of contributing to the total effort. Taken literally, for example, the issues raised by Davis and Fischhoff cover virtually the entire domain of cognitive and social psychology. Were nature to take its course in these various research areas (or even were NASA support to accelerate the overall progress), the odds of learning precisely what needs to be known at critical junctures in the space station's development are quite low. I shall have more to say on this point later. For present purposes, the argument is simply that model building is a useful technique for keeping the research efforts at all levels of generality properly focused. One can study confidence in judgement, or interpersonal tension, or hypothesis generation, or human problem solving tendencies, or what experts know and do, or any of the other general issues identified by the presenters in ways that are more or less likely to generalize to the space station situation. I see no inherent reason why an experiment designed to advance fundamental knowledge in one of these areas cannot be conducted in a space-station context as easily as in terms of weather forecasting, battle planning,

livestock judging, or business management. A model is useful for specifying that context.

A second reason that model building merits the highest priority lies in its contribution to the ultimate development of tasks and procedures. The ways in which this contribution would manifest itself are well described in the two presentations. In essence it boils down to making reasoned design decisions from a system-wide perspective rather than from some parochial or purely traditional point of view--be that an engineering, computer science, cognitive, biomedical, or even a humanistic perspective. It forces early attention to such critical matters as developing a common language and frame of reference within which the various specialists can function interactively. If there is one unique requirement for the successful achievement of this project's goal, it is that barriers to the exchange of information and intelligence among units--human-human, human-machine, machine-machine--be minimized. Systems of the past have generally had to attack such barriers after the fact because of the initial dominance of one or another technical specialty. And they have done so with only limited success. Here the opportunity exists to "design in" features that can minimize barriers. Model development encourages this kind of thinking from the very outset--provided, of course, it is not entrusted to only one technical specialty!

A third argument for the priority of model building is its obvious importance for training, and possibly even personnel selection. True, a model is not a simulation. Nevertheless, simulation at some level of fidelity must ultimately be constructed just as it has been for training on all the earlier projects in the space program. To the extent that the model organizes what is known and unknown at a particular stage, it permits development of simulations that have a greater likelihood of providing training that will transfer positively to the operational tasks. The kinds of uncertainties and unanticipated contingencies the human is apt to encounter in the space station are more likely to arise in a simulator based on a comprehensive modeling effort than they would be in a simulator designed to maximize purely technical fidelity. In the absence of a good conceptual model, the criterion of technical fidelity is almost certain to dominate. To use an extreme example, suppose the modeling effort identified a social phenomenon whose course of development extends over a period of months and whose appearance dramatically alters the way certain kinds of decisions are handled. Naturally, this would argue for incorporating a several month duration requirement into the simulation even if the technical skills could be mastered in weeks. Without this social-process knowledge, the emphasis would almost certainly be on the face validity of the hardware and software components. In other words, comprehensive model development would increase the likelihood that any simulation would capture salient aspects of the operational tasks--even some that cannot be completely anticipated and "programmed in." Similarly, it would provide a better sampling of the overall task domain and hence a more content-valid basis for setting personnel selection requirements.

In citing the virtues of model development for simulation and training, we should never lose sight of Fischhoff's warning against the possibility of overemphasizing the known to the exclusion of the unknown. Training that develops in operators a dependence on routines for handling anticipatable contingencies can be counterproductive when truly novel ones arise. However, thoughtful construction of a model can help obviate this problem by ensuring that the unknown is properly recognized. The real danger lies not in the attempt to build the most complete conceptual models we can, but in the temptation to build simulators that operate only within the domains where our knowledge is most complete.

Finally, model development encourages--indeed forces--the kind of interaction among specialists in the design phase that will have to occur among operational specialists if the program is to be a success. To mount a truly comprehensive modeling effort will demand creation of a shared language and knowledge base; the exercise will serve, in essence, as a case study in multidisciplinary coordination as well as the source of a design product.

In a sense, all the other proposed research directions are subsumed under the objective of model development (or at least are directly related to it). As Davis points out, constructing an appropriately "robust" and "transparent" model requires judicious selection of which properties to include and ignore, and at what level of abstraction. How well that can be done is heavily dependent on our understanding of human cognitive processes in relation to the physical properties of the system. And it is largely to this end that the research suggested by Davis, Fischhoff, and indeed this entire conference is directed. Nevertheless, one can distinguish more narrowly defined issues, and some of these appear more promising or tractable at this point than others. Several that strike me as particularly deserving of a high priority are establishment of institutional values, manual override and standby capabilities, and transfer of training issues.

Establishing Institutional Values

Fischhoff explains that a critical issue facing decision makers in the operational system will be that of representing the organization's values in dealing with non-routine situations. One cannot anticipate all the circumstances that might arise that would require human judgment, but it is possible to define the value parameters along which those judgements would have to be made and the extent to which institutional, crew, or individual value systems would take precedence.

Most decisions incorporate value and expectation considerations in one form or another (Huber, 1980; Keeney and Raiffa, 1976). There are a lot of ways to help objectify or improve the expectation element, but values are inherently subjective. This is why there are political systems, judicial systems, wars, and advertising agencies. Unless we can articulate the value system under which the decision maker is to operate--or at least the general process by which s/he is to assign values--s/he faces an impossible task. It is somewhat akin to that

facing the medical community in its allocation of scarce and costly life-saving resources (such as organ transplants) to a much larger and multifaceted population of worthy recipients. Whose interests take precedence, and how are the value considerations to be weighed?

This issue is not an easy one to address, in part because it gets to the heart of the most sensitive, controversial, and politically charged aspects of any important decision domain. We do not like to make explicit the level of acceptable risk in air safety, nuclear power, or military confrontation (e.g. how many lives we are willing to sacrifice for some larger good). However, there is some implicit value system operating in any such decision, and research over the past decade has produced methodologies for helping to pin it down (Howard, 1975; Huber, 1980; Keeney and Raiffa, 1976; Slovic et al., 1980). Extension of these techniques, and perhaps development of others, to provide a common value framework for crews and individuals to carry with them into space is essential if decision-making is to be of acceptable quality. Indeed, without such a framework the concept of decision quality has no meaning. The options are to face the issue squarely and develop a value framework in advance, or to leave it intentionally vague and ad hoc, thereby offsetting whatever progress is made toward improving decision quality through enhancement of expectation judgments.

Understanding Override and Stand-by Capabilities

Clearly an important set of research issues centers around the idea that human judgment represents the last line of defense against the unanticipated. The ultimate decision that some automated subsystem is malfunctioning, or that some low probability or unclassifiable situation has arisen, and the skill to move quickly from a relatively passive to an active mode in response to it are critical elements of the human's role.

Both presentations address override and standby skill issues albeit in slightly different ways. For Davis, they fall within the category of "making the best of the situation," or what to do when we have no model. He speculates on alternative strategies, and suggests that we need to explore them, but is obviously more concerned with "making the best situation"--increasing the robustness and transparency of the system and its models. For Fischhoff, these issues epitomize a central dilemma in the whole development process--the tradeoff between using everything we know for aiding and contingency planning purposes, and preparing people to deal with the truly unknown. He argues that designing the system to maximize decision accuracy may not really be optimal when one considers the potential costs in human judgment facility. (Here, incidentally, is another instance where the problem of establishing a unified value system becomes critical.)

What strikes me as particularly urgent about research on these issues is that we know just enough to worry, but not enough to say how they should be handled. For example we know about overconfidence bias and can easily imagine its implications for crisis decision-making, but we are far from understanding all the task and individual-difference

parameters that govern its seriousness (Hammond et al., 1980; Howell and Kerkar, 1982). And we know even less about constructs such as creativity in either the individual or group context. Were we able to identify and measure such individual traits, we might include these measures in a personnel selection battery. And understanding group processes might suggest ways to offset deviant individual tendencies. Unfortunately, our present knowledge of group decision making does not allow us to predict with much certainty how group judgments will compare with individual ones (Huber, 1980; Retiz, 1977; Howell and Dipboye, 1986).

Similarly, it is fairly well established, as Fischhoff notes, that stand-by skills suffer from disuse as the human spends more and more time "outside the loop" in a monitoring capacity. This is particularly true for cognitively complex and dynamic systems. But how does one "stay on top of things" when active involvement becomes increasingly rare as more and more reliance is placed on automating decision functions? Is something as elaborate (and costly) as a totally redundant manual back-up ever justified simply for the purpose of maintaining stand-by capabilities? And even if that were done, would the human be able to maintain a serious involvement knowing the status of his or her role? One need only take a look at NORAD operators doing their "canned" training exercises to appreciate the significance of this point! Would some other form of involvement do as well? For what decision tasks should some form of involvement be maintained? To answer questions such as these, more will need to be learned about stand-by capabilities in critical tasks of the sort that are likely to be automated or aided in the space station. Fischhoff's presentation does an excellent job of identifying the key questions.

Issues concerning the override function should be addressed early in the development process at a fairly basic level since more general knowledge is needed before it will be possible to articulate the most critical applied research questions. Stand-by skill maintenance, on the other hand, seems more appropriately addressed at an applied research level after it becomes clear what sorts of functions the human would be asked to back up.

Training for the Known and the Unknown

Issues of training and transfer are closely related to those of standby skill; in fact, the latter are really a subset of the former. The purpose of training is to establish habitual ways of thinking and acting in certain situations that are likely to improve individual or team performance whenever those situations arise. So long as one has at least some idea of what kinds of situations might develop, there is reason to hope that the right habits might be cultivated. But if one guesses wrong, or the situation domain changes, or the habits that work well for the known situations turn out to be counterproductive for the unknown ones, obvious transfer problems arise. Since the unanticipated is by definition inaccessible for simulation or contingency planning, those charged with training development face the dilemma alluded to

earlier. Too heavy an emphasis on the known or suspected task elements could develop habits that prove disastrous when something totally novel comes along. On the other hand, training that emphasizes the flexibility of response necessary to deal with novel situations could undermine the potential advantages of habitual behavior.

Advances have been made toward addressing this dilemma in recent research on fault diagnosis and problem solving (particularly in connection with complex process control systems, e.g. Moray, 1981; Rasmussen and Rouse, 1981). Still, as Fischhoff notes, there are a lot of fundamental questions that remain to be investigated before we can even begin to conceptualize how training ought to be structured in a systems as advanced as the space station. Once again, we have here a set of pressing issues on which some headway has already been made and research directions have been identified. For these reasons, I believe it merits a high priority in the overall research scheme.

To this point, my comments have focused exclusively on priority setting within the domain of research issues raised by the two presenters. To summarize, I believe the modeling effort should be an initial and continuing emphasis--a framework within which many parallel streams of research activity can proceed coherently and purposefully. Of those more narrowly defined issues, I consider the matter of establishing institutional values or value assessment techniques as primary, followed closely by the need to clarify the override function, to find ways to maintain intellectual standby skills (or define an optimal level of automation), and to train operators to deal with changing and unanticipatable circumstances.

There are two other programmatic issues that I would like to comment on briefly that were not an explicit part of either paper: individual differences, and the age-old basic vs. applied research controversy.

On Individual Differences

Both presentations suggest quite correctly that our designs must be geared to typical behavior--of people in general, or potential operators, or "experts". The assumption is that there are commonalities in the way people approach particular decision problems, and our research should be directed toward understanding them. I agree. But I contend there is another perspective that has been all but ignored by decision theorists that might also contribute to the effectiveness of future decision systems. On virtually any standard laboratory problem, subjects will differ dramatically in both the quality of their performance and the way they approach it. True, the majority--often the overwhelming majority--will display a particular bias, heuristic, or preference on cue. But even in the most robust demonstrations of conservatism, or overconfidence, or representativeness, or non-transitivity there will be some subjects who don't fall into the conceptual trap. What we don't know, in any broader sense, is whether these aberrations represent stable trait differences, and if so, what their structure might be and how they might be measured. There has been some work on risk aversion

(Atkinson, 1983; Lopes, in press), information-processing tendencies (Schroder et al., 1967), and decision-making "styles" (Howell and Dipboye, 1986), but very little compared to the vast literatures on typical behavior.

I suspect, though I can't really prove it, that individuals differ consistently in their inclination to attend to, process, and integrate new information into their current judgments. Were this the case, it might be useful to have some means of indexing such tendencies. Speaking more generally, I believe research aimed at exploring the consistent differences in the way people approach decision problem is just as valid as--though considerably more cumbersome than--that concerned with similarities. It should be encouraged.

On Basic and Applied Research Strategies

At various places in the foregoing discussion I have suggested that certain issues might be attacked at a more basic or more applied level given the state of our current knowledge and the demands of the design problem in that area. I should like to conclude my discussion with some elaboration on this general strategic issue.

If there is one limitation on our understanding of judgment/decision processes, in my opinion, it is that of context specificity. Work on judgmental heuristics, diagnosis and opinion revision, choice anomalies, group decision making, individual differences in judgment or decision, etc. each has developed using its own collection of preferred research tasks, strategies, and literatures (Hammond et al., 1980; Schroder et al., 1967). Consequently, it is not always possible to judge how far a particular principle will generalize or whether some human tendency is likely to pose a serious threat to performance in a particular system.

Nevertheless, as the two presentations have clearly demonstrated, these basic literatures provide a rich source of hypotheses and leads for consideration in an evolving program such as the space station. The judgmental heuristics and resulting biases cited by Fischhoff, for example, are indeed robust phenomena, principles to be reckoned with in shaping the space station environment. However, despite their ubiquity, such modes of cognition are more prominent in some contexts and under some conditions than others--a point emphasized by Hammond in his "cognitive continuum theory" (Schum, 1985); and the seriousness of the consequent "biases" depends to some extent on one's definition of optimality (Hammond, 1981; Hogarth, 1981; Schroder et al., 1967; Phillips, 1984, Von Winterfeldt and Edwards, 1986).

Consider the overconfidence bias. One implication of this well established cognitive phenomenon is that decision makers would be likely to act in haste and believe unduly in the correctness of their action, a clearly dysfunctional tendency. Or is it? A common complaint in the literature on organizational management is that managers are all too often reluctant to act when they should (Peters and Waterman, 1982). Perhaps overconfidence may serve to offset an equally dysfunctional bias toward inaction in this setting. Similarly,

decisions must often be made under considerable uncertainty, and this will clearly be no less true of space station than of business or military decisions. However, once a decision is made, albeit on the basis of what objectively is only a 51% chance of success, is there not a certain practical utility in actually believing the odds are better than that? If, as often happens, the decision is not easily reversed, what is to be gained by second-guessing or "waffling", and is there not a potential for benefit through the inspiration of confidence in others? In some cases that alone can increase the "true" odds! The point is, overconfidence, like other human cognitive tendencies, may have functional as well as dysfunctional implications when viewed in a particular context (Hammond, 1981); and even then, its magnitude may be partly a function of that context. Thus the more clearly we can envision the context, the more likely we will be to generate the right research questions, and what that research adds to our basic understanding of overconfidence or other such phenomena will be no less valid than that done in other contexts. All judgment and decision research is done in some context; generalization accrues via convergence of evidence over a variety of contexts.

My basic point is this. The space station offers a very legitimate--indeed, an unusually rich--real-world context within which to explore a variety of "basic" and "applied" research questions concurrently. Properly coordinated, the combined effort holds considerable promise for advancing our understanding of fundamental judgment/decision processes in part because of the shared context. Three considerations would, I believe, promote such coordination.

First, as noted earlier, some effort should be made to encourage basic researchers to consider salient features of the space station situation in the design of their laboratory tasks and experiments. While it could be argued that putting any constraint at all on such work violates the spirit of "basic research," I believe some concessions can be made in the interest of increasing the external validity of findings without compromising the search for basic knowledge. Secondly, research of a strictly applied nature, addressing specific judgment/decision issues that must be answered in the course of modeling, simulation, and ultimately design efforts, should proceed in parallel with the more basic endeavors. In some cases, the question might involve choice of a parameter value; in others, identification of how subjects approach a simulated space-station task. Necessarily, such research would be less programmatic, more responsive to immediate needs, and more narrowly focused than the fundamental work.

Finally, and most importantly, NASA must do everything possible to ensure that the basic and applied efforts are mutually interactive. As hypotheses and generalizations are identified at the basic level they should be placed on the agenda of the applied program for test or refinement; as features are built into the evolving system concept, they should become salient considerations for the basic research effort; as questions of a fundamental nature arise in the course of the applied work, they should be incorporated into the basic research agenda.

This all sounds quite obvious and "old hat." Certainly it is the way DoD research programs, for example, are supposed to work (Meister, 1985). I submit, however, that no matter how trite the notion may seem, having closely coupled research efforts at basic and applied levels must be more than just an aspiration if the judgment/decision challenges of the space station project are to be met successfully. It must be planned and built into the very fabric of the program. The fact that the space station must develop by its own research bootstraps, as it were, permits little slippage and wasted effort. Yet the state of our knowledge does not permit neglect of either basic or applied research domains.

There are, of course, a number of ways this coordination of basic and applied work might be achieved ranging from centralized administrative control to large-scale projects that are targeted to particular sets of issues and encompass both basic and applied endeavors under one roof. I am not prepared to recommend a strategy. Rather, I suggest only that the issue is an important one, and one that deserves special attention at the very outset. How it is managed could spell the difference between enlightened and unenlightened evolution of the whole system regardless of how much resource is allocated to judgment/decision research.

REFERENCES

- Atkinson, J. W.
1983 Personality, Motivation, and Action. New York: Praeger
- Hammond, K. R.
1981 Principles of Organization in Intuitive and Analytical Cognition. Report No. 231, February. University of Colorado, Center for Research on Judgment and Policy.
- Hammond, K. R., McClelland, G.H., and Mumpower, J.
1980 Human Judgment and Decision Making: Theories, Methods, and Procedures. New York: Praeger
- Hogarth, R.M.
1981 Beyond discrete biases: functional and dysfunctional aspects of judgmental heuristics. Psychological Bulletin 90:197-217.
- Howard, R. A.
1975 Social decision analysis. Proceedings of the IEEE 63:359-371
- Howell, W. C., and Dipboye, R. L.
1986 Essentials of Industrial and Organizational Psychology. Chicago: Dorsey

- Howell, W. C. and Kerkar, S. P.
1982 A test of task influences in uncertainty measurement. Organizational Behavior and Human Performance 30:365-390.
- Huber, G. P.
1980 Managerial Decision Making. Glenview, Ill.: Scott Foresman
- Keeney, R., and Raiffa, H.
1976 Decision With Multiple Objectives: Preferences and Value Tradeoffs. New York: Wiley
- Lopes, L. L.
1987 Between hope and fear: the psychology of risk. Advances in Experimental Social Psychology. (In press.)
- Meister, D.
1985 Behavioral Analysis and Measurement Methods. New York: Wiley
- Moray, N.
1981 The role of attention in the detection of errors and diagnosis of failures in man-machine systems. J. Rasmussen and W. Rouse, eds. Human Detection and Diagnosis of System Failures. New York: Plenum.
- Peters, T. J. and Waterman, R. H.
1982 In Search Of Excellence. New York: Warner Books.
- Phillips, L.
1984 A theoretical perspective on heuristics and biases in probabilistic thinking. Humphreys, Svenson and Vari, eds., Analyzing and Aiding Decision Problems. North Holland Publishers.
- Rasmussen, J., and Rouse, W., eds.
1981 Human Detection and Diagnosis of System Failures. New York: Plenum
- Reitz, H. J.
1977 Behavior in Organizations. Homewood, Ill.: Irwin
- Schroder, H. M., Driver, M.J., and Streufert, S.
1967 Human Information Processing. New York: Holt, Rinehart, and Winston
- Schum, D. A.
1985 Evidence and Inference For The Intelligence Analyst. Draft prepared for the Office of Research and Development, Central Intelligence Agency. Copyright: D. A. Schum, 7416 Timberock Rd., Falls Church, Va. 22043.

- Slovic, P., Fischhoff, B., and Lichtenstein, S.
1980 Facts and fears: understanding perceived risk. R. C. Schwing and W. A. Albers, Jr., eds., Societal Risk Assessment: How Safe Is Safe Enough? New York: Plenum.
- Von Winterfeldt, D., and Edwards, W.
1986 Decision Analysis and Behavioral Research. New York: Cambridge University Press

SYNOPSIS OF GENERAL AUDIENCE DISCUSSION

Most of the points raised during Session IV and the general discussion centered around two somewhat related issues:

1. the gap between behavioral (heuristic) and traditional (rule based) approaches to decision making, and
2. how to deal with shortcomings in one or the other that detract from system performance.

The Gap Issue

The observation was made that there seem to be two ways of thinking about decision problems, each with its own philosophy and research agenda, that are proceeding more or less independently. To some extent, it was pointed out, the two papers in the session highlight the differences between the two approaches. The question was whether, and if so how, they should be integrated or linked more closely.

Two conflicting views were offered. One was that since the differences are deeply rooted in their respective traditions and cultures, the barriers will not be broken down easily, and the anticipated payoff for NASA would probably not justify the time and cost necessary to bring about an integration. A number of other issues should take precedence over this one. The opposing view was that the two approaches should be better integrated, probably can be if NASA puts the issue on its research agenda, and in fact is being attempted in a small way through research currently in progress in Fischhoff's lab.

Among the suggestions for an integrative approach were the whole domain of fuzzy logic and the bounded rationality concept (e.g. defining general goals and then "fiddling with the model at the margin as in 'satisficing'"). It was pointed out, however, that in the context of expert systems such approaches reduce to writing a lot of conditional rules over a large number of state variables. Thus one cannot summarize easily what the system will do over the full range of decision problems.

Applications, Or Dealing With Shortcomings

Several options were suggested for minimizing the effect of suboptimalities in human judgment. Training, while not universally effective in overcoming biases, has produced some notable successes (e.g. weather forecasters). The key may well lie in the proper design of training programs (something that merits a continuing research effort). Increasing the trainee's sophistication in statistical concepts, however, is clearly of little help.

Aiding in its various forms and with its inventory of existing models has its place but also has limitations. Multi-attribute utility theory, decision analysis, etc. are useful for solving well defined problems, but "bring no knowledge to the party." Often their logic is not transparent to the user and critical factors may be omitted. Thus their output may not be satisfactory in either an absolute sense or as perceived by the user. When it conflicts with human intuition there is a problem, particularly if the human doesn't understand the logic. User acceptance of even improved decisions becomes problematic.

One approach to dealing with these deficiencies in the aiding models was advocated by Davis: find out what is missing and build it in. Intuition and creative thinking are not magic, but rather, "undiscovered rationality." Research should try to expose that rationality (or reasoning) and apply it in creating more robust models, as well as more transparent ones. To the extent that the research succeeds, it should be incorporated into training as well as aiding applications, and the result could be better decisions and greater acceptance of those decisions by users (who would now be more likely to appreciate the logic).

SESSION V

TELEPRESENCE AND SUPERVISORY CONTROL

TELEOPERATION, TELEPRESENCE, AND TELEROBOTICS:
RESEARCH NEEDS FOR SPACE

Thomas B. Sheridan

INTRODUCTION

The Need and the Dilemma

One of the dramatic challenges posed by space is versatile inspection and manipulation remotely operated by man. Some people within and outside NASA would like to automate everything but cannot--because so many tasks are unpredictable and therefore not doable by special-purpose or preprogrammable machines, or are one-of-a-kind such that dedicated automatic devices to do them are too costly in weight and dollars. So human perception, planning and control are required. But to place man physically there is constrained by hazard and high cost of life support. Remote inspection and manipulation by man, on the other hand, poses serious problems of her getting sufficient sensory information and controlling with sufficient dexterity.

Artificial sensing, intelligence and control can help. Unfortunately we have hardly begun to understand how to integrate human and artificial brands of sensing, cognition and actuation. One thing is clear, however: to cast the problem in terms of humans versus robots is simplistic, unproductive and self-defeating. We should be concerned with how they can cooperate.

Definitions

Teleoperation is extension of a person's sensing and manipulating capability to a location remote from him. A teleoperator includes at the minimum artificial sensors, arms and hands, a vehicle for carrying these, and communication channels to and from the human operator.

Telepresence is the ideal of sensing sufficient information about the teleoperator and task, and communicating this to the human operator in a sufficiently natural way that she feels herself to be physically present at the remote site. A more restrictive definition requires, in addition, that the teleoperator's dexterity match that of the bare-handed human operator.

Robotics is the science and art of performing, by means of an automatic apparatus or device, functions ordinarily ascribed to human beings, or operating with what appears to be almost human intelligence (adapted from Webster's 3rd Intl. Dictionary).

Telerobotics is a form of teleoperation in which a human operator acts as a supervisor, communicating to a computer information about task goals, constraints, plans, contingencies, assumptions, suggestions and orders, getting back information about accomplishments, difficulties, concerns, and, as requested, raw sensory data--while the subordinate teleoperator executes the task based on information received from the human operator plus its own artificial sensing and intelligence. Accompanying the human supervisor is a computer which can communicate, integrate, assess, predict, and advise in human-friendly terms; at the site of the telerobot is a computer which can communicate with the human-interactive computer and effect control using the artificial sensors and effectors in the most efficient way. One human-computer command station can supervise many telerobots.

Supervisory control in the present context is mostly synonymous with telerobotics, referring to the analogy of a human supervisor directing and monitoring the activities of a human subordinate. Supervisory control does not necessitate that the subordinate person or machine be remote.

Early History

Prior to 1945 there were crude teleoperators for earth moving, construction and related tasks. About that time the first modern master-slave teleoperators were developed by Goertz at Argonne National Labs. These were mechanical pantograph mechanisms by which radioactive materials in a "hot cell" could be manipulated by an operator outside the cell. Electrical and hydraulic servomechanisms soon replaced the direct mechanical tape and cable linkages (Goertz, 1954), and closed circuit television was introduced, so that now the operator could be an arbitrary distance away. Soon telemanipulators were being attached to submarines by the Navy and used commercially by offshore oil extraction and cable-laying firms to replace human divers, especially as operations got deeper. By the mid 50s technological developments in "telepresence" (they didn't call it that at the time) were being demonstrated (Mosher, 1964; Johnsen and Corliss, 1967; Heer, 1973). Among these were: force reflection simultaneously in all six degrees of freedom; hands with multi-jointed fingers; coordinated two-arm teleoperators; and head-mounted displays which drove the remote camera position and thereby produced remarkable visual telepresence.

By 1965 experiments in academic research laboratories had already revealed the problems of telemanipulation and vehicle control through time delay (Ferrell, 1965), and the early lunar teleoperator Surveyor demonstrated the problems vividly in an actual space mission. Touch sensing and display research was already underway (Strickler, 1966) though there was little interest in teletouch at that time. Soon thereafter supervisory control was shown to offer a way around the time

delay problem, and also to have advantages even without time delay in the communication channel, where, in order to avoid collision or dropping grasped objects, quicker teleoperator reaction time was needed than the distant human operator could provide (Ferrell and Sheridan, 1967).

Though the NASA nuclear rocket project mounted a major effort in teleoperator development in the 1960s, after that program was cancelled and throughout the 1970s there was little support for space teleoperation or telerobotics. By 1970, however, industrial robotics was coming into full swing, for Unimation, GE and a handful of other American, Japanese and Scandinavian manufacturers had begun using relatively simple assembly-line robots, mostly for spot welding and paint spraying. By 1980 industrial robots had become graced by wrist force sensing and primitive computer vision, and push-button "teach pendant" control boxes were being used for relatively simple programming from the shop floor.

Overview of Current Status

To outward appearances six-degree-of-freedom, force-reflecting, serial-link electrical or hydraulic master-slave manipulators have changed little in forty years. There are a few new and promising mechanical configurations of arms and multi-fingered hands in laboratories, but as yet they are unproven in practical application. Video, driven by a demanding marketplace, is now of high quality and miniaturized, and digitization and simple recognition processing of video images is fast and inexpensive. We have a variety of touch (surface contact and pressure array) sensors available in the laboratory, but as yet little understanding of how to use these sensors. In teleoperation depth perception remains a serious problem, but there is promising research on several fronts. We still have not achieved fine, dexterous telemanipulation with high fidelity feedback as implied by the term "telepresence".

As yet there is no satisfactory control theory of manipulation as an integrated sensory-motor control activity, but new theories have been developed for manipulation task-analysis from an AI perspective, for kinematic-dynamic control of complex linkages, and for force-displacement hand-environment impedance. We still think of controlling manipulator arms and the vehicles which carry them as separate activities; we haven't learned to combine the two (though infants do it with ease). We have demonstrated simple human-supervised, computer-aided teleoperation in a number of ways, but our understanding of human-computer cooperation is very primitive, hardly commensurate with the label "telerobot" we employ with such abandon.

SPECIFIC AREAS IN WHICH NEW RESEARCH IS NEEDED

Research needs are discussed in four categories: (1) telesensing, (2) teleactuating, (3) computer-aiding in supervisory control, and (4) meta-analysis of human/computer/teleoperator/task interaction. Some recent and current research is cited.

Telesensing

My colleague, Dr. Stark, who is an MD and more sense-able than I, will deal more extensively with this category, particularly with vision, the most important human sense, and with the needs and possibilities in virtual displays and controls, depth perception, and other significant needs in teleoperator research.

I would like to comment about resolved force, touch, kinesthesia, proprioception, and proximity—five critical teleoperator sensing needs which must be recognized as being different from one another. These five, together with vision, are essential to achieve the ideal of "telepresence". For each it is important to understand how the human normally functions, and then to understand how the appropriate signals can be measured by artificial transducers and then displayed to the human operator and/or used by artificial intelligence in a way helpful to the human operator.

Resolved force sensing is what the human body's joint, muscle and tendon receptors do to determine the net force and torque acting on the hand, i.e., the vector resultant of all the component forces and torques operating on the environment. In force reflecting master-slave systems this is measured either by: (1) strain gage bridges in the wrist (so-called wrist-force sensors); (2) position sensors in both master and slave, which, when compared, indicate the relative deflection in six DOF (which in the static case corresponds to force); (3) electrical motor current or hydraulic actuator pressure differentials. Display of feedback to the operator can be straightforward in principle; in force-reflecting master-slave systems the measured force signals drive motors on the master arm which push back on the hand of the operator with the same forces and torques with which the slave pushes on the environment. This might work perfectly in an ideal world where such slave-back-to-master force serving is perfect, and the master and slave arms impose no mass, compliance, viscosity or static friction characteristics of their own. Unhappily, not only does reality not conform to this dream it can also be said that we hardly understand what are the deleterious effects of these mechanical properties in masking the sensory information that is sought by the operator in performing telemanipulation, or how to minimize these effects. At least, thanks to computer coordinate transformation, it has been shown that master and slave need not have the same kinematics (Corker and Bejczy, 1985). Force reflection can also be applied to a rate-control joystick (Lynch, 1972) but it is less clear what the advantages are.

Touch is the term used sloppily to refer to various forms of force sensing, but more precisely to refer to differential pressure sense of the skin, i.e., the ability of the skin to detect force patterns, with respect to displacement both tangential and normal to the skin surface, and to time. The skin is a poor sensor of absolute magnitude of force normal to the surface and it adapts quickly. There are now a few instruments for artificial teletouch; most of these have much coarser spatial resolution than the skin, though a few of the newer ones utilizing optics have the potential for high resolution (Harmon, 1982; Schneider and Sheridan, 1984). A major research problem for teletouch is how artificially sensed pressure patterns should be displayed to the human operator. One would like to display such information to the skin on the same hand that is operating the joystick or master arm which guides the remote manipulator. This has not been achieved successfully, and most success has been with displaying remote tactile information to the eyes using a computer-graphic display, or to skin at some other location.

Kinesthesia and proprioception are terms often used together, at least in part because the same receptors in the human body's muscles and tendons mediate both. Kinesthesia literally is the sense of motion and proprioception is awareness of where in space one's limbs are. Telekinesthesia and teleproprioception are particularly critical because, as telemanipulation experience has shown, it is very easy for the operator to lose track of the relative position and orientation of the remote arms and hands and how fast they are moving in what direction. This is particularly aggravated by his having to observe the remote manipulation through video without peripheral vision or very good depth perception, or by not having master-slave position correspondence, i.e., when a joystick is used. Potential remedies are: multiple views; wide field of view from a vantage point which includes the arm base; and computer-generated images of various kinds (the latter will be discussed further below). Providing better sense of depth is critical to telemanipulation in space.

Proximity sensing is not something humans normally do except by vision, but cats do it by whiskers or olfaction (smell), and bats and blind persons do it by sound cues or vibrations felt on their face. Sonar, of course, will not work in space. Electromagnetic and optical systems can be used for measuring proximity (close-in ranging) to avoid obstacles or decide when to slow down in approaching an object to be manipulated (Bejczy et al. 1980). Such auxiliary information can be displayed to the eyes by means of a computer-graphic display, or, if the eyes are considered overloaded, by sound patterns, especially computer-generated speech. We need to understand how best to use such information in space.

TELEACTUATING

It was stated in the previous section that we know relatively little about certain types of remote sensing, i.e., both artificial sensing and display to the human operator controlling the teleoperator (this in spite of knowing a great deal about human sensing *in vivo*). Remote

actuation (in which terms we include control in the conventional sense) poses an even larger problem, since it combines motor actuation with sensing and decision-making, and it can be said we know even less about this, except for the practical knowledge we have from operating the kinds of teleoperators that have been around for a number of years, mostly in nuclear hot-laboratories and for undersea oil operations. Again, comments are offered in a number of specific categories where some research is ongoing but much more needs to be done. The control problems in this category, where computer interaction per se is not the principal issue, apply to both direct and supervisory control.

Multi-degree-of-freedom end-effectors seem a most obvious need, as evidenced by our own human hands, but the sad fact is that these have not been developed beyond a few laboratory prototypes. Commercial manipulators tend to have simple parallel-jaw grippers, and a few have claws, magnetic or air-suction gripping mechanisms, or special purpose attachment devices for welding, paint spraying or other special-purpose tools. Though parallel-jaw gripping seems the most obvious function for a one DOF end-effector, it is not yet clear what a second DOF might be for, or a third, etc. Multi-fingered devices such as those by Salisbury (1986) or Jacobson (1987) will help us answer these questions. At the moment fear of losing objects in space seems to militate against general purpose grippers; that could change in the future. Modern computer-graphic workstations begin to offer the hope of studying problems like these by computer simulation without having to build expensive hardware for every configuration and geometric relationship to be tested.

Two-arm interaction is a necessity for much human manipulation (it has become standard for nuclear hot-lab manipulators), but we rarely see it in industrial or undersea teleoperators. Part of this problem is to get the most out of a given number of degrees-of-freedom. For example, instead of having a single six-axis arm operating on one body relative to a second body (or base), one might accomplish the same by having a three DOF "grabber arm" position the body so that a second, say, three DOF arm can work in coordinated fashion to perform some assembly task. Industrial robot experience shows that two three DOF arms are likely to be simpler and cheaper than one six-DOF arm. This has not been implemented in space applications; the problem needs research.

Redundant DOF Hand-arm-vehicle coordination is a serious problem, and actually a need for any kinematic linkage of more than six DOF which must be controlled in a coordinated way. This is largely an unsolved theoretical problem, at least in part because the number of configurations which satisfy given end-point position/orientation constraints is infinite. One tries to select from among these solutions to minimize energy or time or to avoid certain absolute positions of the joints, or to prevent singularities, etc., but the mathematics is formidable. One arm of three and one of four DOF make for such redundancy, but perhaps even more important, so does a vehicle thrusting in six DOF with an attached arm of even one DOF. We humans coordinate movements of our own legs, arms, and bodies (many redundant

DOF) without difficulty, but just how we do it is still a relatively well-kept secret of nature.

Multi-person cooperative control is one way to control a complex multi-DOF teleoperator--where each of several operators is responsible for maneuvering a single arm or vehicle in relation to others. Is this best or is it better to have a single operator control all DOF of both vehicle and arm? We really don't know. Results from simple tracking experiments suggest that control of multiple independent tasks is very difficult for one person. When the degrees of freedom of a task are closely coupled and/or must be coordinated to achieve the task objectives, that can be relatively easy provided proper control means are provided--but up to how many DOF? It is surprising how little research is available in this area.

Adjustable impedance of master and/or slave is a promising way of making a master-slave teleoperator more versatile than if the compliance-viscosity-inertance parameters remained fixed (Raju, 1986). A carpenter may carry and use within one task several different hammers, and a golfer many clubs, because each provides an impedance characteristic appropriate for particular tasks which are expected. Carrying many teleoperators into space may be avoided by making the impedance between slave and task and/or between human and master be adjustable. We have hardly begun to understand this problem, and have much to learn.

Interchangeable end-effector tools is another way to accomplish versatility, and of course is precisely what carpenters, surgeons or other craftsmen use. Future space teleoperators may have a great variety of special tools for both modifying and measuring the environment. It is not clear how to make the trade between special and general purpose end-effectors.

Task-resolved manipulation means performing standard or preprogrammed operations (e.g., cleaning, inspecting, indexing a tool) relative to the surface of an environmental object (Yoerger, 1986). This means sensing that surface in the process of manipulating and continually performing coordinate transformations to update the axes with respect to which the operations are being done. This is an extension of end-point resolution--ability to command the finger to move in a desired trajectory without having to worry about how to move all the joints in between.

Force-feedback with time delay has been shown both theoretically and experimentally not to work if the force is fed back continuously to the same hand as is operating the controller, for the delayed feedback simply forces an instability on the process which the operator might otherwise avoid by a move-and-wait strategy or by supervisory control (Ferrell, 1966). Yet it seems that forces suddenly encountered or greater than a preset magnitude might be fed back to that hand for a brief period, provided the forward gain were reduced or cut off during that same brief period, and the master then repositioned to where it was at the start of the event with no force-feedback.

Computer-aiding in Supervisory Control

Computers may be used for relatively "low-level" computations in many of the telesensing/display and teleactuation modes described above. There are a number of other teleoperation research problems in which the human-computer interaction is the important part. These include computer simulation, computer-based planning/decision-aiding, and computer-aided command/communication/control in various mixes. All of these are part of supervisory control by a human operator of a telerobot.

Off-line, real-time, human-operable ("flyable") simulation of teleoperation for research, engineering or training has barely begun to be viable. This is because of the complexity of simulating and displaying the vehicle plus the arm and hand plus the manipulated object plus the environment, having all degrees of freedom operate, with removal of hidden lines, and so on. Even nominally high-quality computer-graphics machines have trouble with generation of such complex displays in real time. We can come close today, but since computer power is the one thing that is bound to improve dramatically over the course of the coming few years, we might pay attention to the many possibilities for using computers as a substitute for building expensive hardware to perform man-machine experiments and evaluate new design configurations. There are serious problems to simulate the full dynamics of multi DOF arms and hands. There are problems to be solved to make simulated teleoperators grasp and manipulate simulated objects. There are many problems to get high quality pictures (in terms of resolution, frame rate, gray-scale, color, etc.) Telepresence is an ideal in simulators just as it is in actuality. In fact, to enable the human operator to feel he is "there" when "there" exists nowhere other than in the computer poses a particularly interesting challenge.

On-Line in-situ planning simulators might be used "in the heat of battle" to try out maneuvers just before they are committed for real action (and real expenditure of precious resources in space). In this case commands would be sent to the computer-based model of the vehicle and/or manipulator and these would be observed by the operator prospectively, i.e., before further commands are given (as compared to the retrospective state estimation case to be described below). Commands (supervisory or direct) would be given to the simulation model but not to the actual process, the model results would be observed, and the process could be repeated until the operator is satisfied that he knows what commands are best to commit to the actual process. There are possibilities for having the simulator "track" the movement of the actual process so that any on-line tests could start from automatically updated initial conditions. The problem of what to control manually and what to have the computer execute by following supervisory instruction is something that cannot be solved in general but probably must be decided in each new context; the on-line planning simulator might be a way to make this happen.

On-line simulation for time-delay compensation is appropriate only to direct control, and is not necessary for supervisory control. Here

the commands are sent to the model and the actual system at the same time. The model's prediction (e.g., in the form of a stick figure arm or vehicle) can be superposed on top of the actual video picture delayed in its return from space. The operator can observe the results from the model immediately (before the time delay runs its course), thereby be much more confident in his move before stopping for feedback, and thus save several "move-and-wait" cycles. These techniques have been demonstrated for models of the manipulator arm (Noyes and Sheridan, 1984), but not yet for the manipulator arm and controlled vehicle in combination. When the motion of vehicles or other objects not under the operator's control can be predicted, e.g., by the operator indicating on each or several successive frames where certain reference points are, these objects can be added to the predictor display. With any of these planning/prediction aids, the display can be presented from any point of view relative to the manipulator/vehicle--a feat which is not possible with the actual video camera.

State measurement/estimation/display has potential where all information about what is going on "right now" is not available in convenient form, or where measurements are subject to bias or noise, or multiple measurements may conflict. The purpose is to provide a best estimate of the current situation or "state" (values of key variables which indicate where the telemanipulator end effector is relative to reference coordinates or to environmental objects of interest, what are the joint angles and joint angle velocities, what is the level of energy or other critical resources, and so on) and display this to the human operator in a way which is meaningful and usable by him for purposes of control. This may mean combining information from multiple measurement or data-base sources, then debiasing this information to the extent that can be done (in light of available calibration data), and factoring in prediction of what the state should be based on knowledge of what recent inputs were and what are the likely system responses to these inputs. A complete state estimation yields a "best" probability density distribution over all system states. Much theory is available on state estimation but there has been almost no application to space teleoperation. Some research has shown that human operators are unable to assimilate state information that is too complex, and tend to simplify it for themselves by estimating averages and throwing away the full distribution, or at least by using some simple index of dispersion, or in the case of joint distributions over two or more variables by considering only the marginal distributions, or even simplifying to point estimates on the independent variables (Roseborough, 1986). Research is needed on how to provide the operator all that can be got from state estimation and how to display this in a meaningful way.

Supervisory command languages must be developed especially for space teleoperators. We have a good start from industrial robot command languages (Paul, 1981) and from the few experimental supervisory command languages which have been developed in the laboratory (Brooks, 1979; Yoerger, 1982). We must understand better the relative roles of analogic instruction (positioning a control device in space, pointing,

demonstrating a movement) and symbolic instruction (entering strings of alphanumeric symbols in more or less natural language to convey logic, description, contingencies, etc.). Clearly in everyday discourse we use both analogic and symbolic coding in communicating with one another, especially in teaching craft skills, which seem to relate closely to what teleoperation is. Both communication modes must be used in communicating with a telerobot. The telerobot usually starts with little or no "context" about the world, which objects are which and where they are in space. For this reason, it is necessary to touch objects with a designated reference point on the teleoperator, to point with a laser beam or otherwise to identify objects (perhaps concurrently with giving names or reference information symbolically), and to specify reference points on those objects. Recent progress in computer linguistics can contribute much to supervisory command language.

Voice control and feedback, for all the times it has been suggested as an interesting telemanipulation research topic in recent years, has seen very little systematic research. Voice command probably has the most promise for giving "symbolic" commands to the computer (in contrast to the normal "analogic" or geometric isomorphic commands which the master-slave or joystick provides). Vocal symbolic commands might be used to reset certain automatic or supervisory loops such as grasp force, or to set control gain, master-slave amplitude or force ratio, or to guide the pan, tilt and zoom of the video cameras (Bejczy et al., 1980).

Aids for failure detection/identification/emergency response are particularly important since in a complex system the human operator may have great difficulty knowing when some component has begun to fail. This can be because the component isn't being operated and hence there is no abnormal variable indicated. Alternatively, if it is being operated, the variables being presented as abnormal could have resulted from an abnormality well upstream. Finally, the operator can simply be overloaded. Many new failure detection/diagnosis techniques have been developed in recent years, some of them involving Bayesian and other statistical inference, some involving multiple comparisons of measured signals to on-line models of what normal response should look like, and so on. Failure detection/diagnosis is a critical part of supervisory control, where the operator depends on help from the computer, but himself plays ultimate judge. This may be a prime candidate for the use of expert systems.

Meta-analysis of Human/Computer/Teleoperator/Task Interaction

Abstract theory of manipulation and mechanical tool-using has been surprisingly lacking. Control engineering, as it developed through the 1940-60 period, never really coped with the complex sequential dependencies of coordinating sensory and motor activities to perform mechanical multi-DOF manipulation tasks. Only when industrial robot engineers began to face up to how little they knew about how to do assembly did the need for a theory of manipulation become evident.

Somehow it seems reasonable that the syntax of manipulation is analogous to that of natural language (i.e., tool-action-object corresponds to subject-verb-object, with appropriate modifiers for each term), since both are primitive human behaviors. It then seems a small step to apply computational linguistics to manipulation. But little of this has been done as yet.

Performance measures and assessment techniques need to be developed for teleoperation. At the moment there are essentially no accepted standards for asserting that one telemanipulator system (of hardware or software or both) is better or worse than some other. Of course to some extent this is context dependent, and the success will depend upon specific mission requirements. But there have got to be some generic and commonly accepted indices of performance developed which could be used to profile the capabilities of a teleoperator vehicle/manipulator system, including factors of physical size, strength, speed, accuracy, repeatability, versatility, reliability, etc. One worries whether even terms such as accuracy, repeatability, linearity, and so on are used in a common way within the community. No one is asking for rigid standardization, but some commonality across tests and measures appears necessary to avoid great waste and bureaucratic chaos.

Direct experimental comparisons between astronauts performing hands-on in EVA and teleoperators, performing either in direct or supervisory-controlled fashion, must be done on a much more extensive and scientifically controlled scale, making use of both the manipulation theory and the generic performance measures to be developed. These experiments should be performed first on the ground in laboratories or neutral buoyancy tanks, much as Akin (1987) has begun, then in space on shuttle flights (e.g., EASE experiments), and eventually on the space station itself.

CONCLUSIONS

A number of research topics have been proposed, all seen as critical for the development of needed teleoperator/telerobotic capability for future space station and related missions. These have been presented in the areas of: (1) telesensing (with the longterm goal of telepresence); (2) actuation (with the long term goals of versatility and dexterity); (3) computer-aiding in supervisory control (with the long term goals of providing better simulation, planning and failure detection tools, and telerobots which are reliable and efficient in time and energy); (4) meta-theory of manipulation (with the long-term goals of understanding, evaluation, and best relative use of both human and machine resources).

Telerobotics, as much as any other research area for the space station, has direct research transferability to the non-government sector for use in manufacturing, construction, mining, agriculture, medicine and other areas which can improve our nation's productivity.

REFERENCES

- Akin, D.
1987 Ongoing research at MIT.
- Bejczy, A. K., Brown, J. W. and Lewis, J. L.
1980 Evaluation of "smart" sensor displays for multidimensional precision control of space shuttle remote manipulator. Proceedings of the 16th Conference on Manual Control. MIT, Cambridge, MA, May 5-7.
- Bejczy, A. K., Dotson, R. S. and Mathur, F. P.
1980 Man-machine speech interaction in a teleoperator environment. Proceedings of the Symposium on Voice Interactive Systems. DOD Human Factors Group, Dallas, TX, May 11-13.
- Brooks, T. L.
1979 SUPERMAN: a System for Supervisory Manipulation and the Study of Human-Computer Interactions. SM thesis, Cambridge, MA: MIT.
- Corker, K. and Bejczy, A. K.
1985 Recent advances in telepresence technology development. Proceedings of the 22nd Space Congress. Kennedy Space Center, FL, April 22-25
- Ferrell, W. R.
1965 Remote manipulation with transmission delay. IEEE Transactions, Human Factors in Electronics. HFE-6, 1.
1966 Delayed force feedback. Pp. 449-455 in Human Factors. October.
- Ferrell, W. R. and Sheridan, T. B.
1967 Supervisory control of remote manipulation. IEEE Spectrum 4(10):81-88.
- Goertz, R. C. and Thompson, R. C.
1954 Electronically controlled manipulator. Pp. 46-47 in Nucleonics.
- Harmon, L. D.
1982 Automated Tactile sensing. International Journal of Robotics Research 1(2):3-32.
- Heer, E., ed.
1973 Remotely Manned Systems. Pasadena, CA: California Institute of Technology.

- Jacobson, S.
1987 Ongoing research at the University of Utah.
- Johnsen, E. G. and Corliss, W. R.
1967 Teleoperators and Human Augmentation. NASA SP-5047.
- Lynch, P. M.
1972 A Force Reflecting Joystick. Master's Thesis, Department of Mechanical Engineering, MIT.
- Mosher, R. S.
1964 Industrial manipulators. Scientific American 211(4):88-96.
- Noyes, M. and Sheridan, T. B.
1984 A novel predictor for telemanipulation through a time delay. Proceedings of the Annual Conference on Manual Control. Moffett Field, CA: NASA Ames Research Center.
- Paul, R. P.
1981 Robot Manipulators: Programming and Control. Cambridge, MA: MIT Press.
- Raju, J. G.
1986 An experimental master-slave manipulator system to study the feasibility of operator-adjustable impedance in remote manipulation. Man-Machine Systems Laboratory Memo 86-1, MIT.
- Roseborough, J. B. and Sheridan, T. B.
1986 Aiding human operators with state estimates. Man-Machine Systems Laboratory Report. MIT, July.
- Salisbury, J. K.
1986 ongoing research at the MIT AI Laboratory.
- Schneider, J. and Sheridan, T. B.
1984 An Optical Tactile Sensor for Manipulators, Robotics and Computer-Integrated Manufacturing 1(1):65-74.
- Strickler, T. G.
1966 Design of an Optical Touch Sensing System for a Remote Manipulator. SM thesis, Cambridge, MA: MIT.
- Yoerger, D.
1982 Supervisory Control of Underwater Telemanipulators: Design and Experiment. PhD thesis, Cambridge, MA: MIT.
- Yoerger, D.
Personal communication.

TELEROBOTICS FOR THE EVOLVING SPACE STATION:
RESEARCH NEEDS AND OUTSTANDING PROBLEMS

Lawrence Stark

INTRODUCTION

The definition of telerobotics (TR) has not yet stabilized nor made the standard English language dictionary. I tend to use telerobotics as meaning remote control of robots by a human operator using supervisory and some direct control. Thus, this is an important area for the NASA evolving space station. By robot, I mean a manipulator/mobility device with visual or other senses. I do not name manipulators, as in many industrial automation set-ups, robots even if they can be flexibly programmed; rather calling these programmable manipulators. Our own laboratory at the University of California, Berkeley, has been involved in problems in display of information to the human operator, in problems of control of remote manipulators by the human operator, and in communication delays and band-width limitations as influencing both control and the display. A number of recent reviews have appeared with discussions of the history of telerobotics beginning with nuclear plants and underseas oil rigs.

THREE SIMULTANEOUS RESEARCH DIRECTIONS

I believe that we should engage in triplicate or three way planning. It is important to carry out our research to accomplish tasks (i) with man alone, if possible, such as in EVA (extravehicular activities), (ii) with autonomous robots (AR), and (iii) with telerobotics. By comparing and contrasting the research necessary to carry out these three approaches, we may clarify our present problems. (See Table 1)

There are problems using man alone. The space environment is hazardous. It is very expensive to have a man in space; NASA must have quite adequate cost figures obtained from the demonstration projects that have already been accomplished with the shuttle program. We may also need a higher quality of performance than man alone can provide in terms of strength, resistance to fatigue, vigilance, and in meeting special problems. For example, if the space suit is not of constant volume under flexible changes of the limbs, then a great deal of strength is used up just in maintaining posture.

TABLE 1 Triplicate Planning

 Problems with man alone

Hazardous environment:
 (space similar to nuclear plants, underseas)
 Expensive (i.e. EVA in space)
 Need increased quality in
 Strength
 Fatigue resistance
 Vigilance
 Performance

 Problems with Autonomous Robots

Not yet available
 Design not fixed
 Feasibility not certain
 Reliability not tested

Therefore: TR is a viable leading edge technology

All three directions should be supported for evolving space station planning, research, and development.

Problems with autonomous robots lie in our not having mastered the technology to build them and have them perform satisfactorily. They are not yet available! Indeed, designs are not yet fixed and it is not certain how feasible they will be, especially in terms of robustness and reliability.

Therefore, we can see that telerobotics is a viable leading edge technology. However, all three directions should be intensively pursued in research and development, especially for the next stages of the evolving space station planning.

SPACE STATION TASKS

One of the major roles that NASA can play is to hypothesize tasks for the evolving space station. In this way research regarding the design of telerobots to accomplish these tasks can be guided. For a list of seven groups of tasks see Table 2.

TABLE 2 NASA should Hypothesize TASKS f. r Evolving Space Station

Housekeeping

Life support systems
 Inventory control, access and storage
 Record keeping
 Garbage disposal

Protection

From space garbage
 From meteorites
 From traffic flow

Maintenance

Satellite
 Vehicles
 Space station itself

Construction

Additional space station structures

Manufacturing

Crystal growth, biopharmaceuticals

Mobility

Automatic piloting
 Navigation
 Path planning

Scientific

Landsat type image processing for agriculture
 Meteorology
 Astronomy
 Human factors research
 Scientific record keeping

As I will consider later, it is important to distinguish between those tasks unique to the NASA/evolving Space Station and those with "industrial drivers" that will accomplish development of new technologies in hopefully a superior fashion and thus enable conservation of limited NASA resources.

PROBLEMS IN TELEROBOTICS

First I overview problems in telerobotics: those concerning displays, vision and other senses (Table 3) and those dealing with control and communication (Table 4).

In each table, I start with basic properties of the human operator and end up with planned capabilities of autonomous robots. In between, I try to cover what knowledge exists now in our field of telerobotics.

Experimental Set-Up for Three-Axis Pick-and Place Tasks

A teleoperation simulator constructed with a display, joysticks, and a computer enabled three-axis pick-and-place tasks to be performed and various display and control conditions evaluated (Figure 1). A vector display system (Hewlett-Packard 1345A) was used for fast vector drawing and updating with high resolution. In our experiments, displacement joysticks were mainly used, although in one experiment a force joystick was used to compare with a displacement joystick. An LSI-11/23 computer with the RT-11 operating system computer was connected to the joystick outputs through 12-bit A/D converters, and to the vector display system through a 16-bit parallel I/O port.

TABLE 3 Display Problems for the Human Operator

Display graphics (raster/vector)
On-the-screen enhancements
On-the-scene enhancements
Other senses displayed
Inputs to other senses
Perspective and Stereo Displays
Task performance criteria
Helmet Mounted Display
Telepresence; space constancy
Human Operator (H.O.) Performance
Fatigue, effort, vigilance
Robotic Vision
LLV - Chips
MLV - blockworld and hidden lines
HLV - ICM, AI

TABLE 4 Control and Communication Problems for the Human Operator

Basic properties of H.O., especially for EVA task performance
Nerve, muscle, AG/AT model
Sampled-data (SD) and adaptive control
Prediction, preview, optimal control--Kalman filter
H.O. control of vehicles, manual control
H.O. control of TR
H.O. special control:
Preview, delay, bilateral, homeomorphic control
Locomotion (human, robotic):
Navigation--pathways
Potential field algorithms
HLC (high level control):
Supervisory control
Multiperson cooperative control; RCCL; fuzzy sets
Autonomous robotic (AR) control
Sensory feedback, adaptive control, AI

A typical presentation on the display screen for three-axis pick-and-place tasks included a cylindrical manipulator, objects to pick up, and boxes in which to place them, all displayed in perspective (Figure 2). Since perspective projection alone is not sufficient to present three-dimensional information on the two-dimensional screen, a grid representing a horizontal base plane and reference lines indicating vertical separations from the base plane are also presented (Ellis et al., 1985; Kim et al., 1985 submitted). The human operator controlled the manipulator on the display using two joysticks to pick up each object with the manipulator gripper and place it in the corresponding box. One hand, using two axes of one joystick, controls the gripper position for the two axes parallel to the horizontal base plane (grid). The other hand, using one axis of the other joystick, controls the gripper position for the third axis (vertical height) perpendicular to the base plane. Picking up an object is accomplished by touching an object with the manipulator gripper. Likewise, placing an object is accomplished by touching the correct box with the manipulator gripper.




FIGURE 1 Experimental arrangement.

Puma Arm Simulator

In addition to the cylindrical manipulator simulation, the kinematics and dynamics of a six degree-of-freedom Puma robot arm were simulated. Each of these degrees of freedom were controlled simultaneously using two joysticks. Although no experiments have yet been performed with the puma simulation, it is hoped that it will be a step toward experiments with more complex manipulators. A low-bandwidth telephone connection to control two Puma arms at Jet Propulsion Labs in Pasadena is planned. The simulation will allow prediction of the robots' motion to provide a preview display to help overcome the communication delays inherent in such a low bandwidth connection, or as in transmissions to manipulators in space.




FIGURE 2 Ames-Berkeley visual enhancement display.

Helmet Mounted Display Design

Motivation

The motivation of the HMD system is to provide the human operator with a telepresence feeling that he is actually in the remote site and controls the telemanipulator directly. The HMD system detects the human operator's head motion, and controls the remote stereo camera accordingly. In our current system, the remote telemanipulation task environment is simulated and the pictures for the display are generated by the computer.

Head Orientation Sensors

A two-axis magnetic Helmholtz coil arrangement was used as a head orientation sensing device, to detect horizontal and vertical head rotations (Figure 3). By assuming that the pan and tilt angles of a remote stereo camera are controlled in accordance with the horizontal and vertical head rotations, respectively, the computer generates the corresponding stereo picture for the HMD. The head orientation sensing

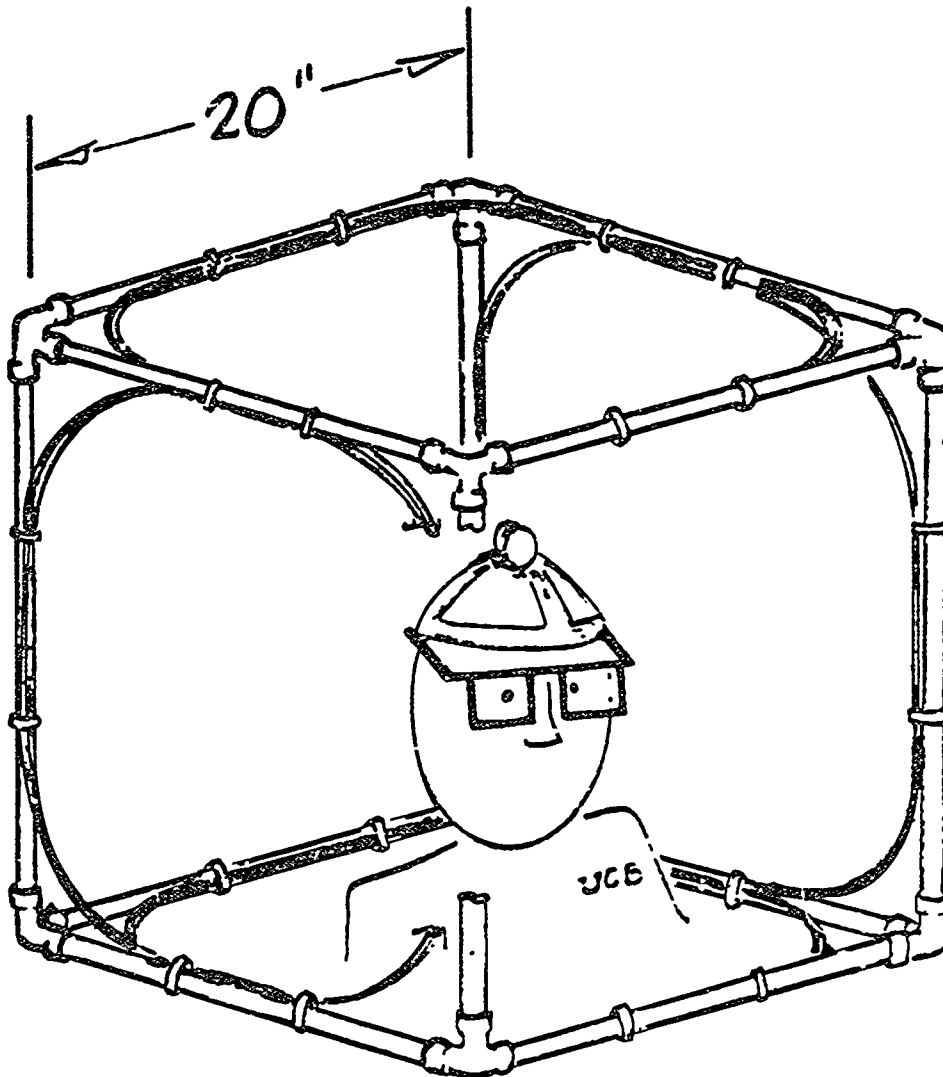


FIGURE 3 Head orientation sensor device.

device is composed of a search (sensing) coil mounted on or beneath the helmet and two pairs of field coils fixed with respect to the human operator's control station. The right-left pair of the field coil generates the horizontal magnetic flux of a 50 KHz square wave. The up-down pair of the field coil generates the vertical magnetic flux of a 75 KHz square wave. The search coil detects the induced magnetic flux, which is amplified and separated into 50 and 75 KHz components. The magnitude of each frequency component depends upon the orientation of the search coil with respect to the corresponding field coil (Duffy, 1985).

LCD Display

An early configuration of the HMD had a flat-panel LCD (liquid crystal display) screen (a commercially available portable LCD television) mounted on the helmet for the display (Figure 4). However, the picture quality of the LCD screen was poor due not only to low resolution but also to poor contrast.

CRT Display

A new design of the HMD that we currently have, mounted a pair of Sony viewfinders (Model VF-208) on the helmet (Figure 5). Each viewfinder has a 1-inch CRT (cathode ray tube) screen and a converging lens through which the human operator views the CRT screen. The computer-generated stereo picture pair (stereogram) is displayed on the CRT screens; one for the left eye and the other for the right. The converging lens forms the virtual image of the stereogram behind the actual display screen. When the CRT screen is 4.2 cm apart from the lens whose focal length is 5 cm, the virtual image of the CRT screen is formed at 25 cm apart from the lens with an image magnification of 6. Thus, a 1-inch CRT screen appears to be a 6-inch screen to the viewer. At appropriate geometrical and optical conditions, the right and left images overlay, and most people can fuse the two images into a single three-dimensional image. The stereoscopic display formulas used to generate the stereogram for the helmet mounted display are described in references (Kim et al., 1987).

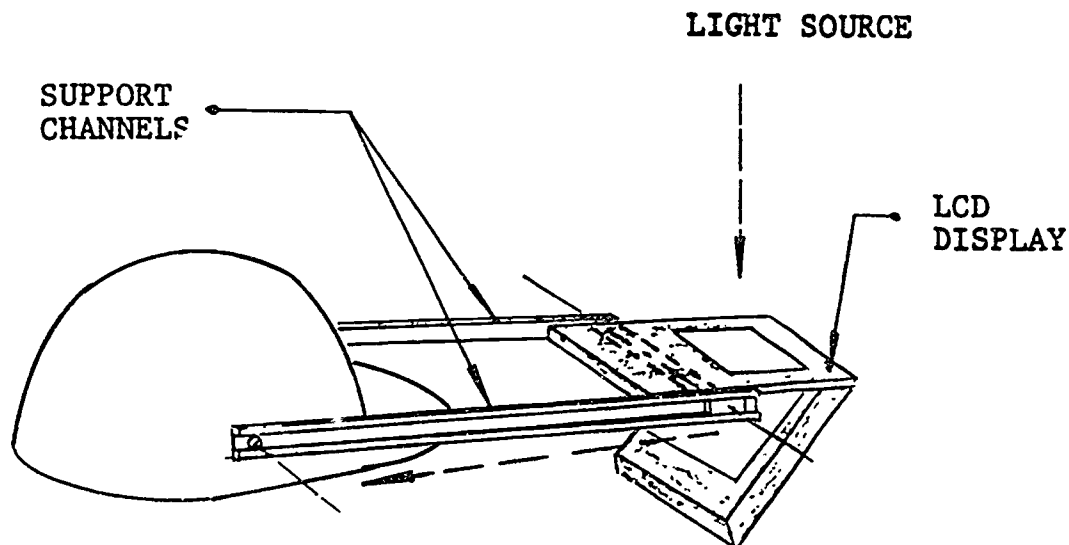


FIGURE 4: Early HMD design with LCD screen.




FIGURE 5 Current HMD design.

Mechanical Design

Five degrees of freedom were provided for the mechanical adjustment of the position and orientation of each viewfinder, allowing three orthogonal slidings and two rotations (Figure 5). A 1 lb. counterweight was attached to the back of the helmet for counter-balancing.

Communication Delay and Preview

Communication delay is a significant constraint in human performance in controlling a remote manipulator. It has been shown (Sheridan et al, 1964, Sheridan, 1966; Tomizuka and Whitney, 1976) that preview information can be used to improve performance. Stark et al. (1987)

demonstrated that preview can significantly reduce error in tracking experiments with imposed delay.

Experiments were performed to investigate whether a preview display could improve performance in pick-and-place tasks with delay. A single bright diamond-shaped cursor was added to the display to represent current joystick position. This was a perfect prediction of what the end effector position would be after the delay interval. Thus, the task was the same as if there were no delay, except that the H.O. had to wait one delay period for confirmation that a target had been touched or correctly placed (in the non-previewed display, the target letter was doubled when picked up, and became single again when placed in the correct box).

Preview improved performance at delays up to 4 seconds so that it was almost as good as for a small delay of 0.2 seconds (Figure 6). While task completion time in the delayed condition increased greatly with delay, there was only a small increase in the preview case. This is because the H.O. must compensate for delays by using a "move-and-wait" strategy, making a joystick movement and waiting to see the resultant and effector movement. In the preview case, this

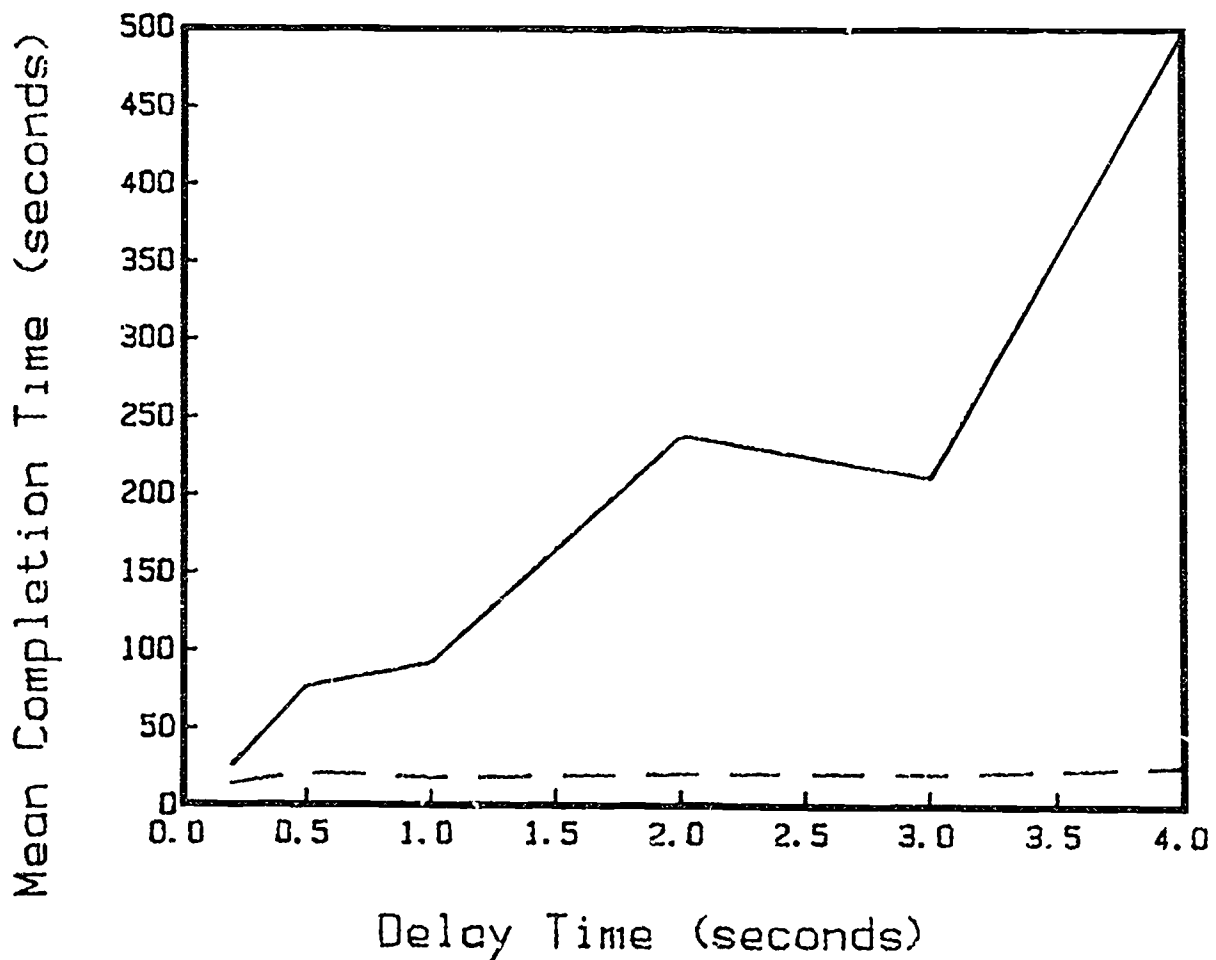


FIGURE 6 Performance affected by delays and by preview control mode.

strategy is only necessary when very close to the target or box to wait for confirmation that the goal has indeed been touched.

Control Mode Experiments

Position and rate controls are the two common manual control modes for controlling telemanipulators with joysticks (or hand controllers) (Johnsen and Corliss, 1971; Heer, 1973). In the position control the joystick command indicates the desired end effector position of the manipulator, whereas in the rate control the joystick command indicates the desired end effector velocity.

In our three-axis pick-and-place tasks, the human operator controls the manipulator hand position in the robot base Cartesian coordinate by using three axes of the two displacement joysticks. In pure (or ideal) position control, the system transfer function from the joystick displacement input to the actual manipulator hand position output is a constant gain G_p for each axis. In pure rate control, the system transfer function is a single integrator G_v/s for each axis. In the rate control, a 5% dead-band nonlinearity is introduced before the pure integrator in order to inhibit the drift problem associated with the pure integrator.

Comparison of Pure Position and Rate Controls

Three-axis pick-and-place tasks were performed with both pure position and rate control modes for various gains (Figure 7). The mean completion time plot clearly shows that pick-and-place performance with pure position control (mean completion time 2.8 seconds at $G_p=2$) was about 1.5 times faster than that of the pure rate control (mean completion time 4.3 seconds at $G_v=4$).

Trajectories of Joystick and Manipulator Movements

In order to examine why the position control performed better than the rate control, several trajectories of the joystick displacement input and the manipulator hand position output during the pick-and-place operation were observed. Typical trajectories from the start of trying to pick up an object to its accomplishment were plotted to illustrate position, rate, and acceleration controls (Figure 8). Components only for the x-axis (side-to-side) are plotted, since components for the other two axes are similar. Observation of several trajectories indicates that a precise re-positioning of the manipulator hand is achieved by a combination of quick step re-positioning operations and slow smooth movement operations. In position control one quick step re-positioning of the manipulator hand from one position to another requires one joystick pull or push operation, whereas in the rate control it requires a pair of operations; pull-and-push or push-and-pull operations (Figure 8). This is a major reason why the




FIGURE 7 Performance comparison of position and rate control.

position control yielded better performance than the rate control for our pick-and-place tasks. It should be noted, however, that the pick-and-place task is a positioning task. If the task is following a target with a constant velocity, then velocity (rate) control would perform better.

Acceleration Control

Three-axis pick-and-place tasks were also tried with acceleration control. It turned out, however, acceleration control was not adequate to perform stable, safe pick-and-place operations. In acceleration control, the manipulator tends to move almost all the time even though the joystick is at the center position. Note that in pure rate




FIGURE 8 Position, rate and acceleration control.

control, the manipulator does not move when the joystick is at the center position regardless of previous history of the joystick displacement.

Human Adaptation to Gain Change

Mean completion time did not change much for the various gains tested (Figure 7), which means that the human operator adapted well to the gain change (McRuer et al., 1965; Young, 1969; Stark, 1968). Both lower and higher gains relative to the optimal gains caused slight increase in the mean completion time. A reason of slightly longer mean completion times with lower gains is because lower gains demand wider joystick displacements and it takes longer for the finger or hand to displace the joystick wider. A reason for slightly longer mean completion times with higher gains is that higher gains demand more minute joystick displacements, degrading effective resolution of the

joystick control. An additional major reason for longer mean completion times with lower gains for the rate control is due to the velocity limit.

Force Joystick

The two common joystick types are the displacement and force joysticks. The output of the displacement joystick is proportional to the joystick displacement, whereas the output of the force joystick (isometric or stiff joystick) is proportional to the force applied by the human operator. The advantage of the force joystick is that it requires only minute joystick displacements (a few micrometers) in contrast with the displacement joystick (a few centimeters).

Pick-and-place tasks were performed for pure position and rate controls with displacement and force joysticks. The experimental results for two subjects (Figure 9) shows that in the rate control, task performance with force joystick was significantly faster than that with displacement joystick. This is mainly because the force joystick senses the applied force directly, requiring only very minute joystick displacements. In the position control, however, the force joystick performed no better than the displacement joystick. In fact, all three subjects preferred to use the displacement joystick in this mode, since the force joystick required more force to be applied than the displacement joystick, especially when the manipulator hand is to be positioned far away from the initial center position. Position control also performed better than the rate control regardless of joystick types, and furthermore the position control with the displacement joystick performed best for our pick-and-place tasks (Figure 9).

Resolution

The experimental results demonstrate the superiority of position control when the telemanipulator has a sufficiently small work space (Figures 7, 8, & 9). Note that our three-axis pick-and-place tasks used in this experiment implicitly assumes that the manipulator work space is small or at least not very large, since our task allows the human operator to perform successful pick-and-place operations with a display showing the entire work space on the screen. Examples of small work space telemanipulators can be found in nuclear reactor teleoperators, surgical micro-telerobots, or small dexterous telerobotic hands. Position control can also be utilized during proximity operations in conjunction with the force-reflecting joysticks for enhanced telepresence (Bejczy, 1980). When the telemanipulator's work space is very large as compared to human operator's control space, position control of the entire work space suffers from poor resolution since human operator's control space must be greatly up-scaled to accommodate the telemanipulator's large work space (Flatau, 1973). One way of solving this poor resolution problem in position control is using indexing (Johnson and Corliss, 1971; Argonne National Lab,




FIGURE 9 Displacement and force joystick control.

1967). In the indexed position control mode, the control stick gain is selected so that the full displacement range of the control stick can cover only a small portion of the manipulator work space, and large movements of the manipulator hand can be made by successive uses of an indexing trigger mounted on the control stick. Note, however, that rate control can inherently provide any higher degree of resolution by mere change of control stick gain without use of indexing.

Homeomorphic Controller

Most of our pick-and-place and tracking experiments were performed with joysticks as the input device through which the human operator

controlled the simulated manipulator. The operator's movements when using joysticks are non-homeomorphic, so that the movements he must make to produce a desired manipulator response do not match the movement of the manipulator end effector. Thus, he must mentally convert the desired end effector position to Cartesian coordinates and use the joysticks to input these coordinates.

To attempt to study whether a truly homeomorphic input device could improve performance in tracking tasks, an apparatus of identical form to our simulated cylindrical manipulator was built. A vertical rod was supported by bearings on the base to allow rotation, θ . A counterweighted horizontal arm was attached to the rod with sliding bearings to permit rotation and translation in the r and z axes respectively. The human operator could control position through a handle on the end of the arm corresponding to the end effector of the simulated manipulator. Potentiometers measured movement in each axis to determine input r , θ , and z . The LSI-11/23 computer read these values through A/D channels and displayed the manipulator in the identical position.

Three-dimensional tracking experiments were performed with the homeomorphic controller and with joysticks for gains varying from 1 to 5 to compare performance (Figure 10). The results do not show a significant difference between the homeomorphic controller and joysticks over the range of gain values. Although the larger movements required for the homeomorphic controller, with greater inertia and friction than the joystick, may have limited performance, we believe that human adaptability minimizes its advantages.

Training by Optimal Control Example

A simplified simulation of the manned maneuvering unit, MMV, enabled study of training of human control performance (Jordan, 1985). Only three translatory degrees-of-freedom, x , y and z , were used. Thrusters generating pulses of acceleratory control were controlled via a keyboard and the task was to accelerate simultaneously in x , y and z to a maximum velocity, transit to the desired new location, and decelerate again simultaneously. Two displays were used--a perspective display of a minified model of the MMV, or two two-dimensional projectors of that model with a small inset of the perspective display.

Subjects generally performed poorly during the few hundred seconds allowed for the tasks (Figure 11a). It was decided to allow the subjects to view this control problem carried out by a simple optimal control algorithm (see Figure 11b). This experience was of considerable help and several subjects then performed quite well (Figure 11c).

This experiment, learning-by-example, illustrates a strategy that perhaps may be effective in more complex and realistic tasks as well.

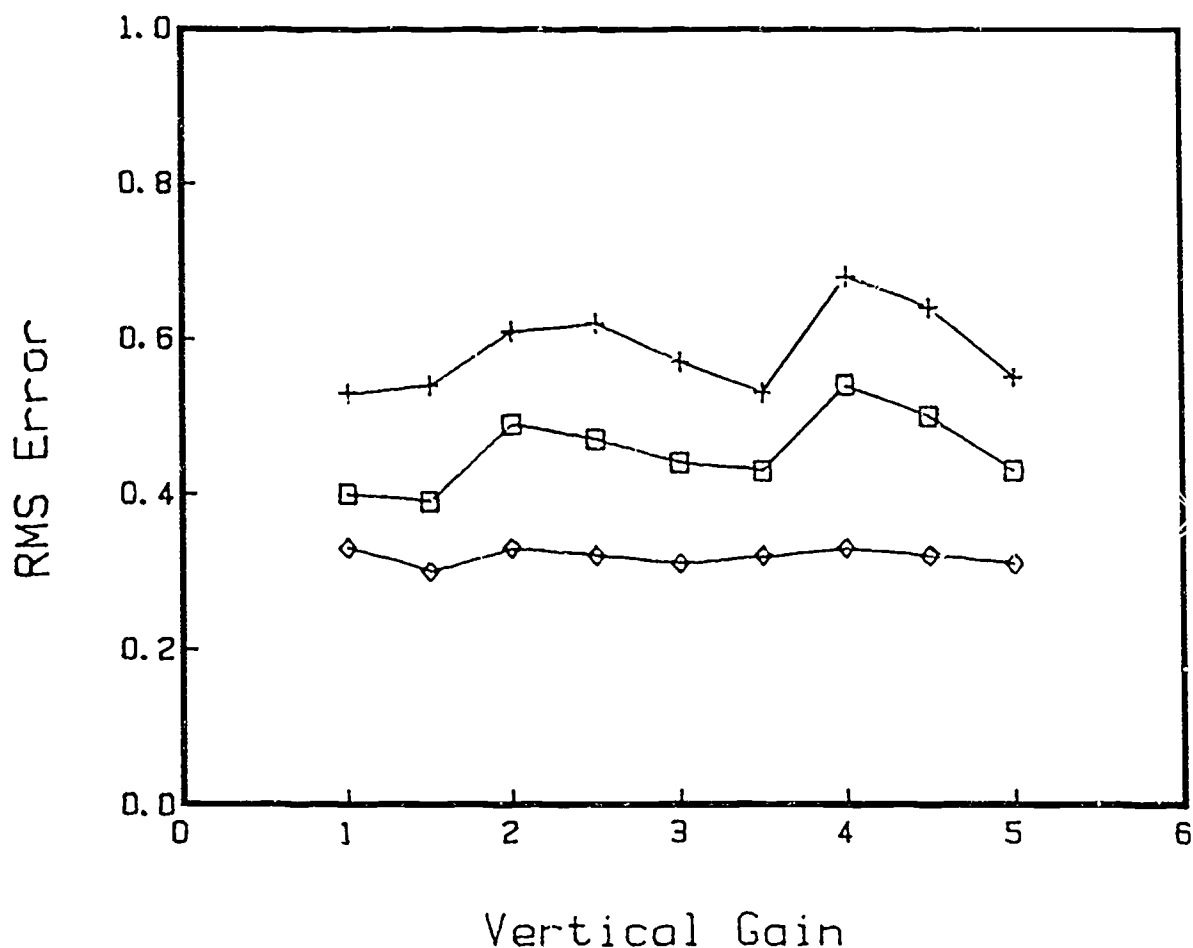











FIGURE 10 Homeomorphic Controller.

INDUSTRIAL DRIVERS FOR CERTAIN NECESSARY SPACE STATION TECHNOLOGIES

This next section deals with the future, and especially with "industrial drivers" other than NASA for new technologies which may be required in the evolving Space Station. In Table 5 I list nine components of a telerobotics system that certainly seem to be driven by important industrial hardware requirements, research and development. Therefore, it seems reasonable for NASA to sit back and wait for and evaluate these developments, saving its resources for those necessary technologies that will not be so driven.

Looking at these figures gives us some concept of how industrial development may provide various types of technologies for the evolving Space Station; indeed, NASA may be able to pick and choose from off-the-shelf items! For example, the most powerful computers on the last space shuttles were the hand-held portable computers that the




315

FIGURE 11a: Training by example (human operator control before automatic)

310

313










311


317

FIGURE 11b: Training by example (human operator control after automatic)


318



REAL TIME



REAL TIME



REAL TIME

312

319

FIGURE 11c: Training by example (automatic control results)

TABLE 5 Drivers other than NASA for Nine Needed Technologies

Robotic Manipulator and Control Scheme

Joystick - Aircraft

AR Manufacturing Industry, Nuclear Industry, Mining Industry,

Sensors: Force and Touch; compliant control

ROV and Mobility

Military, tanks and other vehicle plans?

Undersea ROV - Oil and Communications Industry

Locomotion - University Research

Shipping Industry: Ships at Sea [AR, TR, Man]

TV Camera

Entertainment Industry - commercial device

Security Industry

Need mounts, controls and motors for PAN, TILT and for Stereo VG

Graphics

Entertainment industry is a better driver than companies building

Flight Simulators;

HMD as an example.

EM sensors research/Head-Eye Mouse

ICM

Landsat

Security

Medical Industry - CT and MRI

Industrial Production Lines

TD - Image Understanding

Computer

Computer Industry

(HDW) and (SFW)

Computer Science research base is now very broad

Communication

Communication Industry is huge

Ships at Sea

BW Compression

Remote Oil Rigs

Arctic Stations

Plans and Protocols to Combat H.O. Fatigue and to Promote H.O.Vigilance

Office Automation Forces

Air Traffic Control Needs

Security Industry

Cooperative Control

Military - submarine control

Helicopter flight control

Air traffic controllers

Nuclear industry

Chemical plant industry

astronauts brought aboard which contained much greater capability than the on-board computers; those had been frozen in their design ten years ago in the planning stages for the space shuttle.

NECESSARY TELEROBOTICS TECHNOLOGIES TO BE SPARKED BY NASA

However, there are several areas in telerobotics that may likely not be driven independently of NASA, or where NASA may have an important role to play. Indeed, the Congress has specifically mandated that 10% of the Space Station budget should be used for Automation and Robotics development, and that this in some sense should spearhead industrial robotics in the United States (Table 6).

TABLE 6 Areas Sparked by NASA not Industrially Driven

Visual Enhancements for Graphic Display

Telepresence with Stereo Helmet Mounted Display (HMD)

Multisensory Input Ports:

Worry about H.O. overload condition
(especially with cooperative control and communication)

Higher Level Robotic Vision:

Example—Image Compression by Modeling (ICM)
(to require less information flow and faster update)

Special Control Modes for H.O.

Homeomorphic control
Bilateral control
Time delay and preview control for time delay
Compliant control

Higher Level Control Languages

(such as RCCL; fuzzy control; path planning by potential field construction)

Remote operating vehicles (ROV) special control problems:

Navigation, orientation, obstacle avoidance for ROV

Cooperative Control:

Cooperation amongst humans, telerobots, and autonomous robots

Compliant, Flexible, Homeomorphic Manipulators

Grasp versus tool using

Homeomorphic Dual Mode Control

Impedance Control

UNIVERSITY NASA RESEARCH

I now would like to make a plea that NASA should expand and stimulate telerobotics research conducted within the university environment. Of course, as a professor I may have a bias in this direction and I am willing to listen to contrary arguments! In addition to the benefits of the research accomplished by universities, NASA also gets the education and training of new engineering manpower specifically directed towards telerobotics, and focused on the evolving Space Station.

What kind of university and educational research should be funded in general by NASA. I believe there are two levels of cost (with however three directions) into which these educational research labs should be classified.

(i) First are Simulation Telerobotics Laboratories. Here we need graphics computers, perhaps joysticks, perhaps higher level supervisory control languages, cameras, image compression techniques and communication schemes. I would guess that our country needs at least thirty such systems for education and training. These systems should be very inexpensive, approximately \$50,000 each. They need not even be paid for by NASA, since universities can provide such research simulation laboratories out of their educational budgets or from small individual research grants. Our Telerobotics Unit at Berkeley has been thus funded. A good deal of exploratory research can be carried out inexpensively in this manner.

(ii) Second, we need Telerobotic Laboratories with physical manipulators present as important research components. In this way, experiments with various robotic manipulators, especially those with special control characteristics such as flexibility, homeomorphic form, new developments in graspers, and variable impedance control modes, other than are found in standard industrial manipulators, would be possible. I guess that there are about five such laboratories in some stage of development at major universities in the country. I would further estimate that these laboratories could each use an initial development budget of \$300,000 to enable them to purchase necessary hardware in addition to software as existent in the simulated Telerobotics Laboratories.

Another set of costly laboratories would be Telerobotics Laboratories with remote operating vehicles (ROV). Here again, we need about five laboratories at universities with first class engineering schools. Again, I estimate about \$300,000 each for the initial hardware support of these ROV labs. They could then study transfer vehicles, local Space Station vehicles, Moon/Mars Rovers, and even compare MMU vs. telerobotic controlled vehicles.

The university laboratories would contrast with and serve a different function than ongoing aerospace industrial laboratories, and NASA and other government laboratories. These latter assemble hardware for demonstration and feasibility studies. Then unfortunately they are somehow unable to carry out careful human factors research dealing with the changing design of such pieces of equipment. In the university setting, this apparatus could be taken apart, changed,

revitalized, modified and the flexibility would inform our current capability. I would like to contrast the Gossamer Condor and Gossamer Albatross with the NASA program. It was clear that if McCready was ever to be successful, he had to build an experimental plane which was expected to break down each experimental day. But the plane could be repaired in a few minutes! This "laboratory bench" concept is so different from twenty-year-ahead-planning currently controlling our space program that has been effectively eliminated at NASA. I think it is important to reintroduce rough and ready field laboratories back into the space program.

NASA PRIZES

Another role that NASA might play is to offer demonstration contracts or, even better, prizes for accomplishment of specific tasks. Again I turn to the Kremer Prize; here a private individual donated prize money to be awarded to the first to build a man-powered aircraft conforming to certain carefully laid out specifications.

Communication channels for controlling remote vehicles and remote manipulators are already set up. Thus we could have prize contestants demonstrating at differing locations on earth at one "g"; next demonstrations using elements capable of operating in space, or even more stringently, of having that minimum mass capable of being lifted into space; and then we might have true shuttle and space station demonstrations.

INTELLECTUAL PROBLEMS IN TR FOR THE SPACE STATION

Finally, I would like to leave you with the thought that the list of to-be-sparked-by-NASA problems in Table 6 contains many important intellectual problems facing the area of telerobotics. Although these areas are being approached in our research community at the present time, it may not be possible to foresee what novel kinds of challenges will face the evolving Space Station in twenty years. Even though I may not predict accurately, I certainly hope I am there in person to watch telerobotics playing a major role in operating the Space Station.

SUMMARY

The telerobotic, TR, system is a simulated distant robot with vision and manipulator and/or mobility subsystems controlled by a human operator, H.O. The H.O. is informed mainly by a visual display, but also by other sensors and other sensory displays, i.e. auditory, force or tactile. His control can be direct via joysticks, or supervisory via command and control primitives effected by partially autonomous robotic functions. Delays and bandwidth limitations in communication are key problems, complicating display and control (Stark et al., 1987).

Class experiments enabled our Telerobotic Unit at the University of California, Berkeley to explore in a number of research directions. The HMD direction has now been greatly extended and is a major focus in our laboratory. On the other hand, the homeomorphic controller did not seem to be a productive project to continue because of the adaptability of the H.O. to many configurations of control. Also, our interest in supervisory and other high level controls is leading us away from the direct manual control. The students taking a graduate control course, ME 210 "Biological Control Systems: Telerobotics," during the fall semester, 1985, in which the helmet mounted display, HMD, is emphasized, were enthusiastic and felt the course stimulated their creativity and provided an opportunity for them to engage in relatively unstructured laboratory work--a good model for subsequent thesis research.

ACKNOWLEDGEMENTS

We are pleased to acknowledge support from the NASA-Ames Research Center (Cooperative Agreement NOC 2-86) and the Jet Propulsion Laboratory, California Institute of Technology (Contract #956873).

We would also like to thank visiting lecturers from NASA-Ames; Mark Cohen, Stephen Ellis, Scott Fisher, Arthur Grunewald, John Perrone and Mordecai Velger; Drs. Won Soo Kim and Blake Hannaford, and Frank Tendick, Constance Ramos and Christopher Clark of University of California, Berkeley.

REFERENCES

- Argonne National Laboratory
1967 Manipulator Systems for Space Applications. Technical Report, Argonne.
- Bejczy, A. K.
1980 Sensors, controls, and man-machine interface for advance teleoperation. Science 208(4450):1327-1335.
- Duffy, M. K.
1985 A Head Monitor System Using the Search Coil Method. Master's thesis, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley.
- Ellis, S. R., Tyler, M., Kim, W. S., McGreevy, M. W., and Stark, L.
1985 Visual enhancements for perspective displays: perspective parameters. Pp. 815-818 in IEEE Proceedings of the International Conference on Systems, Man, and Cybernetics.

- Flatau, C. R.
1973 The manipulator as a means of extending our dexterous capabilities to larger and smaller scales. Pp. 47-50 in Proceedings of 21st Conference on Remote Systems Technology.
- Heer, E.
1973 Remotely Manned Systems: Exploration and Operation in Space. California Institute of Technology.
- Johnsen, E. G., and Corliss, W. R.
1971 Human Factors Applications in Teleoperator Design and Operation. Wiley-Interscience.
- Jordan, T.
1985 The Simulated Manned Maneuvering Unit and Pursuit Experiments. Master's thesis, Department of Mechanical Engineering, University of California, Berkeley.
- Kim, W. S., Ellis, S. R., Tyler, M., Hannaford, B., and Stark, L.
1987 A quantitative evaluation of perspective and stereoscopic displays in three-axis manual tracking tasks. IEEE Trans. on System, Man, and Cybernetics 16:61-72.
- Kim, W. S., Ellis, S. R., Tyler, M., and Stark, L.
1985 Visual enhancements for telerobotics. Pp. 807-811 in IEEE Proceedings of the Int. Conf. on System, Man, and Cybernetics.
- Kim, W. S., Tendick, F., and Stark, L.
1986 Visual enhancements in pick-and-place tasks: human operator's controlling a simulated cylindrical manipulator. Submitted to the IEEE Journal of Robotics and Automation.
- McRuer, D., Graham, D., Krendel, E., and Reisener, W.
1965 Human Pilot Dynamics in Compensatory Systems: Theory, Models, and Experiments with Controlled Elements and Forcing Function Variations. U. S. Air Force, AFFDL-TR-65-15.
- Sheridan, T. B.
1966 Three models of preview control. IEEE Transactions Human Factors in Electronics HFE-7:91-102.
- Sheridan, T. B., Merel, M. H., et al.
1964 Some predictive aspects of the human controller. Progress in Astronautics and Aeronautics. Vol. 13. Academic Press Inc., New York.
- Stark, L.
1968 Neurological Control Systems: Studies in Bioengineering. Plenum Press.

- Stark, L., Kim, W. S., Tendick, F., et al.
1987 Telerobotics: display, control and communication problems. IEEE Journal of Robotics and Automation RA-3(1):67-75.
- Tomizuka, M., and Whitney, D. E.
1976 The human operator in preview tracking: an experiment and its modeling via optimal control. Transactions in the ASME Journal of Dynamic Systems, Measurements, and Control 98:407-413.
- Young, L. R.
1969 On adaptive manual control. IEEE Transactions Man Machine Systems MMS-10(4):292-331.

DISCUSSION: COMMENTS ON TELEPRESENCE AND SUPERVISORY CONTROL

Antal K. Bejczy

Telepresence and supervisory control technologies, as Professor Sheridan pointed it out, represent advancements or refinements of existing teleoperator technology capabilities. Both technologies are strongly driven by and rely upon increased computer and computing capabilities and are regarded as substantial contributors to evolving space station capabilities in the sense of reducing EVA astronaut involvement in assembly, servicing and maintenance operations. Moreover, both technologies carry the promise of substantial spin-off for advancing capabilities of the U.S. production and service industries.

Professor Sheridan and Professor Stark enumerated and elucidated many specific topics and issues in sensing, controls and displays for telepresence and supervisory control which need research attention to advance the state of the art in the two technologies. In my discussion and comments, I would like to focus attention on the same research topics and issues from the following viewpoints: (a) In what sense and to what extent can we expect the enhancement of human operator capabilities through telepresence and supervisory control? (b) What specific conditions and constraints are imposed by the space application environment on the evolving telepresence and supervisory control technologies? (c) The multidisciplinary nature of the required research effort since neither telepresence nor supervisory control are intrinsically separate science or engineering disciplines. A brief description of the basic objectives of telepresence and supervisory control technologies may help illuminate the questions that arise from the above three viewpoints.

The basic objective of telepresence technology is to alleviate the human operator's sense of remoteness in the control station by providing sufficient information to the operator on the activities of the remote machine in usable form. The content of the last attribute "usable form" heavily depends on human capabilities under given conditions, on the capabilities and characteristics of machines to be controlled, and on the nature of tasks to be accomplished. Also implied in this technology is the operator's enhanced control response ability to the perceived remote events. Briefly, telepresence technology is aimed at providing—so to speak—a more intimate,

sensitive and high fidelity input and output connection between operator and remote machine.

The basic objective of supervisory control technology is to provide sufficient capabilities for the human operator to tell the remote machine what to do and, eventually, how to do it, without involving the operator in continuous control coordination of a multitude of machine actuators needed to execute a task (note that a dual-arm system contains fourteen or more actuators). Thus, in supervisory mode of control, the operator controls the task instead of controlling the individual degrees of freedom and associated actuators of a multi-degree-of-freedom complex machine. Implied in this technology are two important technical capabilities: (a) flexible automation of actions of a multi-degree-of-freedom complex mechanical system, and (b) flexible language-like or menu-type interface to, or interaction with, the automated mechanical actions of a remote machine.

Several notes should be added to the objective descriptions of telepresence and supervisory control technologies. First, none of them eliminates the human operator from the operation, but both change the operator's function assignments and employ human capabilities in new ways. Second, both technologies promise the performance of more tasks with better results, but, in doing so, both technologies also make a close reference to human capabilities of operators who will use evolving new devices and techniques in the control station. Third, both telepresence and supervisory control technologies make reference to evolving capabilities of other technologies like sensing, high performance computer graphics, new electro-mechanical devices, computer-based flexible automation, expert systems for planning and error recovery, and so on. Thus, the progress in both technologies are tied to rich multidisciplinary activities. Fourth, both technologies require the evaluation and validation of their results relative to the application environment. For space station scenarios this implies the effect of zero-g on human operators, restricted local resources (like power, work volume, etc.) for a control station in Earth orbit, limited communication bandwidth and some communication time delay between a control station and remote machines, fragile and sensitive nature of space systems a teleoperator machine will be working on, changes in visual conditions in Earth orbit relative to visual conditions on Earth, and so on.

The above notes, together with the objective description of telepresence and supervisory control technologies, motivate a few important conclusions.

First, the high fidelity, human operator referenced, man-machine coupling--hardly worked on in telepresence technology--suggests we revisit anthropomorphic machine technology. The primary reason for the revisit is not a declaration of some intrinsic optimality of anthropomorphic machines, but a recognition of their potentially easy and natural interface to human operators to physically extend the rich human manipulative capabilities, embodied in the dexterity of the human hand, to remote places. One may visualize a backdrivable glove-type device on the operator's hand connected through bilateral control to a controllable mechanical replica of the human hand equipped with some

sensing capabilities. This vision may not seem too strange when capabilities of component technologies needed for the development of this anthropomorphic machine are considered.

Second, the performance of nonrepetitive, singular or unexpected teleoperator tasks in space may benefit from the development of shared manual and automatic computer control techniques whenever application scenarios permit their use. These techniques intend to combine the best attributes of human operators and computer control under restricted conditions.

Third, the operator is facing a very rich environment in the control station in terms of decision, command, control and information processing even with increased telepresence and supervisory control capabilities. Due to the nature and time scale of activities in telemanipulation, the operator's mental status and readiness can be compared to an airplane pilot's functional situation during take-off or landing. Thus, proliferation of control and information hardware in the control station does not serve the best interest of the human operator. The more computer technology is employed at the control and information interface in the control station in a clever way, the better off is the human operator to make control decisions efficiently.

Fourth, the R&D effort for advancing telepresence and supervisory control technologies should be accompanied by systematic work on developing a human factors data base and models for understanding and utilizing the results of these evolving technologies. It is apparent from the nature of these evolving technologies that the limits or limitations rest not so much with the technologies themselves but with the human capabilities to absorb and use these technologies.

Fifth, final evaluation and validation of telepresence and supervisory control technologies for space station naturally require experiments and manifests in space whenever human perception, decision, control and other activities are influenced by space conditions. Simulations are useful research and development tools, and they can pave the way towards performance evaluation and validation. But a comprehensive simulation of true space conditions on Earth for developing a human factors data base and models in telepresence and supervisory control technologies does not seem feasible.

Professor Stark make a strong case for NASA-University research in this arena. The benefits of NASA-University connections in human factors research in the field of telepresence and supervisory control can indeed be manifested through past and present examples. Particularly appealing are cases when graduate students carry out the experimental part of their thesis research at NASA-supported, unique laboratories like ARC, JPL, JSC, and so on, or when students spend some working time at NASA laboratories as cooperative students or as academic part-time employees working on topics related to their university studies.

SYNOPSIS OF GENERAL AUDIENCE DISCUSSION

QUESTIONS AND ANSWERS

The first question was focussed on a comment, made by Professor Larry Stark, that wide-field-of-view displays are particularly needed in flight simulators. The question was prefaced with the suggestion that this is a limiting technology for anyone who is interested in robotics applications in space, where (a) the location of the observer is likely to be moving, and (b) the observer needs to be concerned, not only about the orientation of the object being manipulated, but also about his or her own orientation with respect to some larger coordinate range. It was noted that there are some state-of-the-art wide field of view displays that cost millions of dollars and proposed that some kind of research to lower the cost of wide-field-of-view displays might be in order at this point in time.

Professor Stark replied that, in this opinion, wide-field-of-view technology is very important. He provided the following example:

- When people lose their wide field-of-view (e.g., have tunnel vision due to some neurological disease) they find that they can read and their visual acuity is 20-20; they find, however, that it is hard for them to merely walk through a doorway because they are lacking a functional flow-field, the lateral and vertical expansion flow-fields, which are directly connected by primitive neuro-pathways to the vestibular system and are coordinated in the foveolus of the cerebellum as shown in some brilliant studies by Jerry Simpson and other neurophysiologists recently; the lateral and vertical expansion flow-fields give us our orientation.
- On the other hand, when people lose their foveal vision while retaining their flow fields, they are legally blind (with a vision rating of 20/200); they may not be able to read, however, they can still walk through rooms, get into a car, and drive (patients say---"You know, Doctor, I can drive very well, I just can't read the freeway signs, so I don't know when to get off").

Professor Stark concluded that, when people are doing some tasks (manipulating, inspecting) they need foveal vision. In other cases (moving about within an area) they may need a wide-field-of-view. The human visual system is a dual system--we have both--and it should be possible to design something (perhaps using inexpensive TV cameras) to provide wide-field-of-view for gross movement tasks, and high resolution (like reading glasses) for manipulation tasks.

The second question was directed at Professor Sheridan's comment that there is yet no good way of describing (or representing) the process of manipulation. It was suggested that something like the notation system used by choreographers, to represent complex dance motions, might be useful in this context.

Professor Sheridan agreed that "labanotation" (dance scoring) or musical scoring (which is more thoroughly developed), is the kind of thing that might be useful--given a substantial amount of additional development. One problem discussed in relation to the use of this type of notation, was the fact that, for a given instrument, the range of manipulations (speed or fingering) is fairly constrained.

In teleoperations and robotics manipulations, the notation system would have to be able to cope with continuous geometry, hyperspace, and time. In this type of manipulation, considerations include: multiple degrees of freedom (six degrees of freedom for any object, plus maybe the six derivatives, plus the six accelerations--and that is just the beginning) and multiple objects/components in motion (when three or four things are moving in relation to one another you immediately get into a twelve or twenty-four dimensional space and problems of dealing with trajectory in state-space to describe a manipulation). It is a very big order to develop a notational scheme which is both sufficiently complex, and sufficiently comprehensible, to be useful.

Professor Newell noted that the problem of telepresence (generating a feeling, on the part of a remote operator, of "being there" at the work site) is an interesting example of a situation where researchers are working with only a seat-of-the-pants notion of the underlying concepts. He suggested an immense need for a theory and a plausible model of presence--a theory of what happens to humans (and why) when they "project" themselves to a remote work site.

Professor Sheridan suggested caution in the use of terms like "project oneself". He noted that it might be possible to project oneself through drugs, or some other method, which would not be particularly helpful in terms of performance. In addition, he suggested that "being in control of" a remote operation might not require a feeling of "being there"--that telepresence by itself is not the goal--it is really performance that makes the difference.

These caveats notwithstanding, Professor Sheridan agreed that the development of a cognitive theory of presence would be a highly desirable goal. He suggested that "pieces of it are lying around" (e.g., the work of Murray and others in image rotation, etc.).

Professor Stark suggested that "teleprojection" is a very natural phenomenon. He noted, for example, that when an athlete swings a baseball bat, that he or she as an operator/tool user is able to "project" kinesthetically and visually to the end of the bat. He

pointed out that people automatically develop models for activities that they do on a regular basis (e.g., picking up a pen, using tweezers), and suggested that persons operating remote equipment (e.g., a robotic arm 200 miles away) would develop the same sorts of models--as long as there is some sort of causal relationship between their behaviors and the behavior of the remote system.

CONCLUDING REMARKS

In conclusion, one should note that telepresence and supervisory control are not mutually exclusive. Telepresence is needed in supervisory control. The supervisory control language, for example, represents only one abstract operator output interface to the remote system. The perceptive element in supervisory control, that is, the information input to the operator from the remote system, should be in the form of telepresence "frames" in order to help the operator to determine the necessary abstract commands.

We should also note that telepresence has both qualitative and quantitative aspects. The qualitative aspects of telepresence are useful for stabilizing a control situation. The quantitative aspects of telepresence are not well understood (as indicated by control experiments). For instance, when I am working in a force field, and I have active force feedback to my hand, then I am stable--but I have a poor quantitative perception of the acting forces. However, if I show the values of the acting forces on a display simultaneously with the active force feedback to my hand, then I am stable and reasonably good quantitatively. This type of cross modal reference should also be considered in creating telepresence capabilities.

SESSION VI

SOCIAL FACTORS IN PRODUCTIVITY AND PERFORMANCE

SOCIAL STRESS, COMPUTER-MEDIATED COMMUNICATION SYSTEMS,
AND HUMAN PRODUCTIVITY IN SPACE STATIONS: A RESEARCH AGENDA

Karen S. Cook

INTRODUCTION

The sheer complexity of the space station program is enough to boggle the mind of any academic trained in a single discipline. Certainly, space station design requires the ultimate in interdisciplinary teamwork and integration of basic and applied programs of research. In this sense, the project demands knowledge and insights not easily produced in an isolated discipline, be it engineering, aeronautics or sociology. It is a challenging task and one that should call forth the best efforts of those touched by the allure of extending the boundaries of human knowledge.

For a sociologist there are a myriad of research problems which come to mind in even a cursory glance into the window of the future as envisioned by those currently planning the space station program. Clearly, a wide range of processes and factors must be taken into account when considering the more social aspects of this enterprise. These include technological constraints, environmental pressures, physiological limits, psychological processes (including cognitive capacities and motivational factors), and the many interfaces between "man" and machine required by the intense interdependencies of human and technological forces in space. Such intense interdependencies in this extreme are much less often observed on earth (with the possible exception of certain medical contexts in which life is tenuously maintained by sheer technological support).

Given this reality, one cannot extrapolate easily from what is known about society as we experience it on earth and "life aloft." It has even been said that humans may become a very different species while in space. Similarly, social systems which emerge to support and maintain life in this context may deviate along many dimensions from those social structures and processes that are a part of our daily existence and often so "routine" that they are taken for granted. Nothing must be considered as "routine" in a novel environment. It must be said at the outset that what we transport from earth in the way of social, psychological and organizational adaptive mechanisms (e.g. norms, rules, shared expectations, roles, etc.) may prove much less functional than we envisioned given a completely altered social and technological environment. Because we have virtually no scientific evidence

concerning the parameters of life after eighty-four days in space (that is, there is no U.S. experience to rely on), one is forced to engage in speculation and extrapolation despite the potential pitfalls.

My reading of the documents we have been supplied with concerning the space station program in the 1990s and beyond and my very limited exposure to NASA through a two-day symposium, lead me to several tentative conclusions regarding the most critical social contingencies (besides the issue of conflict addressed by Michener) confronting NASA as it plans for the extended duration existence of groups of individuals in space with limited opportunity for replacement or exit. These critical contingencies include the social and psychological management of stress (regardless of the nature of the stressors) and determination of the most efficient and socially productive mechanisms for handling interpersonal communications (e.g. within the crew, between crews of different modules, and between the crew and the "ground," including family members and friends). The successful management of both stress and interpersonal communications is critical to individual and group-level performance, productivity and ultimately, "mission success." While there are many other issues which could be investigated profitably from a sociological perspective, time and space limit the scope of this first foray into life as currently envisioned on space stations.

STRESS, INDIVIDUAL PERFORMANCE AND GROUP PRODUCTIVITY

Stress has been identified as a contributing factor in the etiology of certain acute and chronic illnesses (e.g. ulcers, high blood pressure, heart attacks, nervous disturbances, etc.). It has been demonstrated to have consequences not only for the health status of individuals, but also for individual performance, decision-making and productivity. With respect to space-related research Foushee (1986) states that an important goal is "to understand and minimize the effects of acute and long-duration stresses on group functioning." Although there is enormous literature on the effects of stress on individuals, researchers have been slow to address the impact of stress on groups. Furthermore, the bulk of the existing research examines the physiologic and psychological consequences of stress. There is much less work on the antecedents of stress, in particular the stresses created by social factors (Pearlin, 1982). Another limitation to existing research is the tendency for investigators especially in experimental work to focus on single, isolated stressors. This work is extremely important, but it does not inform us about the interactive and/or cumulative effects of multiple stressors.

Defining and Measuring Stress

The most commonly cited definition of stress is Hans Selye's, "the nonspecific, that is, common, result of any demand upon the body, be the effect mental or somatic." In the tradition of research initiated

by Selye (e.g. 1936, 1956, 1974) this "result" or reaction of the body to stress is referred to as the "general adaptation syndrome" (GAS) or "biologic stress syndrome." It consists of an alarm reaction biologically detectable in such organs as the adrenal glands, thymus, lymph nodes and stomach, followed by the stage of resistance accompanied also by marked physiologic responses, then the stage of exhaustion at which point Selye argues the acquired second-stage adaptation is lost.

Other researchers emphasize the significance of the "cognitive appraisal" of stressors (see Breznitz and Goldberger, 1982, etc.), noting the importance of the "subjective, phenomenological experience of stress" which lies between the stressor and its effects. Some definitions of stress include reference to cognitive appraisal, others, like Selye's, do not. Currently, there is no agreed upon definition of the term and existing differences reflect major unresolved theoretical debates in the field. Though they disagree on the significance of cognitive appraisal, researchers do agree on the common goal of understanding adaptations to stress or the nature of coping mechanisms. Much of the current research focuses upon specifying the nature of these mechanisms. Before discussing adaptations however, let us examine the problems associated with the measurement of stress.

Various approaches have been adopted to the problem of measuring stress; none of them completely satisfactory. One of the most common approaches to measurement, popular over the past two decades because it can be applied outside experimental settings, is the "life-events" scale (e.g. Holmes and Rahe, 1967) or the modified life-events scale (Dohrenwend and Dohrenwend, 1974a, 1974b). Life-events typically mean "objective events that disrupt or threaten to disrupt the individual's usual activities" (see Dohrenwend and Dohrenwend, 1974b:133, 1984). Events listed on such scales include both health-related (onset of chronic illness, major illness or accident, etc.) and non health-related events such as divorce, separation, increase in family income, retirement, death of a spouse, pregnancy or remarriage, etc. (see Thoits, 1981, for a cogent critique of the life-events approach).

The main debate in this research tradition has been over whether or not only undesirable events contribute to stress or whether events that require change either desirable or undesirable produce stress. The latter has been referred to as the "total change" approach to measuring stress, the former, the "undesirability" approach (Thoits, 1981). Thoits (1981) identifies several studies suggesting that only the undesirable changes significantly affect stress levels, although she goes on to critique these studies as well as many of the total change studies for failing to include independent indicators of their independent and dependent variables. Her findings also suggest that "when health-related events are controlled, other undesirable events have small and nonsignificant effects upon psychophysiological distress" (as measured by reports of psychosomatic symptoms using the MacMillan Health Opinion Survey Index). The main conclusion she draws relevant to current research is that "previously well-established correlations between undesirable events and distress may have been inflated due to the operational confounding of health-related items on

the independent and dependent variable scales." A major contribution of new research on stress would be to refine existing measures of stress and to develop more sensitive and reliable measurement techniques.

Laboratory research employs quite different methodologies than survey research, however, as Holroyd and Lazarus (1982:26) point out, "lab paradigms in biological science have tended to isolate stress responses from the psychological and social context." Though measurement problems are reduced in this way, little knowledge is gained concerning the interplay of physiological, psychological and social mechanisms. Holroyd and Lazarus (1982:30) call for "field research that examines stress in the psychosocial context" and more descriptive work on the sources of stress "that operate in naturalistic settings." The space station environment is a "natural" laboratory for this type of research.

Multiple Stressors in Space

The reality of space station existence includes the potential for continual and intermittent exposure to multiple stressors. In this regard it is not at all clear that much of the existing research, except that done in analogous environments, can be extrapolated to apply to the space station. Both the number and the magnitude of stressors in the space environment is likely to be at the high end of existing scales, and quite possibly off the scale. Only research in rare, high stress situations will contain insights of direct relevance.

Sources of potential stress in space stations include sensory deprivation, environmental factors like noise level, crowding, spatial arrangements, and invasion of privacy, as well as isolation, confinement, and the possibility of life-threatening dangers or crisis situations. Nickerson in his chapter for this volume includes in the category of potential stressors: weightlessness, unfamiliar motion, motion restriction, sensory and perceptual restriction as well as sleep interference and acute medical problems. Work-related factors like variety and intensity of assigned tasks, and workload, etc. may also be stressors in the space station environment. Cooper (1983) indicates that in many work environments work or job overload is a major stressor. There is some indication that workload intensity and time pressure were factors that contributed to the problems experienced by crew members aboard the Skylab 4 Mission. According to Holroyd and Lazarus (1982:24), "the individual who is constantly challenged by even relatively innocuous occupational and social demands and who is, as a result, repeatedly mobilized for struggle may be particularly vulnerable to certain disorders (Glass, 1977)." Given the duration of planned space station missions, the cumulative physiological, psychological and social impact of intermittent and continual exposure to multiple stressors must be investigated.

Another significant factor in space stations related to multiple stressors is the recognition that the stressors will be produced by quite different types of events and forces. Stressors may be produced both by the astronaut's home environment, to the extent that s/he has

information about significant events occurring on earth (e.g. in the lives of his/her close relatives and friends, etc.), and by life aloft. Within the space capsule, factors contributing to stress are environmentally induced resulting in both physiological and/or psychological distress as well as socially induced, created by factors associated with the interpersonal environment, especially the intense interdependence of the crew members. Since both physiological and psychological factors have been given more consideration in the existing literature, I will emphasize the social forces likely to induce stress.

Identifying Socially Produced Stressors

Outside of the life-events tradition and research focusing upon occupational stress (e.g. Cooper and Payne, 1978), there have been few investigations of stress produced by interpersonal factors in small group settings (Levine and Scotch, 1970). Potential causes of stress in settings requiring intense interdependence among group members include basic personality conflicts, incompatibilities in interpersonal orientation and style, an inefficient or inequitable division of labor, a lack of perceived legitimacy concerning the allocation of leadership responsibilities or authority, the inequitable allocation of individual or collective rewards, lack of a clear definition of role or task responsibilities, uncertainty regarding the timing, coordination or sequencing of related tasks especially when synchronization is a critical factor, and the arbitrary or inappropriate exercise of authority or influence (i.e. violating role prescriptions or norms concerning the use of private time). Many of these factors have been demonstrated to have significant impacts upon group functioning in non-stressful situations and may or may not be exacerbated in situations of high stress. Research on mountain-climbing teams indicates that under periods of high stress many of these problems become extremely salient and in some cases result in aborted attempts to reach the summit. Interpersonal conflicts appear to be a major problem for many expeditions especially when the goal of reaching the summit is highly valued by all and where there is a great deal of uncertainty about achieving the goal. Connors (1985:147) also notes that in simulation research, "members of isolated and confined groups who were incompatible showed increased stress, withdrawal, and territorial behavior."

Many of these potential stressors have not been examined in the context of group functioning primarily because the predominant model in this area of inquiry has been one of individual functioning. I will comment more upon the limitations of such a perspective in a subsequent section of the paper.

Monitoring Stress

Related to the problem of measuring stress and identifying the antecedents of stress is the problem of monitoring stress. Unobtrusive mechanisms for monitoring stress at both the individual (physiologic and psychological) level and the group level need development, given the potential deleterious consequences of high levels of stress for individual and group functioning. An important byproduct of such monitoring is that it will give us some insight into the interactive and cumulative impact of various stressors. Furthermore, it will enable us to address issues still under debate regarding the extent to which the effects are linear, curvilinear, or approximate a step-function (or threshold function). It may also be the case that the effects of certain stressors are compensatory given that not all the effects are potentially negative. The positive impact of stress has been given little attention in the literature.

Personal Characteristics, Crew Composition and Stress

As several authors have suggested, the "right stuff" may be the "wrong stuff" when it comes to the selection of crew members who will not only have the necessary technical and professional skills, but will also have the psychological and social competencies required for the creation of effective interpersonal relations and relatively smooth group functioning on space station "missions." According to Biersner and Hogan (1984:495), veterans believe that "social compatibility is as important as technical skills for overall Antarctic adjustment" to isolation. Social competence will become even more critical as a basis for selection and training in the future as NASA envisions shorter training periods for some astronauts (e.g. teacher and congressmembers in space programs). The potential for commercial joint ventures with NASA not only increases crew heterogeneity, but also means that some space station members in the U.S. module will in all likelihood not have the benefit of intense NASA training (and selection).

Intriguing research by Helmreich and his colleagues (e.g. Helmreich et al., 1980) on this basic topic suggests that at least one characteristic typically associated with the "right stuff" constellation of traits, interpersonal competitiveness, may be dysfunctional for producing smooth group functioning depending upon the mix of personnel and their traits in any particular crew. As Connors (1985:155) notes, Helmreich et al. (1980) "hypothesize that the combined interests of task accomplishment and social compatibility will be best served if crew members show a strong work and mastery orientation, but relatively little competitiveness."

Group Decision-Making Under Stress

Research of particular interest to NASA is the research on the relationship between stress and decision-making which indicates that

the experience of stress generally interferes with psychological processes related to effective decision-making. Janis (1982), for example, reports the following reactions associated with stress during decision-making:

- (1) narrowing of attention span and range of perceived alternatives,
- (2) reduction in problem-solving capabilities,
- (3) oversight of long-term consequences,
- (4) inefficiency in information search,
- (5) premature closure, and
- (6) with intense fear, there is also temporary loss of perceptual acuity and perceptual-motor coordination (Duffy, 1962).

Evidence further suggests that accelerating time pressure increases the probability of these reactions, although clearly more research is needed on the temporal aspects of stress reactions as well as situation specific/individual difference interaction effects. (Individuals in certain situations are likely to respond differently both to stress and to the demands of the decision-making task.)

Janis (1982) also specifies five basic patterns of decision-making under stress. The first four patterns in the list represent "defective" patterns of response, the fifth is the term Janis uses for the most adaptive response pattern. Observed patterns of response under stress include:

- (1) unconflicted inertia
- (2) unconflicted change
- (3) defensive avoidance
- (4) hypervigilance, and
- (5) vigilance.

Of the four defective response patterns, hypervigilance is found to be the dominant reaction under conditions of high stress or near-panic. As Janis (1982:77) notes, "Excessive alertness to all signs of potential threat results in diffusion of attention...one of the main sources of cognitive inefficiency whenever someone becomes hypervigilant, and it probably accounts for some of the failures to meet the criteria for effective decision-making." Results also suggest that other problems emerge in high stress situations. "Along with cognitive constriction there is a marked tendency toward stereotyped thinking in terms of oversimplified categories and reliance on oversimplified decision rules" (Janis, 1982:78). Two conditions appear to enhance the probability of hypervigilance occurring in stressful situations: unconflicted inertia (or the failure to react to early warnings) and defensive avoidance (e.g. procrastination). Additional factors identified by Janis as associated with the antecedents of hypervigilance are the lack of contact with family members or other support persons, lack of perceived control and lack of preparatory information about potential stressful events.

The prevention of "defective" patterns of response in threatening situations has focused in recent years upon several strategies including "benign preexposure to the threatening situation, stress inoculation via preparatory communications" and various types of relaxation techniques designed to mitigate physiologic reactions (Janis, 1982:82; see also, Janis et al., 1982). Research on these techniques and the extent to which they are successful under specific circumstances continues. Extrapolation to situations likely to be encountered in space stations must be done carefully. Some techniques may be effective for single stressors, but less effective in the face of multiple stressors. Again, further research is needed. Certainly, however, this research gives us some clues as to problems associated with decision-making in highly stressful contexts.

A Comment on the Limits of Medical and Psychological Models of Stress: The underlying framework a researcher adopts to the analysis of a problem often circumscribes both the nature of scientific inquiry as well as conceptions of potential solutions. Thus it is not surprising that medical research on stress tends to examine primarily physiologic response patterns and the impact of drugs on the functioning of the individual undergoing stress. Psychologists similarly focus on cognitive and emotional factors, examining individual differences associated with cognitive appraisals of stress and reactions. The solutions they consider include biofeedback, stress "inoculation", and various types of individual training and therapeutic techniques. All of this research is necessary since the problem entails both physiologic and psychological dimensions. What is missing, however, from much of the current work is the investigation of the system properties of stress and examination of solutions to the problems created by multiple stressors at the group or collective level (also sometimes called the system level). Inquiry of this type would examine the interpersonal dynamics related to stress responses and adaptive strategies rather than treating the problem purely from an intraindividual perspective. Adoption of an interpersonal or system level perspective would lead to quite different conceptions of adaptive mechanisms. In Connors (1985:146) words, "Given that future missions will require increased levels of cooperative functioning, selection and training procedures must not only yield effective individuals, they must yield effective groups."

The dominant characteristic of space station missions in the near future involving 6-8 crew members marooned in space for approximately ninety day intervals of the high degree of interdependence among the group members (and possibly between groups in different modules at some point). Stressors which significantly impact any single group member will, of necessity, influence group functioning--even if it simply entails the reassignment of duties or tasks for brief periods of time or temporary isolation of a group member. In addition, group members may be impacted similarly by stressors and thus collective solutions should be explored. Strategies might be developed, treating the group as a social system (as Michener does) of interdependent parts and group members might be trained in specific response patterns through a division of labor. For example, roles could be assigned such that each

attends to a specific problem associated with inefficient decision-making under high stress. One crew member might be assigned the task of vigilance with respect to only alternatives, another to long-term consequences, etc. and coordination might be achieved either by an assigned group leader or some sort of computerized decision-aide.

Computer-aided systems could be developed which help to meliorate common deficiencies observed in cognitive processing during peak periods of stress. Coping strategies of this type are more like Janis' suggestion that an appointed "devil's advocate" be used to mitigate the negative consequences of "groupthink." They have the possible advantage that "failure" is not localized in a single individual (typically, the "leader") who must assume full responsibility for group decisions in "crisis" or intensely stressful situations. Furthermore, a clear division of labor also reduces the workload on any single individual under stress. The work on distributed decision-making by Fischhoff and others may well provide models for this type of coping mechanism. Relevant work on computer-aided decision-making should also be explored.

Mediators of Stress and Adaptation

In the words of Holroyd and Lazarus (1982:25), "It has been increasingly acknowledged that health outcomes are a product of effective coping rather than simply a consequence of the presence or absence of stress." Identifying factors that result in effective coping is an important research agenda item, however current investigations focus more on drug therapy, biofeedback and "cognitive-behavioral" interventions to modify responses to stress and facilitate coping. The social and organizational management of stress, as noted above, has not been examined. Psychological approaches take us one step beyond physiologically focused management strategies, but even they have not been evaluated extensively.

Coping mechanism and adaptation responses form one axis of current research, the second axis is extensive work on factors that "mediate" the stress response. Such factors include individual differences which relate not only to susceptibility, but also to cognitive appraisal and effective coping. Variables incorporated into these investigations are ethnicity, age, gender, occupation, income, level of education, marital status, health status and access to social support (i.e. personal resources and network supplied resources), among others. Access to social support, for example, has been demonstrated to mitigate some of the effects of stressful events (e.g. Caplan and Killilea, 1976). Much of this work is useful for general medical and scientific purposes, but caution must be exercised when attempting to generalize these findings to astronauts and the space station environment. The range of variation on some of these variables is quite restricted in the astronaut population, although increasing heterogeneity must be assumed along many of these dimensions (i.e. gender, age and ethnicity) in the future.

Research linking gender to stress, for example, indicates in a variety of studies that women are more susceptible to stress (e.g. Kessler and McRae, 1981); given certain levels of stress they report higher levels of distress as reflected typically in symptomatology (primarily self-reports). Research discussed by Kessler and McLeod (1984) documents that women tend to be more affected by undesirable life events than men even though they do not report significantly more such events. Kessler and McLeod (1984) present findings that indicate that women are more vulnerable to "network" events, events that happen to significant others in their networks, than men, and it is this difference that accounts at least in part for previously observed sex differences in responses to stress. Thus, they argue that women are not "pervasively more vulnerable than men to stress," but vulnerable specifically to stress linked to the important people in their lives as a result of their "greater emotional involvement in the lives of those around them." Belle (1983) refers to this fact as the "stress of caring".

There are many unanswered questions concerning the link between gender and stress. The extent to which female astronauts are more vulnerable to stress than male astronauts is an open question. Few of the existing studies include in their samples women in such high stress occupations and it may well be that women with high capacities for coping with stress self-select into these occupations (e.g. as is likely the case for women mountain climbers). It should also be noted that many of these studies reporting sex-related stress differences are based on sample data obtained in the 1950s, 1960s and early 1970s; little evidence exists based on more recent data including samples of women in more varied occupational contexts and roles.

Impact on Productivity: Individual and Group-Level Effects

The link between stress and productivity has been demonstrated to be somewhat complex. Mandler (1982:94) argues that "the problem of stress is twofold; both the initial autonomic signals and the conditions that generate these signals require some conscious capacity...and therefore interfere with the performance of targeted tasks." What is not clear is specifically how and under what conditions performance is impaired. In fact, as Mandler (1982:96) indicates, like noise, stress reduces "attentional capacity and narrows it to central tasks," thus if the target task is central, "then autonomous arousal may improve performance." This depends upon both the centrality of the target task and specific characteristics of the task, or task sequence which requires performance. Early research on this topic seemed to suggest that there is a curvilinear relationship between arousal and performance such that performance is enhanced by moderate levels of arousal, but impaired significantly at both very low and very high levels of arousal. The generality of this effect is still under debate. Mandler (1982:95) concludes that "understanding the relation between efficiency and stress requires an analysis of specific stressors, an approach to arousal that assigns it definable

properties..., and knowledge about the requirements of the task." Research by Baddeley (1972) and others indicates that stress associated with dangerous environments "affects performance through its influence on the subject's breadth of attention...but we still do not know what mechanisms mediate the effect of arousal" on attention span or even what is entailed in the adaptation to fear.

Evidence suggests that problem-solving abilities are affected by stress in much the way Janis indicates that decision-making is impacted. In particular, "if much of problem-solving involves the manipulation in consciousness of alternatives, choices, probable and possible outcomes and consequences, and alternative goals," then stress interferences with efficient problem-solving. Few alternatives are actually considered and the thought process is guided more by habituation and stereotyping than by the conscious weighing of alternative strategies. What is needed, he argues, is "fine-grained" analyses of these processes. "Preoccupation with the unstressed mind has restricted experimental work on these problems" (Mandler 1982:101). A related shortcoming is the failure to consider the social context of problem-solving behavior. The bulk of the research deals with individual tasks, not collective or highly interdependent tasks.

A Research Agenda: System-Level Responses to Stress

In the previous era when highly trained male pilots were selected as astronauts on the basis of physical stamina, high tolerance for stress, psychological stability and technical competence for space missions involving relatively short-term exposure to multiple stressors in dangerous environments, less attention was paid to research on stress. In fact, Mandler (1967) noted in his early studies of highly trained astronauts a lack of anticipated stress responses; these men had been "trained to have available response sequences, plans and problem-solving strategies for all imaginable emergencies" thus emergencies were transformed into "routine situations" and therefore not experienced as stressful. At this stage in the space program endurance was the primary focus of both selection and training. Even space capsule design decisions were not frequently made in order to minimize environmentally induced stress or to increase "habitability" (Clearwater, 1985).

The future holds forth a different scenario. First, astronaut selection procedures have changed to include non-white males and scientific personnel as well as pilots. There is greater diversity among potential astronauts in occupational training, gender, age, ethnicity, and personality traits. Given this heterogeneity and the increased complexity and duration of space station missions, emphasis must now be placed (as Helmreich, 1983; Foushee, 1984; and other social scientists have argued) on the selection and training of highly compatible crews especially as group size increases to eight or more in relatively small modules. In addition, habitability has recently become an integrated aspect of the space station design process. Alterations in selection processes to maximize crew compatibility and

design decisions to improve habitability are essential ingredients. But as Danford et al. (1983) note in their chapter, "Humane Space Stations", social and organizational factors must also be considered. Two specific factors have been isolated for consideration in this paper: (1) the social management of stress and development of interpersonal coping mechanisms, and (2) the socially efficient and productive management of interpersonal communications.

Development of a specific research proposal is beyond the scope of this chapter, however, research recommendations to NASA would include examination of existing data on crew interactions especially under stressful conditions to isolate effective interpersonal strategies for coping with stress and to identify particular interaction sequences which either exacerbated or mitigated stress responses. These data should be examined in relation to individual performance, group performance and interpersonal climate. Variation in interpersonal strategies by type and duration of stressors should also be investigated. In the early stages of the mission stressors may be predominantly physiological (e.g. resulting from space adaptation sickness or initial bodily responses to micro-gravity, etc.), however, as duration of the mission progresses psychological and social stressors may become more pronounced (i.e. intensification of the sense of isolation and confinement, monotony of the physical environment, and increased sensitivity to interpersonal incompatibilities, etc.). The most promising data sources for such analyses are likely to be tapes from the Skylab Missions given that they provide some insight into flights of analogous duration to planned space station missions.

Another useful focus of research would be investigation of group decision-making under stress. Existing data could be mined for insights into the impact of stress on predicted cognitive and behavioral responses (e.g. the possible occurrence of hypervigilance), in decision-making situations of varying types. A separate research strategy would be to simulate group decision-making under stressful circumstances. One model for this type of research is the work by Foushee and his colleagues (e.g. Foushee and Helmreich, forthcoming) on crew performance under stress in aircraft flight simulations. Again, the aim would be to identify successful interpersonal strategies for coping with critical deficiencies resulting from stress. One potential byproduct of this research would be identification of the characteristics of computer decision-aides which would facilitate group functioning under conditions of high task interdependence and high stress. Information-seeking behavior, for example, could be isolated and analyzed for inefficiencies which could be meliorated by the proper use of expert systems or computerized search procedures. As Nickerson concludes in his chapter, "Stress is likely to be an important factor in the Space Station...Exactly how these factors, especially in combination, will affect performance and productivity is not known."

MEDIATED COMMUNICATION AND CREW PRODUCTIVITY

In a 1983 NASA-ASEE final report entitled "Autonomy and the Human Element," the authors state that the "general transmission and processing of information lies at the heart of almost every aspect of space station activity." Over the past decade information processing and communications have engaged more and more of the design capabilities of NASA both in terms of hardware and software development efforts. Rapid advances in technology make this aspect of space station design especially volatile and vulnerable to obsolescence. While cost understandably plays a major role in design decisions, other factors affecting crew morale and productivity must be taken into account. Communication modality is also a critical factor in the coordination of activities aboard the space station. An intensive examination of the benefits and disadvantages of different modes of communication for within crew interactions, as well as for interactions between crew members and "ground" or mission control personnel, and for crew interactions with significant others is required. Morale, efficiency, productivity, the potential for conflict, the exercise of authority and control, and, ultimately, mission "success" are all affected by communication modality, access to information, and the structure of the communication channels.

Computer-Mediated Communication as Primary Modality

As Connors et al. (1985) put it "mediated communication systems must be developed to meet the needs of the crew throughout an extended mission." Such communication systems are not only vital to the ongoing mission of the space station, but may also be critical in maintaining social contact between station crew and ground personnel and thus contribute to the reduction in stress created by the sense of isolation and confinement. Maintenance of good communication links between the ground staff (e.g. "mission control" and other base operations) and the members of the space station crew are essential to the smooth functioning of the space station. Currently, one of the primary modalities for communication processes is computer-mediated interaction (Simes and Sirsky, 1985). This section of the chapter includes a brief review of some of the relevant research on the impact of computer mediation on group interaction and decision-making. Other modalities for mediated communication are mentioned; however, cost factors necessitate heavy reliance upon computer-mediation.

Studies of the Effects of Computer-mediated Interaction

Siegel et al. (1986), in experimental studies contrasting the effects of face-to-face versus computer-mediated communication, find that with certain types of group problem-solving tasks there are marked differences between communication modes. Three types of communication modes were examined in the studies they report: face-to-face,

simultaneous computer-mediated discussion and computer mail. While the results are not definitive, they suggest that communication mode affected the speed required to reach a group decision, the equality of participation rates of group members, communication rates, nature of the interpersonal communications, as well as the degree to which the group's decision deviated from individual's initial choices. The results indicate that there are certain advantages and disadvantages to computer-mediated communication systems which are relevant to plans for space station communications, although more systematic research is required.

Specific results of interest include the fact that computer-mediated simultaneous communication appeared to retard group decision-making when contrasted with face-to-face communication. In addition, this mode of communication fostered greater equality in participation rates among group members, increased the number of inflammatory or "uninhibited" remarks and resulted in group decisions which deviated to a greater extent from initial individual choices than was the case when communication was face-to-face. (It should be noted that the subjects who participated in these groups had no prior association with one another.) Findings from the condition in which subjects communicated by computer mail were similar in most respects to the computer-mediated "conference" mode.

Implications for Space Station Communication Systems

The implications of the findings of Siegel et al. (1986) for decision-making and group problem solving aboard the space station are intriguing, though speculative. First, it would appear that complex problem-solving tasks, especially when time to solution is critical, are facilitated most by face-to-face communications even though this modality increases inequality in participation rates. The role of video connections in approximating face-to-face communication where physical copresence is not possible (as between crew members and family members or between crew members and mission control) has yet to be fully investigated. Limited research suggests that video contact (which is available to both parties) reduces perceived "social distance," but the role of perceived social distance in complex group problem-solving is not clear. Research varying both the complexity of the task and the degree to which face-to-face contact is mediated is needed.

Results concerning the effects of communication mode on participation rates also requires further investigation in relation to task complexity and degree of task interdependence. The greater equality in participation rates fostered by computer-mediation may be functional for tasks requiring creative solutions (or during the "brainstorming" phase of group problem-solving) when maximization of input is essential. Computer-mediation may also mitigate to some extent the effects of status differences on participation rates (a well-established finding in the small groups literature, see Bales work on the link between status and power and prestige orders and

participation rates). Though the finding concerning the impact of computer-mediation on participation rates and its implications for the reduced effect of status differences is speculative, it certainly requires further investigation. Studies in which clear status differences exist among group members need to be conducted in computerized settings. Computer-mediation may facilitate the "upward flow" of negative information or information that challenges the positions of those in high status roles in the group. This effect is important since under time pressure or in stressful situations information is often critical to effective decision-making. Experimental research and simulation studies could be conducted on this topic. It appears that computer-mediation may mitigate the inhibiting effects of face-to-face communication when "subordinates" have access to critical information and may need to challenge authority or the group's dominant decision strategy (see Foushee, 1982, 1984, etc.). Connors (1985:174), for example, cites research indicating that "correctable pilot errors have gone uncorrected because of unquestioning attitudes, a lack of assertiveness, or deficient communication skills." Another intriguing result cited by Connors (1985:197) was obtained by Champness (1971) indicating that people are more likely to change their established positions on issues and reach a compromise with other group members when communication is not mediated. This may have important implications for both the process and outcome of group decision-making aboard the space station.

Alterations in the norms surrounding communication content under varying communication modes also need further investigation. The normative restraints of face-to-face interaction on communication content are lessened in the more anonymous condition in which computers mediate interaction. As Siegel et al. (1986) note, computer-mediated communications included more inflammatory remarks. If this finding is observed in groups which have a history of interaction, then computer-mediation could foster interpersonal conflict and mechanisms to meliorate this possibility would have to be developed. A related concern is the protection of privacy in communications meant for family and friends, especially communications high in socio-emotional content. All forms of mediated communication raise issues of access as well as privacy which need careful examination in relation to individual morale, group cohesiveness and other dimensions related to the interpersonal environment within the space station. Connors (1985:197) cites studies indicating that mediated communication contains "reduced socio-emotional content," and thus is less effective for certain types of tasks such as negotiation or getting acquainted in contrast to tasks which require "the giving and receiving of information, asking questions, or exchanging opinions." Research on space station communications and the impact of computer-mediation on the performance of different types of tasks, as well as the nature of the interpersonal dynamics within the crew and between crew and ground is needed.

Individual and Group Level Impacts of Computer-Mediated Communication Networks

Kerr and Hiltz (1982) discuss the potential impacts of computer-mediated communications on individuals and groups focusing on cognitive, affective and behavioral dimensions. They are concerned with broad effects at the organizational and societal levels, many of which go far beyond the scope and size of the space station. Some of the hypothesized effects have been verified in research discussed above by Kiesler and her colleagues (Siegel et al., 1986), but many of the topics raised by Kerr and Hiltz have not been subjected to systematic research. Furthermore, much of the evidence they cite is anecdotal, based on the experiences of those in positions to evaluate existing computer-mediated communication networks. Though computer-mediated communication networks of various sizes have existed for at least a decade, research examining the effects on specific variables related to group functioning and organizational effectiveness is fairly recent.

With respect to individual performance, Kerr and Hiltz (1982) discuss such issues as information overload, new skill requirements and improvements, expansion of learning opportunities, etc. as potential cognitive impacts of computer-mediated communication systems. Hypothesized affective impacts include: enhancement of the candor of opinions, potential "addiction" and heavy usage, increased network size and possible sources of social support (from kin, friends, and professionals), the ability to maintain friendships despite lack of geographical proximity, etc. Negative potential consequences discussed include increased isolation from non-mediated communication relations, new sources of stress related to changes in existing patterns of work and communication as well as alterations in social networks, and the frustration created by the lack of immediate feedback, etc. Hiltz (1979), however, notes that in some cases, "The desire to have truly synchronous conferences seems to almost totally disappear as experience is gained on the system."

Of the individual-level behavioral impacts discussed, several are of primary interest. First, it is clear that computer-mediated networking increases connectedness among individuals thus expanding the scope and range of social relationships. According to Kerr and Hiltz (1982:114), computer-mediated communication systems lead to "increased collegial contacts, an increase in the number of contacts that can be maintained, and create the opportunity for regular connections with many people." Expansion of the actual or perceived social network through computer-mediated communication systems may help mitigate the sense of isolation experienced by space station inhabitants. Results indicate that a major strength of such systems is the ability to "keep in touch with others" (see Kerr and Hiltz, 1982:114, Vallee et al., 1978:111-115). In addition, such systems seem to alter the centrality of individuals by allowing those geographically (or for other reasons) on the periphery to regain a sense of centrality through increased communication contact.

Group-level impacts are especially relevant to space station design. Kiesler's work addresses some of the issues related to group

decision-making contrasting computer-mediated communication with face-to-face groups. However, Kerr and Hiltz (1982:121-122) identify a wide range of other group and organizational level impacts, some of which correspond to Kiesler's concerns. The group-level hypothesized cognitive impacts include: (1) the creation of "on-line" groups or "communities of interest", (2) improved group decisions, and (3) an increase in "knowledge-based authority," etc. With respect to group decisions, the findings cited are mixed. On the positive side results suggest that the capabilities of data base searches, increased access to information and access to decision-aides enhance group problem-solving and decision-making. As Turoff and Hiltz (1980:123) indicate "the computer can aid in gathering subjective estimates within a group" and facilitate the resolution of disagreements.

While Kerr and Hiltz (1982) indicate some empirical support for "at least the same quality of solution" when comparing computer-mediated to face-to-face groups (Turoff, 1980; Hiltz et al. 1981); Kiesler et al. (1984) and Siegel et al. (1986) report a decrement at least with respect to time to solution for the computer-mediated groups. Others, Kerr and Hiltz (1982) note, (see Johansen et al., 1979) argue that more conflict may result from the increased access to alternative views and that a "false sense of group consensus" may arise (Kerr and Hiltz, 1982:125).

On group problem-solving Kerr and Hiltz (1982:124) cite the work of Lipinski et al. (1980:158-159) which suggests that when considering, the "task-focused communications required by groups involved in joint problem solving, computer-based communication systems are appropriate in the structuring, evaluating, and documenting phases of problem solving, since time delays are acceptable, written responses are appropriate, and face-to-face contact is not essential." However, they go on to state that the "implementing, searching, and conceptualizing stages of problem solving are less amenable to this technology." More research is needed concerning the phases of problem solving and the effects of computer mediation.

The following list includes some of the hypothesized behavioral impacts on groups identified by Kerr and Hiltz (1982:132-133). Many have not been sufficiently investigated to provide definitive evidence. (Adapted from Kerr and Hiltz:)

1. Computerized communication increases cross-group communication.
2. It increases lateral network linkages among organizations.
3. It increases lateral network linkages within organizations.
4. Computerized communication may change social structures from pyramid or hierarchical to network-shaped.
5. It changes the centrality of members within groups.
6. It increases the possible span of control.
7. It can increase the effective limits on the size of working groups.
8. It increases the density of social networks, increasing connectedness.
9. It increases opportunities for decentralized communication.
10. Computerized communication may increase informal communications.
11. It changes who talks to whom.

12. Groups take longer to reach agreement and consensus is less likely.
13. Computerized communication sometimes makes it difficult to focus discussions.
14. Regularity of individual participation is sometimes difficult to enforce.
15. There is greater equality of participation than in conventional media.

Communication Network Structure, Centrality and Power

Prior research on communication networks in the social sciences provided evidence that the specific configuration or structure of the network affected the efficiency of problem solving groups. But more recent research tends to indicate that these results may not be valid for mediated communication systems. Subjects in various four-person network structures, given telephone contact capabilities, were able to come to consensus on group decision problems without much variation in degree of consensus or time to achieve consensus across structures (see Friedkin and Cook, 1987). Results from the computer-mediated version of this experiment are not yet complete.

Centrality has been linked to power in various studies of communication and in networks in which resources other than information are exchanged (see Freeman, 1979; Cook et al., 1983). In computer-mediated communication networks centrality is linked to access to information and control over the flow of information. To the extent that computer-mediation alters these parameters decentralization of power may occur. Research is needed which examines the relationship between the structure of the communication network and control over information channels. Certainly as Kerr and Hiltz (1982:150) indicate "opportunities for decentralized communication are increased" in computer-mediated networks, "because it is easier to keep all those concerned with issues informed and up to date." Thus the efficient flow of information is enhanced. But efficient decision-making in groups in which communication is computer-mediated may require structured access to information rather than open access during the final stages of decision-making. Levels of access to information rather than the availability of communication channels becomes the critical determinant of positional centrality and thus power in this circumstance. Further research on these topics is needed.

Communication Networks, Authority and Control

Kerr and Hiltz (1982:125), among others, predict that computerized communication increases the "appreciation of knowledge-based rather than hierarchical authority." If this result is general, it will be important to study the conditions under which conflict can arise between knowledge-based and hierarchical authority structures. Efficient group functioning and problem solving is likely to be enhanced when there is minimal conflict between these sources of

authority. Furthermore, hierarchical authority and command systems must be designed in such a way that information flow is not tightly hierarchically structured.

As noted above, in particular, in systems involving highly trained professionals the upward flow of critical information must not be circumvented by bureaucratic procedures or restricted communication channels. Maximization of group productivity and problem solving efficiency is likely to occur under conditions of open access to communication channels rather than strict hierarchical access under conditions of complex tasks, high uncertainty and a highly professionalized staff. Specific research on optimum alternative authority structures under varying communication network structures and task conditions is required.

With respect to authority and control in systems using computer-mediated communication networks, two additional impacts cited by Kerr and Hiltz (1982:150-151) are relevant. They argue (p. 150) that "greater delegation of authority is possible with the capacity for accountability and reviewing decisions in a timely and orderly manner." Second, they argue (p. 151) that it "increases the possible span of control" and "allows more centralized control over geographically dispersed units." Computerized decision-aides have the potential to alter both accountability and review procedures, but the specific extent and optimum role of these systems in human decision-making has yet to be determined.

Extension of the span of control and the degree of centralized control over units dispersed in space may become more important considerations during the post-IOC phase of the space station program. Some of these issues as they relate to the potential for intergroup conflict have been addressed by Michener in his chapter in this volume.

The Impact of Computer-Mediated Interaction: Research Needs

Research on the impacts of computer-mediated interaction on individual and group-level functioning is relatively new. There are major limitations to existing knowledge in this area; results are more often based on anecdotal reports than systematic research or are derived from very limited observations over limited time spans in situations in which there is little control over the relevant variables. A major research program is required. Of particular importance in the design of space station configurations and communication systems is research on the links between information access channels and the exercise of authority and control. Various factors make the space station unique: the high degree of professionalization of the staff, the complexity of the tasks involved, the high degree of interdependence and uncertainty surrounding many of the tasks to be accomplished, the enormous information requirements, the difficulty and complexity of continual on-line monitoring, the spatial separation of the ground-based crew and command personnel from the space crew, and the potential existence of multiple authority structures.

Existing research is focused on earth based communication networks primarily among colleagues or remote members of interest groups where the exercise of authority is rarely an issue. Information exchange is frequently the primary or sole goal of the interaction. Thus extrapolation from the results of studies on these networks must be treated as highly speculative. New research must be designed around the specific problems and parameters facing crews in space. Simulations could be designed which would mirror some of the most critical circumstances and used to evaluate alternative network structures, systems of controlled versus open access to information, given different types and levels of complex tasks. Problem solving efficiency and group productivity would be a primary focus of the research, although other issues such as increased social communication between crew members and ground personnel would also need to be addressed in terms of the impact on mission success, broadly defined. Priority should be placed on the development and evaluation of on-line data collection systems for post-IOC space station missions and other long-duration, "manned" missions concerning the multiple impacts of computer-mediated communication systems.

Summary Statement Concerning Research Needs

The 1986 Challenger disaster was as much a failure in organizational decision-making as a technical failure in the right rocket booster on the shuttle. This fact attests to the tendency in organizational contexts for scientists and managers to focus attention primarily on the technological aspects of systems rather than the social aspects of system design. Historically, in the social sciences, as well as the physical sciences, productivity has been viewed fundamentally as a problem of technical system or organizational design and innovation. Those who design and evaluate complex systems which require human participation, however, must eventually recognize the significant role of psychological and social factors in productivity. Human factors are now incorporated in NASA's research program, but this is a recent and fairly small beginning given the time frame within which research commitments are necessarily made.

My recommendations assume that technical and social systems can not be designed in isolation of one another and that interdisciplinary research which crosses the invisible boundary between the physical and social sciences is required. Designing space stations which are maximally habitable and which optimize human comfort, satisfaction and productivity and minimize the sense of isolation and the stresses associated with risk and uncertainty, as well as the potential for intra-group and inter-group conflict is as critical a goal as the flawless design of structures which will provide the technical support for "life aloft".

Research on many critical aspects of social system design is simply not available. In part this is because the technologies under consideration are new (e.g. computer-mediated networks to facilitate interpersonal communication are relatively recent); but also in part,

this state of the art is a function of national priorities and budgetary constraints. Hopefully, this situation will change. The quality of life in space in the twenty-first century will hinge upon decisions we make during this decade as to what research is necessary to maximize not only productivity, the bottom line for many, but also less tangible qualities such as habitability, sociability and liveability. The space station is, after all, a place to be inhabited, a mini-society which at some not too distant time in the future must begin to cope with not only the technological requirements of its environment, but also the psychological and social needs of its inhabitants and the social constraints and requirements of an emerging society. Recruitment, selection, training, sustenance and replacement of persons will be as critical as the maintenance and replacement of parts.

The following is an abbreviated list of research needs (see Table 1) which I have suggested in the text of this report related to social factors involved in space station design during the post-IOC phase. The emphasis in this report has been placed on issues related to stress, its causes and consequences, and the impacts of computer-mediated communication systems (since that is currently the primary modality envisioned.) I have only scratched the surface.

In conclusion, it is important to note that as with many of the research programs of NASA and University-based scientists, the benefits to be derived from the proposed research extend far beyond the limited purposes of future space station missions. Improved methods for coping with multiple stressors in hostile environments and a better understanding of the social and psychological effects of computer-mediated communication systems have great potential applicability in a wide range of human social contexts. The payoffs for society as we know it on earth are potentially even greater than the payoffs for life as we envision it on space stations in the next century.

TABLE 1 Selected Research Needs: Social Factors and Productivity on Long-Duration Space Station Missions

Social Stress, Human Productivity and Group Functioning:

- (1) Develop more comprehensive and precise measures of stress levels for situations involving multiple stressors.
- (2) Research and develop stress monitoring systems, on-line data collection procedures, and more unobtrusive measures of stress.
- (3) Conduct research on personal characteristics (e.g. personality dimensions, gender, etc.) and specific responses to stress and adaptations to stress on long-duration space station missions.
- (4) Examine group composition factors which maximize efficient group functioning under multiple stressors.
- (5) Research the specific impacts over time of multiple stressors on individual and group decision-making processes. Assess the effectiveness of different coping strategies and decision aides under varying levels of stress and combinations of stressors.
- (6) Expand research on the causes of stress to include as well as psychological and physiological factors social factors such as group size, group composition, division of labor, workload, perceptions of equity in the assignments of tasks and responsibilities, styles of leadership, type and degree of contact with significant others, etc. on long-duration missions.
- (7) Begin to develop process models which relate stress to individual performance and group-level functioning and specify the conditions under which the impairment of individual performance seriously compromises group functioning.

Computer-Mediated Communication Systems, Human Productivity and Group Functioning:

- (1) Extend existing research on the social impacts of computer-mediated communication systems on individual decision-making and group problem solving.
 - (2) Investigate the effects of computer-mediation in relation to the phases of group problem solving, complexity of the tasks and variations in the levels of environmental stress and uncertainty.
 - (3) Conduct research on computer-mediated communication systems and the distribution of power and authority. Investigate in particular the potential for conflict between knowledge-based and hierarchical authority structures and the link between centrality and the exercise of power and influence.
 - (4) Investigate the potential consequences of computer-mediated communication between crew members and significant others on earth attending to issues of privacy, social support and the effects on responses to isolation, confinement and other stressors on space station missions.
 - (5) In the future, research the differential impacts on individual performance and group functioning of various types of mediated communication systems (including audio and video channels).
 - (6) Examine factors related to communication modality and access to communication channels which inhibit the upward flow of critical information (especially negative information) and mechanisms which circumvent this problem.
 - (7) Consider the effects of computer-mediated communication on the relations between crew members and ground personnel and between crews of different modules with respect to the potential for intergroup conflict and develop mechanisms to mitigate conflict.
-

REFERENCES

- Raddeley, A. D.
1972 Selective attention and performance in dangerous environments. British Journal of Psychology 63:537-546.
- Belle, D.
1982 The stress of caring: women as providers of social support. Pp. 496-505 in L. Goldberger and S. Breznitz, eds., Handbook of Stress: Theoretical and Clinical Aspects. New York: The Free Press.
- Piersner, R. J., and Hogan, R.
1984 Personality correlates of adjustment in isolated work groups. Journal of Research in Personality 18:491-496.
- Breznitz, S., and Goldberger, L.
1982 Stress research at a crossroads. Pp. 3-6 in L. Goldberger and S. Breznitz, eds., Handbook of Stress: Theoretical and Clinical Aspects. New York: The Free Press.
- Broadbent, D. E.
1971 Decision and Stress. New York: Academic.
- Caplan, G., and Killilea M., eds.
1976 Support Systems and Mutual Help: Multidisciplinary Explorations. New York: Grune and Stratton.
- Champness, B. G.
1971 Bargaining at Bell Laboratories. Paper E/71270/CH, Communication Studies Group, London, England.
- Clearwater, Y.
1985 A human place in outer space. Psychology Today. July, 34-43.
- Cook, K. S., Emerson, R. M., Gillmore, M. R., and Yamagishi, T.
1983 The distribution of power in exchange networks: theory and experimental results. American Journal of Sociology 89:275-305.
- Cooper, C. L., ed.
1983 Stress Research: Issues for the Eighties. Chicester: John Wiley & Sons.

- Cooper, C. L., and Payne, R.
1978 Stress at Work. Chichester: John Wiley & Sons.
- Connors, M. M., Harrison, A. A., and Atkins, F. R., eds.
1985 Living Aloft: Human Requirements for Extended Spaceflight. Washington D.C., NASA.
- Danford, G. S., Ericson, J. E., Freeman, H. J., Hays, D. G., Lewis Jr., W. C., Schoonhoven, C. K., Sims, H. P., Walsh, P. J., Wellens, A. R., and Wise J. A.
1983 Humane space stations. Pp. 63-95 in Autonomy and the Human Element in Space. Final Report of the NASA/ASEE Summer Faculty Workshop, NASA.
- Dohrenwend, B. S. and Dohrenwend, B. P.
1974a Overview and prospects for research on stressful life events. Pp. 313-331 in B. S. Dohrenwend and B. P. Dohrenwend, eds., Stressful Life Events: Their Nature and Effects. New York: Wiley.
- Dohrenwend, B. S. and Dohrenwend, B. P., eds.
1974b Stressful Life Events: Their Nature and Effects. New York: Wiley.
- 1984 Stressful Life Events and Their Contexts. New Jersey: Rutgers University Press.
- Duffy, E.
1962 Activation and Behavior. New York: McGraw-Hill.
- Foushee, C.
1982 The role of communications, socio-psychological and personality factors in the maintenance of crew coordination. Aviation, Space, and Environmental Medicine 53:1062-1066.
- 1984 Dyads and triads at 35,000 feet: factors affecting group process and aircrew performance. American Psychologist 39:885-893
- Foushee, C. and Helmreich, R. L.
1987 Group interaction and flightcrew performance. To appear in E. L. Wiener and D. C. Nagel, eds., Human Factors in Modern Aviation. New York: Academic Press (forthcoming).
- Freeman, L. C.
1979 Centrality in social networks: conceptual clarification. Social Networks 1:215-239.

- Friedkin, N., and Cook, K. S.
1987 Communication Network Structures and Social Influence Processes in Small Problem-Solving Groups. To be presented at the annual meetings of the American Sociological Association, Chicago, August, 1987.
- Glass, D. C.
1977 Behavior Patterns, Stress, and Coronary Disease. Hillsdale: Erlbaum
- Helmreich, R. L.
1983 Applying psychology in outer space: unfulfilled promises revisited. Pp. 445-450 in American Psychologist. April.
- Helmreich, R. L., Wilhelm, J. A., and Runge, T. E.
1980 Psychological considerations in future space missions. Pp. 1-23 in T. S. Cheston and D. L. Winters, eds., Human Factors of Outer Space Production. Boulder, Colorado: Westview Press.
- Hiltz, S. R., Johnson, R., and Turoff, M.
1981 The quality of group decision-making in face-to-face versus computerized conferences. Presented at the annual meetings of the American Sociological Association, Toronto, Canada.
- Holmes, T. H., and Rahe, R. H.
1967 The social readjustment rating scale. Journal of Psychosomatic Research 11:213-218.
- Holroyd, K. A. and Lazarus, R. S.
1982 Stress, Coping, and Somatic Adaptation. Pp. 21-35 in L. Goldberger and S. Breznitz, eds., Handbook of Stress: Theoretical and Clinical Aspects. New York: The Free Press.
- Janis, I. L.
1972 Victims of Groupthink. New York: Harcourt, Brace.
1982 Decisionmaking under stress. Pp. 69-87 in L. Goldberger and S. Breznitz, eds., Handbook of Stress: Theoretical and Clinical Aspects. New York: The Free Press.
- Janis, I. L., Defares, P. B., and Grossman, P.
1982 Hypervigilant reactions to threat. In H. Selye, ed., Selye's Guide to Stress Research. Vol. 3. New York: Van Nostrand.
- Johansen, R., Vallee J., and Spangler, K.
1979 Electronic Meetings: Technological Alternatives and Social Choices. Addison-Wesley: Reading, Mass.

- Kerr, E. B., and Hiltz, S. R.
1982 Pp. 114-151 in Computer-Mediated Communication Systems: Status and Evaluation. New York: Academic Press.
- Kessler, R. C., and McLeod, J. D.
1984 Sex differences in vulnerability to undesirable life events. American Sociological Review 49: 620-631.
- Kessler, R. C., and McRae Jr., J. A.
1981 Trends in the relationship between sex and psychological distress: 1957-1976. American Sociological Review 46:443-452.
- Kiesler, S., Siegel J., and McGuire, T. W.
1984 Social psychological aspects of computer-mediated communication. American Psychologist 39:1123-1134.
- Levine, S., and Scotch N. A., eds.
1970 Social Stress. Chicago: Aldine Publishing Co.
- Lipinski, H., Spang S., and Tydeman, J.
1980 Supporting task-focused communication. Pp. 158-160 in A. R. Benenfeld and E. J. Kazlauskas, eds., Communication Information: Proceedings of the 43rd ASIS Annual Meeting. Knowledge Industry Publication: White Plains, New York.
- Mandler, G.
1967 Invited commentary. In M. H. Appley and R. Trumbull, eds., Psychological Stress: Issues in Research. New York: Appleton Century Crofts,
1982 Stress and thought processes. Pp. 88-104 in L. Goldberger and S. Breznitz, eds., Handbook of Stress: Theoretical and Clinical Aspects. New York: The Free Press.
- Pearlin, L.
1982 The social context of stress. Pp. 367-379 in L. Goldberger and S. Breznitz, eds., Handbook of Stress: Theoretical and Clinical Aspects. New York: The Free Press.
- Selye, H.
1936 A syndrome produced by diverse nocuous agents. Nature 138:32.
1956 The Stress of Life. New York: McGraw-Hill.
1974 Stress Without Distress. Philadelphia: Lippincott.
- Siegel, J., Dubrovsky, V., Kiesler, S., and McGuire, T. W.
1986 Group processes in computer-mediated communication. Organizational Behavior and Human Performance 37:157-187.

- Simes, D. K., and Sirsky, P. A.
1985 Human factors: an exploration of the psychology of human-computer dialogues. In H. R. Hartson, ed., Advances in Human-Computer Interaction 1:49-103. Norwood, N.J.: Ablex Publishing Corp.
- Thoits, P. A.
1981 Undesirable life events and psychophysiological distress: a problem of operational confounding. American Sociological Review 46:97-109.
- Turoff, M.
1980 Management Issues in Human Communication Via Computer. Presented at the Stanford Conference on Office Automation. Stanford, California.
- Turoff, M., and Hiltz, S. R.
1980 Structuring Communications for the Office of the Future. Presented at the National Computer Office Automation Conference, Atlanta, Georgia.
- Vallee, J., Johansen, R., Lipinski, H., Spangler, K., and Wilson, T.
1978 Group Communication Through Computers 4:Rep. R-40. Institute for the Future, Menlo Park, California.

CONTROL, CONFLICT, AND CRISIS-MANAGEMENT
IN THE SPACE STATION'S SOCIAL SYSTEM (YEAR 2000)

H. Andrew Michener

THE SPACE STATION'S CREW AS A SOCIAL SYSTEM

This paper discusses the organization of the crew on board NASA's Space Station in the year 2000. In line with the work of Sells and Gunderson (1972), the perspective adopted here is that the crew of the Space Station is not just as a collection of people but a functioning social system. Crew members are viewed not just as individuals, but as interdependent parts in a larger structure.

Under current plans, the Space Station will evolve from its earliest form (called the Initial Operating Configuration, or IOC), which will exist approximately in year 1993, to a complex form (herein called the Second-Stage Operating Configuration, or SSOC) in year 2000. In the IOC (1993), the crew of the Space Station will be small (i.e., 6-8 persons). As the Space Station evolves over time, the crew will grow in size, and by SSOC (2000) it will have grown to 20-30 persons. It is possible, of course, to view the crew as a system even when there are only 6-8 people on board, as in IOC. However, it becomes increasingly useful to view human relations in system terms when there are more persons on board, as in SSOC.

NASA has traditionally placed great emphasis on careful selection and intensive training of its crews, and the outstanding performance of NASA crews aloft attests to the success of this approach. Selection and training will continue to play an important part in IOC and SSOC Space Station operations. Nevertheless, as the Space Station evolves from IOC to SSOC, NASA will find that it must rely less on selection and more on intentional design of the on-board social system to achieve adequate performance by the crew. This will occur because the growth in size will render the crew increasingly less a collection of individuals and increasingly more a system with emergent properties. During the evolution from IOC (1993) to SSOC (2000), important changes will occur in the social system on board. Not only will the system increase in size, but it will become differentiated into distinct subgroups and more complex in structure. These evolutionary changes will not only affect the Space Station's performance, but also determine the types of problems and failures that occur within the social system on board.

The main purpose of this paper is to assist NASA in developing a research agenda for the SSOC social system. It must be recognized, however, that neither the IOC nor the SSOC social systems exist today. This means that research is problematic, because there is no way that one can directly observe these systems or take measurements on them at this point in time. Since the IOC and SSOC social systems are yet to be developed, the essential question is not research, but planning and design--what shape and structure will these systems have and how will they function. Research becomes useful primarily as an adjunct to the design problem; that is, it becomes useful to the extent that it improves some social system designs or eliminates some candidate designs from further consideration.

To develop research ideas for SSOC, this paper first describes ways in which the SSOC social system will differ from the IOC social system. Next, it discusses three operating problems that may be more troublesome in SSOC than in IOC. These are (a) supervising and controlling the diversity of payload activities, (b) handling the relationship between differentiated subgroups of crew members, with its potential for intergroup conflict, and (c) responding to environmentally-induced crises. Finally, some avenues of research are suggested regarding these operating problems.

COMPARING THE IOC AND THE SSOC SOCIAL SYSTEMS

Social Systems in Space

Social systems in space operate under parameters different from social systems on Earth. These parameters, which apply to both the IOC and SSOC social systems, include:

(a) Perilous Environment. In contrast to most Earth-based social systems, the crew on board the Space Station (and on any space vehicle) will face a perilous environment (microgravity, no oxygen) and require complex life-support. Crew members will face significant hazards and risks to life.

(b) Relative Isolation. The social system on the Space Station will be isolated from other social systems and (in many respects) self contained. It will be in contact with Earth only via telecommunications, and hence it potentially has some degree of independence from Mission Control on Earth.

(c) Long Duration. The social systems on board the Space Station, while transitory compared with those on Earth, will remain in space for increasingly long durations. Space Station crew members will fly missions that endure 90 days. (The Space Station itself may continue usefully in orbit for 20-30 years.) From the standpoint of individual crew members, long-duration missions may entail stress, psychological depression, and diminished performance (Bluth, 1980, 1981; Cunningham, 1977; Oberq, 1981).

Exogenously Mandated Changes in SSOC

The environment faced by the Space Station's crew in SSOC will be just as perilous as that in IOC. However, the Space Station's social system will not remain constant. NASA has already mandated certain changes in the social system that are to occur between IOC (1993) and SSOC (2000). These changes include:

Change in Crew Size

One difference between IOC and SSOC is the size of the crew on board the Space Station. In IOC, the crew will be small (5-8 persons). In SSOC, the crew size will be larger, perhaps 20-30 or even more. This increase in size will be made possible by the physical expansion of the Station. Most of the added crew members in SSOC will be Payload Specialists, not Astronaut Pilots.

Change in Crew Composition

Several important changes in the composition of the crew will occur between IOC and SSOC. First, the Japanese and European Space Agencies will attach modules to the Space Station in SSOC and place their own Astronauts aboard. Whereas the IOC crew will consist of USA-NASA personnel, the SSOC crew will include substantial numbers of several distinct nationality subgroups: USA, Japan, Europe.

A second change to occur concerns the skill mix of the crew. In IOC, most crew members will be Astronaut Pilots. In SSOC, there will obviously still be some Astronaut Pilots on board, but the crew will include many more Payload Specialists than in IOC. Some calculations illustrate this point. If it takes two Astronaut Pilots to fly the Space Station at one time, then a total of four persons will be needed to fly the Space Station around the clock (assuming that flight operations are never left unattended and that Astronauts work 12 hours at a stretch.) The implication is that, in IOC, at least half the crew members will spend their time flying the Space Station, not conducting payload operations. The situation in SSOC will be more favorable, because the number of persons needed to fly the Space Station will presumably remain about the same (despite the larger physical size of the Station); most of the additional persons on board in SSOC will be Payload Specialists, who can devote their time to scientific or manufacturing productivity.

A third change, less well defined at this point, concerns the gender mix of the crew in SSOC. NASA has shown that it intends to put women in space, although missions to date have been male dominated. Presumably the crew of the Space Station will include some women. With the move from IOC to SSOC, and the accompanying increase in crew size, there may be opportunity to move the ratio of females/males on board closer to 1.00, should NASA opt to do this.

Change in Mission Statement and Goals

In IOC, the primary mission goals will be, first, to fly the Space Station and, second, to construct large space structures, i.e., expand the physical structure of the Space Station using components flown up via the Shuttle (Danford et al., 1983). These goals will doubtless apply to SSOC as well.

In SSOC, however, the increased number of Payload Specialists on board will permit other goals to be pursued. These goals may include manufacturing and materials processing under conditions of micro-gravity, and tending and repairing communications satellites. Other objectives may include conducting scientific experiments, carrying out remote sensing and meteorological monitoring, and engaging in flight support (assembly, maintenance, checkout, launch, recovery) for manned or unmanned LEO transfer missions (Danford et al., 1983). Overall, the goals pursued by the crew members in SSOC will be more complex and diverse than those in IOC. Expressed more formally, the SSOC social system will be attempting to optimize what may be construed as a highly complex multi-objective function (Keeney and Raiffa, 1976).

Change in Onboard AI and Computerization

Current plans for the Space Station call for an increasing use of artificial intelligence (AI) and expert systems over time. The extent to which AI can be used in IOC and SSOC depends both on the capabilities of the Space Station's computers and on the software itself.

In past missions, the computers on board NASA's space vehicles have not been powerful, due in part to limitations imposed by physical size and weight. The situation will be somewhat better in IOC. Plans indicate that IOC will include some AI systems, although these will be small-to-moderate in size. NASA will, of course, use mainframe computers on Earth, and these may supplement the AI routines of the Space Station's smaller onboard computers. Some AI systems on board will probably serve as consulting devices for the diagnosis of hardware failures. Other onboard computerization may involve scheduling of crew activities and maintenance of databases (e.g., materials inventory).

By SSOC, the computers on board the Space Station will be faster and capable of running large AI programs. Moreover, the software will have evolved with experience on board the Space Station, and will become more wide-ranging in its capacities. Thus, AI and expert systems will be more prominent in SSOC than in IOC, and SSOC will be more automated.

From the standpoint of the social system, the evolution of computerization is relevant because AI will become integral to onboard decision-making. By SSOC, the AI software will be able not only to diagnose hardware failures, but also to schedule human activities and perhaps even to resolve conflicts among humans regarding priority of objectives.

Induced Structural Changes in SSOC

The exogenous changes mandated by NASA for SSOC, as listed above, will bring about many changes in the internal organization of the SSOC social system. Of course, because neither the IOC nor the SSOC social systems exist today, one cannot draw firm conclusions about their structural properties or performance under specified conditions. Nevertheless, by considering the proposed systems in light of research findings on Earth-based social systems and earlier space-flight social systems, some plausible conjectures can be made regarding their structure and performance. It seems fairly clear that the SSOC social system, as contrasted with the IOC system, will be more complex, more differentiated into subgroups, and more decentralized with respect to decision-making.

Complexity

The SSOC social system will be far more complex than that in IOC. The SSOC social system will include more members (20-30, rather than 6-8), and the complexity of the system will increase nonlinearly with crew size. The primary source of this increased complexity is not just larger crew size per se, but rather the fact that the system's growth will occur via differentiation (elaborated subgoals and subgroups) and not via segmentation (Sutherland, 1975; Casti, 1979).

This increase in complexity is reflected, for instance, in the number of communication channels in IOC as contrasted with that in SSOC. With 8 crew members in IOC, there are 28 channels (assuming that each channel is 2-way and that a crew member does not require a channel to communicate with himself); with 30 crew members in SSOC, there are 435 channels. Thus, a 4-fold increase in crew size produces a 16-fold increase in channels. Of course, it may be the case in SSOC that every crew member will not have a need to communicate with all others, but the increase in structural complexity is nevertheless clear.

Increased complexity will show up not merely in structural measures but also in functional ones. For instance, complexity might become apparent in slower response to emergencies or crises. Today there is no way to measure the response-time performance of the SSOC social system. Could one do this, however, the SSOC social system might emerge as slower (and less predictable) than the IOC system when responding to such emergencies as fire on board or a collision with space debris. To mobilize 20-30 persons scattered in several modules (in SSOC) will probably take more time than to mobilize 6-8 in one module (in IOC).

Differentiation

The social system in SSOC will be far more differentiated--that is, composed of subgroups with distinct identities--than the social system in IOC. The bases for this differentiation will be national origin and

task specialization; there may also be some subgroup differentiation based on gender.

Under current plans, NASA will add physical modules to the Space Station between IOC and SSOC, causing an evolutionary expansion in size. NASA itself will supply some modules, but others will come from foreign space agencies (Japan, Europe). Hence, the crew on board the SSOC Space Station will consist of persons from all three space agencies (USA, Japan, Europe), possibly in proportion to the financial contribution by various participating nations. This means the SSOC crew will consist of subgroups that (a) have different national origin (US, Japan, Europe--Britain, France, Germany, Italy), (b) have different native languages, (c) have different skin color and racial characteristics, making group membership readily visible, (d) have different moral and religious belief systems, and (e) perhaps have different goals and long-term agendas. This SSOC crew profile differs sharply from the far more homogeneous IOC crew profile; in IOC the crew will be single nationality (primarily or entirely USA), single language, consonant beliefs, unitary goals, single command structure on the ground (NASA), etc.

Crew members from the three space agencies will, at least to some degree, constitute distinct subgroups on board the SSOC Space Station. Of course, the use of a single language (English) on board will help to lessen subgroup differentiation. Nevertheless, an extrapolation from research on Earth-based social systems suggests that differences in the factors noted above (nationality, skin color, native language, belief systems), reinforced by NASA's plan to house together persons from a given country in their own module, will cause the subgroups to have at least a moderate degree of in-group identification and well-defined boundaries (Tajfel and Turner, 1986; Wilder, 1986; Brewer and Campbell, 1976).

Another basis for subgroup differentiation present in SSOC (but not in IOC) is task specialization. As noted above, both IOC and SSOC will have Astronaut Pilots, but SSOC will have many additional Payload Specialists. The SSOC crew, for instance, may include such diverse specialists as a university astrophysicist, a commercial materials engineer, and a national security intelligence analyst.

The Astronaut Pilots in SSOC may view themselves as a distinct subgroup within the larger social system. They will have similar backgrounds, perform similar activities, and work for the same employer on the ground (NASA). Whether the Payload Specialists in SSOC will view themselves as a second distinct subgroup is less clear, because they may differ significantly among themselves. That is, the Specialists will come from a range of educational backgrounds, work for different employers on Earth, pursue a diversity of objectives while on board the Space Station, and perhaps even operate under orders to keep their activities secret from others on board. If some Payload Specialists work interdependently on tasks or report to similar commands on Earth, there is the possibility that they will form identifiably distinct, functioning subgroups on the SSOC Space Station.

Decentralization

The social system in SSOC will be more decentralized than that in IOC. In other words, decision-making will be distributed more widely across persons in SSOC than in IOC. Supervisory control over various functions will shift away from a central command and reside instead with a diversity of specialists.

Pressures toward decentralization of decision-making and control in SSOC will come from several sources. First, as the Space Station evolves from IOC to SSOC, there will be a change in the Station's mission. Payload operations will become more prevalent and important. As a result, the activities on board will become more differentiated and specialized (e.g., materials processing under microgravity, satellite servicing, and conduct of experiments). Most of these new activities will be expertise-based, and they will be controlled by the only persons on board who know how to do them (i.e., Payload Specialists, not Astronaut Pilots). The expansion of expertise on board in SSOC will coincide with decentralization of decision-making.

Many Payload Specialists in SSOC will be employees not of NASA, but of other organizations on Earth. One implication is that the Payload Specialists presumably will report to different supervisors on the ground. This fact will conduce toward more decentralization of decision-making on board the Space Station.

OPERATING PROBLEMS FACING THE SSOC SOCIAL SYSTEM

As detailed above, the social system on board the Space Station will undergo significant structural changes from IOC to SSOC. The system will experience a change in mission statement, grow in complexity, differentiate into subgroups, and decentralize in decision-making. These shifts will produce operating problems for the SSOC social system that were not present in IOC. Although one can doubtless identify many such problems, three are of special interest here. These are singled out not only because they pose special threats to overall mission performance, but also because they potentially can be mitigated (if not eliminated) through design and research efforts. The three are:

(a) The SSOC system will face problems with supervisory-control functions that were not present in IOC. The burden of coordination will be greater, because the SSOC system will include distinct national subgroups as well as more task-specialization subgroups than IOC. Coordination of activities will be more problematic in SSOC, in part because decision-making will be more decentralized.

To some degree, the problems with supervisory-control functions can be addressed through design efforts prior to SSOC. The broad research/design issue for NASA is what type of supervisory-control structure will best serve the SSOC system, in the sense of providing greatest efficiency and highest probability of mission success.

(b) The SSOC system will pose risks of intergroup conflict that were not present in IOC. The presence on board of several distinct

subgroups, with potentially opposing interests and objectives, increases the prospect of conflict.

The broad research/design question for NASA is what safeguards to build into the system to reduce the probability of overt conflict occurring. A related question is what can be done to assure that any conflicts that do arise are resolved constructively.

(c) The SSOC system may have more difficulty than the IOC system in coping with crises (e.g., fire on board, collision with space debris, etc.). The SSOC social system will probably have more resources than the IOC system for coping with many crises. At the same time, the SSOC system—with its greater degree of differentiation and decentralization—may be worse-off organizationally than IOC and have more difficulty mobilizing to deal with crises.

The broad research/design question for NASA is how best to structure the SSOC social system so that it can mobilize adequately to deal with various crises.

The following sections discuss each of these problems in turn. Primary focus is on the nature and genesis of the problems. Attention is also given to design issues—that is, to what research might be done by NASA prior to SSOC to mitigate these problems.

SUPERVISORY-CONTROL AND OPTIMAL PERFORMANCE

The topic of supervisory control by humans on board the Space Station has several dimensions. First, there is the matter of humans' reliance on and control over machines. Under current plans, the Space Station's physical subsystems will include many sensors and control devices to monitor and regulate automatically a variety of outcomes, including life-support, power sources and management, flight control, thermal control, and traffic control. Thus, when interfacing with machines, the crew members on board will enter the Space Station's control process only in a high-level monitoring, troubleshooting, and decision-making capacity (Kurtzman et al., 1983; Von Tiesenhausen, 1982).

A second aspect of supervisory control on the Space Station is the regulation of crew members' activities by other crew members. This topic is of interest here because there will be a shift in the Space Station's onboard supervisory-control structure during the evolution from IOC to SSOC. The following discusses some aspects of this change.

The Supervisory-Control Structure

As used here, the term supervisory-control structure refers to that functional subsystem on board the Space Station which (a) regulates crew activity in the interest of attaining system goals, (b) makes choices among collective behavioral alternatives, and (c) handles dissent, including the treatment of noncompliance by crew members.

In social systems on Earth, supervisory-control structures (often called "authority" systems) typically specify who makes what decisions,

who evaluates whose performance, and who influences (gives orders to) whom. No doubt the supervisory-control structure on the Space Station will entail such specifications, with the added characteristic that some prerogatives will reside with crew members on the Space Station while others will inure with NASA personnel on the ground.

Supervisory-control structures can assume a wide variety of forms. For instance, at one extreme there is the archetypical military command model with hierarchical lines of authority and command. In pyramidal structures of this type, control flows from the top down, while information flows up (Mesarovic et al., 1970). At another extreme there is the equalitarian model with a flat authority structure. In the Space Station context, such a model might consist of equally-ranked Astronauts aloft, not taking orders from a crew member on board, but each reporting to someone on Earth. A third supervisory-control structure--falling between the extremes of hierarchy and equality--is the heterarchy. A heterarchical structure is one that resembles a network, the nodes of which are relatively independent control systems and the arcs of which are the lines of communication passing between the nodes (Sutherland, 1975). On the Space Station, the nodes in such a structure might be individual Task Specialists, or possibly teams of Specialists.

It follows that one important research/design issue is exactly which supervisory-control structure should be deployed on board the Space Station. Since this issue is important both in IOC and in SSOC, it is useful first to look briefly at the IOC situation.

Supervisory-Control Structure in IOC

The main objectives of the Space Station crew during IOC will be to fly the Station and to expand its physical structure (add new habitation modules and platforms). Any of several alternative supervisory-control structures might suffice in IOC to accomplish these objectives, although some structures are probably better than others. The question, then, is which to deploy. NASA might base its choice on such procedures as trial-and-error or extrapolation from previous experience with space flight supervision. Alternatively, systematic research could be used to narrow the choice by eliminating some candidate structures.

More specifically, NASA might conduct simulations on the ground to test various outcomes from different supervisory-control structures. Simulations might be done under conditions that closely replicate those found in space--e.g., high stress, high noise, restricted communication, 90-day duration, tasks similar to those done in space, and so on. Important outcome measures include productivity levels, crew satisfaction, lack of conflict, adequacy of response to emergencies, etc. Multiple replications could be run on each of several alternative supervisory-control structures using standard experimental designs. The results should provide a fair idea of how the alternative supervisory-control structures will perform.

Without the results of such research, it is hard to know what type of structure will eventually be deployed. A plausible conjecture, however, is that the Space Station's IOC supervisory-control structure will, at least to some degree, resemble a standard "military command model" with hierarchical lines of authority and command. There is a general tendency for groups facing perilous environments to organize themselves hierarchically, primarily because it strengthens their capacity to respond to emergencies and crises (Helmreich, 1983; Harrison and Connors, 1984). This pattern occurs not only in space missions, but in submarines, underseas research vessels, North Sea oil rigs, and polar expeditions. Most likely, the IOC system will be no exception.

Thus, the supervisory-control structure on board during IOC will presumably involve a designated "Mission Commander" (or some such title) with authority to issue orders to subordinates. Of course, the 6 or 8 Astronauts on board during IOC are going to be competent, skilled, and resourceful persons. They will have been selected via a rigorous screening process, and there will be little reason to doubt their capacity for decisive action. Nevertheless, their roles will be fairly restrictive and afford little independence, and they will essentially be taking orders from Mission Control on Earth and from their Mission Commander on board the Space Station.

Relations Between the Crew and Mission Control

Both in IOC and in SSOC, one research/design issue deserving consideration by NASA is the exact allocation of control between Mission Control on Earth and the crew on the Space Station. The viewpoint taken here is that the Space Station will not be "autonomous" or independent of Mission Control. Because many monitoring and control functions are better performed on the ground than in space, Mission Control will exert considerable influence over a wide range of crew members' activities and decisions throughout IOC. Crew members, however, will probably retain control over such things as the inventory of items on board the Space Station and the flow of traffic in and around the Space Station.

More problematic is whether crew members will have control over the scheduling of their own day-to-day activities. On one hand, Mission Control needs assurance that crew members are performing adequately and thus may wish to exercise strong supervision over schedules. On the other hand, tasks which are easy to perform on Earth may consume great time and energy under microgravity in space (Sloan, 1979). This may cause Mission Control to expect too much and could lead to overscheduling of daily activities by personnel on the ground.

Excessive regulation of crew schedules by Mission Control can produce role overload on space missions (Helmreich et al., 1979). Even worse, lock-step regulation of the crew's schedule by Mission Control might result in such labor problems as the well-publicized one-day "strike in space" that occurred during the 1973 Skylab mission (Balbaky, 1980; Cooper, 1976). To achieve a workable balance, what the

Space Station needs is an arrangement whereby Mission Control can specify (longer-range) goals to be achieved, while crew members can express and to some degree enforce their preferences regarding local work flow and task-allocation.

One approach to such an arrangement is based on experience in earlier space missions. Both the Russians and Americans have reported some success with task-assignment procedures whereby decisions regarding mission and related tasks are made under the hierarchical model, and decisions regarding off-duty activities and living arrangements are made democratically (Leonov and Lebedev, 1975; Nelson, 1973). Although promising, these results pertain primarily to short-duration missions, and their applicability to longer-duration missions is still an open question subject to further research.

Another approach to the issue of day-to-day task scheduling is to rely heavily on computer software. This approach will be relevant in SSOC, and may also be applicable in IOC. Many large projects of various types are managed on Earth today via project planning software. Task scheduling on the IOC Space Station will probably not be so complex as to require software more elaborate than that available today. In fact, computer software for project management on the Space Station will not only be useful in achieving optimal allocation of tasks to crew members, but may even emerge as a tool for conflict resolution between the Space Station crew and Mission Control.

Supervisory-Control Structure in SSOC

As noted above, the social system in SSOC will be larger, more complex, more differentiated, and more decentralized than that in IOC. In consequence, the supervisory-control structure in SSOC will be more elaborate than that in IOC and probably will assume a fundamentally different form.

Functions of Supervisory-Control in SSOC

The SSOC supervisory-control structure must be geared to handle many of the same functions as the IOC system. These include flying the Space Station, coordinating with Mission Control on Earth, and building (expanding) the Space Station. In addition, it will have to handle other functions, such as processing materials and servicing satellites, as well as serving as a node in a larger communication and transportation network in space.

The SSOC social system will include not only Astronaut Pilots, but also a large number of Payload Specialists (perhaps as many as 20 of them). Regulation of these Specialists may prove a complicated task. Most Payload Specialists will be highly educated professionals knowledgeable in their respective specialties. Many will be accustomed by prior employment to working under supervisory-control structures permitting a high degree of independence and autonomy. On the Space Station, they may be performing activities (such as research) that are

best accomplished under decentralized decision-making, and they will probably be working for someone other than NASA (reporting to non-NASA authority on Earth). One implication of these facts is that a straightforward extrapolation of the hierarchical IOC military command model to SSOC will not suffice.

Morphology of Supervisory-Control in SSOC

It was suggested above that NASA might use experimental research (simulations) to design the initial IOC supervisory-control structure. A similar approach would be applicable to the design of the SSOC control structure. In the absence of such research, however, it is a plausible conjecture that the supervisory-control structure in SSOC not resemble a military hierarchy to the same extent that the IOC structure did (Helmreich, 1983; Danford et al., 1983; Schoonhoven, 1986). Instead, it may more nearly resemble an industrial heterarchy. This is a structure in the form of a network, the nodes of which are relatively independent control systems.

Due to task specialization, decision-making within SSOC will be more decentralized than in IOC. Interaction will be more collaborative, collegial, and advisory. To a significant degree, influence will flow in many directions (not just top-down) and will be based on expertise and control of information as well as on organizational status. Despite all this, however, Payload Specialists in SSOC will not be truly autonomous or independent. They may have more decision-making prerogatives than Specialists in IOC, but their discretion will nevertheless be circumscribed and their performance will doubtless be subject to administrative regulation and review. Much of this administration will originate from (non-NASA) personnel on the Earth, not from other persons on board the Space Station.

On the Space Station itself, many Payload Specialists in SSOC may be organized into small teams (2-, 3-, 4-persons, etc.) working on specific tasks. This team structure will capitalize on the added productivity that results from such processes as social facilitation (Zajonc, 1965; Henschel and Glass, 1968; Marcus, 1978); at the same time, it will permit the Space Station's crew to tackle a diversity of unrelated tasks requiring different competencies (research, materials processing, satellite servicing, construction, etc.) The teams constituting the nodes of the heterarchy will each have decision-making authority regarding work-flow on their own task (doubtless with the consent of supervisors on Earth).

In addition to this structure, the SSOC system will likely include a small administrative staff--e.g., a Mission Commander and several lieutenants who will be responsible for coordinating relations among the diverse projects on board the Space Station. These administrators will have the power to halt or reschedule activities on one project in order to facilitate another. (Again, coordination of this type will require the concurrence of Mission Control on Earth.) Moreover, these administrators will also have the capacity, if an emergency or crisis

arises on board, to halt all task activities and to mobilize the entire crew to cope with the emergency.

In sum, it is suggested that the supervisory-control structure in SSOC will probably differ from that in IOC, and may assume the form of a heterarchy or quasi-heterarchy. This statement, however, can be no more than a conjecture. It has been proposed that NASA might use simulation research on alternative supervisory-control structures as a basis for developing the design of the SSOC system.

Leadership Roles in SSOC

The model of the supervisory-control structure sketched here assumes that, in SSOC as in IOC, there will be an overall Mission Commander on board the Space Station. The exact nature of the Mission Commander's role is an open research/design issue. It seems clear, however, that his role during SSOC will be different from that during IOC, for he will coordinate and oversee rather than give directives, military-style. Although he will have the skills to fly the Space Station, he will not handle the minute-to-minute task of piloting the Station. Nor will he carry out many payload operations per se. Instead, his major role will be to coordinate flight operations and payload operations, as well as coordinate relations among nationality subgroups on board and with Mission Control on Earth.

Beyond the nature of the Mission Commander's role, there is the question of what persons might be candidates for that role. Whether the position of Mission Commander in SSOC will be restricted to NASA Astronauts or open to crew members from Japan and Europe is yet another research/design issue for NASA to address. A similar issue, too, arises with respect to the lieutenants and other officers on board the Space Station.

THE POTENTIAL FOR CONFLICT IN SSOC

Risks of Conflict in SSOC

Conflict in social systems can manifest itself in diverse forms. Typical forms include argumentation, social "friction," interpersonal disliking, attitudes of distrust, passive refusal to cooperate, and so on. Of course, when conflict becomes severe it will emerge in still other forms such as physical violence.

Although the evidence on this point is largely anecdotal, relations among crew members in earlier NASA space flights have been harmonious. There is little evidence of serious conflict or disagreements among crew members themselves. There is, however, some evidence that disagreements have occurred between space crews on one hand and Mission Control on the other (Pogue, 1985; Balbaky, 1980; Cooper, 1976). The source of these conflicts appears to have been task overload or lock-step regulation of crew activities imposed by Mission Control.

Except for the longer flight duration, there is no reason that conflict in the IOC social system will be greater than that during previous NASA space flights. There may again be some disagreements between the crew and Mission control, but probably not much conflict among crew members themselves.

In SSOC, however, the situation is different. There is more potential for interpersonal and intergroup conflict in SSOC than in IOC.

Sources of Conflict in SSOC

The risks of conflict are higher in SSOC than in IOC because the crew will be differentiated into subgroups and decentralized with respect to decision-making. First, SSOC will include many more Payload Specialists than IOC. Each such person will have his or her special goals, which means that the SSOC Space Station will be pursuing more complex (multi-objective) goals and that decision-making will be more decentralized than in IOC. These diverse goals may be (somewhat) incompatible, and coordination will be more problematic.

Just as significantly, the inclusion in SSOC of several nationality groups with distinct identities (USA, Japan, Europe) creates the potential for intergroup conflict. Whether conflict actually erupts among members of different subgroups depends on incompatibilities among the different roles, values, and goals of these persons. The fact that these subgroups will be housed in distinct physical modules will probably heighten cohesiveness within the subgroups and thereby increase the likelihood of friction between subgroups. The added fact that Americans may be in the minority (or, at least, not in the majority) on board the Space Station in SSOC could make the situation even more volatile.

Although it may be only partially relevant to SSOC, research on Earth-based systems shows that integration is problematic in social systems having many subgroups. Systems of this type are more vulnerable to higher levels of conflict, mis-coordination, lack of cooperation, and mistrust than are systems having no subgroups. Not surprisingly, conflict between subgroups is especially likely to occur when they have divergent objectives or interests (Campbell, 1965; Sherif et al., 1961; Diab, 1970). Moreover, when conflict does occur in social systems of this type, it often is more severe (i.e., more rancorous, more divisive, more difficult to resolve) than that occurring in systems having no distinct subgroups. This happens because, in systems with distinct subgroups, conflict is not just disagreement among persons as individuals, but among persons as agents of subgroups.

In sum, NASA has chosen to deploy a heterogeneous, differentiated SSOC social system in which the risks of conflict are higher than would be the case in certain other types of social systems. The risks would be less severe, for instance, had NASA chosen to deploy an SSOC system more like that in IOC--i.e., a system where crew members have a uniform

nationality (USA), single native language, unitary goals, centralized command structure on the ground (NASA), single living module, and so on.

The Importance of Avoiding Conflict

No assertion is made here that conflict, mistrust, and lack of coordination are inevitable in the SSOC social system. It is merely being suggested that conflict is more likely in SSOC than in IOC. Conflict occurring in SSOC will probably be of low-to-moderate intensity (not severe intensity), and will probably appear in such forms as argumentation, friction, and distrust (not physical violence). There will be no need to install a jail on the SSOC Space Station.

It is assumed here that NASA will wish to avoid conflict in SSOC. The primary reason for this is cost. The dollar expense per crew man-hour aloft is very high (est. \$40,000 per man-hour), and it is obviously undesirable to waste time through lack of coordination or, worse, through the need to resolve open conflict.

A second reason for avoiding conflict in SSOC is that conflict in social systems often feeds on itself. That is, an initial conflicted encounter may lead to hard feelings, disliking, and attitudes of distrust toward out-group members, as well as the development of cognitive biases and stereotypes (Wilder, 1981; Brewer, 1986; Wilder and Cooper, 1981). This makes subsequent cooperation harder to achieve, and may even intensify the problem (i.e., "escalation of conflict"). Interpersonal conflict changes the attitudes and beliefs of people involved, and this change is often for the worse when viewed from the standpoint of system performance (Cooper and Fazio, 1986; Michener et al., 1986; Pruitt and Rubin, 1986).

In the following sections, then, consideration is given to various means by which NASA, through its design efforts, can reduce the risk of conflict among the crew in SSOC. These means include the alignment of goal structures, patterning of social interaction, and selection and training of crew members. The fundamental research/design issue underlying this discussion is how to design the SSOC social system to avoid or minimize interpersonal conflict; a related issue is how to equip the crew with techniques to resolve conflict (if it occurs) in a manner that is constructive from the standpoint of the larger system.

Conflict Avoidance via Goal Structure

Various approaches are available to NASA for avoiding and/or reducing conflict in the SSOC social system. One of the more effective is to give close attention to the design of, and alignment among, subgroup goals.

The Importance of Goal Alignment

As noted above, opposition of interests among subgroups in differentiated social systems is an important factor producing conflict. With opposition of interests, overt conflict frequently occurs; without it, there is no reason for conflict to occur (Campbell, 1965; Sherif et al., 1961)

In IOC, there will not be much opposition of interests among crew members. The Space Station will have a single coherent goal (i.e., an objective function that specifies what should be maximized by system performance). The main mission will be to fly the Space Station and to carry out evolutionary expansion of the Station via construction. Crew members will not be working at cross-purposes. In contrast, during SSOC the Space Station will have a more complex objective function. It may even have more than one objective function because, in addition to the function for the entire system, there may exist separate sub-functions for each of the subgroups on board. Conflict might arise, for instance, over manpower scheduling or over allocation of scarce resources such as electrical power. In SSOC there will be at least some risk that one or several subgroups on board may have (or develop) goals that do not mesh smoothly with those of other subgroups.

An important research/design issue for NASA is to specify objective function(s) for the SSOC crew such that the attainment of goals by one subgroup does not prevent the attainment of goals by some other subgroup(s). Well-conceived objective functions will promote harmony and productivity; conversely, ill-conceived or misaligned goals will doubtless generate interpersonal and intergroup conflict.

The Superordinate Goals Approach

One approach to aligning goals among SSOC subgroups is to establish objective functions that embody what are called "superordinate goals" (Sherif et al., 1961; Blau and Mouton, 1968, 1976, 1984). A superordinate goal is one that (a) is held to be important by each of the subgroups comprising the larger social system and (b) can be attained only through cooperative interaction among subgroups (i.e., cannot be attained by a single subgroup acting alone). Superordinate goals induce a high coincidence of interest among diverse subgroups.

Research on Earth-based social systems has shown repeatedly that superordinate goals inhibit conflict among subgroups. Moreover, in social systems where the subgroups are already engaging in open conflict, the introduction of new superordinate goals can mitigate or resolve conflict (Sherif et al., 1961). Superordinate goals reliably improve cooperation and reduce conflict among subgroups in a larger system.

There may be several ways to incorporate superordinate goals in the design of the SSOC social system. One particularly interesting possibility is to include such goals in the computer software used on board the Space Station. This becomes especially viable if NASA uses some kind of "project scheduler" software to assign tasks to crew

members. Software of this type entails optimization in some form or another; when designing this software, NASA will have to decide exactly what is to be optimized. It is suggested here that what should be optimized in SSOC is not merely "productivity," but also system integration. Both concerns are important. The design and use of project scheduler software provides an opportunity to expressly incorporate goals that bind the subgroups together.

The Game-Theoretic Approach

An alternative approach to goal design is to treat the relations among subgroups in SSOC as a set of n-person games (Shubik, 1982, Owen, 1982, Vorob'ev, 1977; Leitman, 1976). That is, the subgroups in SSOC might be viewed as players having (somewhat) opposing interests in n-person non-constant-sum games. These games could be analyzed to identify points of contention between subgroups and likely outcomes of conflict.

Specifically, one might first identify a set of scenarios (situations) that could arise on board the Space Station, and then treat each of these as a distinct n-person game. These scenarios might include such events as EVAs, health emergencies, payload experimentation, space debris emergencies, etc. Each could be analyzed in terms of the likely equilibrium outcome under some solution concept (e.g., the Nash non-cooperative equilibrium). Results of such an analysis would show the extent to which the subgroups have opposing interests and indicate whether they would play a strategy leading to an outcome that is not desirable collectively (i.e., not Pareto optimal).

The point of conducting such an analysis is not only to anticipate issues over which conflict might erupt, but eventually to design the subgroups' objective functions to assure that the payoff matrices for most n-person games played on board lead to a benign equilibrium.

Persons within NASA are familiar with the game theoretic approach; NASA used game theory to resolve conflict among groups of engineers with competing demands regarding equipment to be placed on the Mariner spacecraft. There may be opportunity again to use it advantageously in SSOC.

Conflict Avoidance via Patterned Social Interaction

Another broad approach to avoidance of conflict in SSOC entails intentional structuring or channeling of social interaction among crew members. In particular, NASA might (a) design the supervisory-control structure so that it detects and resolves conflict readily, (b) structure the interpersonal contact on board the Space Station to minimize the probability of conflict occurring, and (c) structure communication on board so that message-type maps into media-types in a way that lessens the probability of conflict. Each of these is discussed below.

Conflict and the Supervisory-Control Structure

Usually it is better to prevent conflict before it arises than to attempt to resolve it after it has escalated. For this reason, when designing the onboard supervisory-control structure for SSOC, NASA may wish to include what are termed "boundary-spanning roles" (Adams, 1976; Wall, 1974; Katz and Kahn, 1966; Holmes et al., 1986). These are roles the occupants of which perform functions that link subgroups together. For instance, persons in boundary-spanning roles may communicate across groups on sensitive issues, or serve as representatives in decision-making that affects the relations between subgroups. Because the SSOC social system will contain several subgroups, the inclusion of boundary-spanning roles in the larger system may help to avoid conflict between groups and to resolve conflict should it occur.

In systems without boundary-spanning roles, one typical consequence of conflict is a reduction or cessation of communication between the parties. Any such reduction of communication would obviously be undesirable in SSOC. The use of boundary-spanning roles in SSOC may be a way of establishing--and of keeping open--channels between the nationality groups on board. In addition, occupants of boundary-spanning roles can also serve as negotiators with respect to points of contention between subgroups.

In sum, the use of boundary-spanning roles in SSOC may provide a mechanism for avoiding conflict. The research/design issues for NASA are exactly what boundary-spanning roles, if any, to include in SSOC, and how to interface these roles with the activities of the Space Station's Mission Commander and other administrators. One possibility in this regard is to design the role system such that persons who will serve as lieutenants to the Mission Commander will also function as boundary-spanners.

Conflict and Interpersonal Contact

A related research/design issue is how best to structure interpersonal contact among regular crew members to promote cohesive, non-polarizing relations among subgroups in SSOC.

Research on Earth-based systems suggests that NASA might reduce the probability of conflict between groups by assigning tasks to crew members with an eye not just to getting work done, but also to promoting cooperative contact and interdependence among persons from different subgroups (Amir, 1969; Worchel et al., 1977, 1978; Deutsch, 1973; Worchel, 1986). For instance, NASA might assign tasks such that persons from different nationality groups work on an interdependent basis. Under such an arrangement, both Americans and Europeans would do EVA, both Japanese and Europeans would do payload operations (experiments), and so on. The situation to avoid is one where the Japanese do all the EVA, the Europeans do all the payload operations, the Americans do all the flying, etc. The key is to create task-interdependence and cross-linkages among nationality groups.

Another potential overlap is that between Astronaut Pilots and Payload Specialists. If there are only four or six Astronaut Pilots on board in SSOC, there may not be much opportunity for task overlap between these groups. If there are many Astronauts on board, however, tasks can be assigned to promote collaboration. Some Astronaut Pilots might be assigned to conduct payload experiments on an interdependent basis with Payload Specialists. Again, the objective is to create ties across subgroups.

Beyond task interaction, NASA may also find it possible to structure non-task activities among crew members in such a way as to develop ties across subgroup boundaries. Of course, most waking hours each day will be spent on tasks (12 hours/day); crew members will have little time for non-task activities. Yet, non-task interaction may prove important in creating and maintaining positive attitudes and trust across subgroups, in part because the size of the SSOC crew will preclude all members from interacting with one another in a task mode.

Some research on Earth-based systems shows that informal contact across subgroups is most effective in strengthening intergroup bonds when it is conducted on an equal-status basis (Amir, 1969, 1976; Norvell and Worchel, 1981). Exactly how to do this in SSOC is an open issue. For instance, it may be desirable to assign spatial living quarters to create cross-linkages among nationality groups. That is, assign some USA astronauts to sleep in the Japanese module and the European module, assign Japanese and European astronauts to one another's modules and to USA module, etc. Alternatively, it may prove desirable to have crew members of different subgroups eat together (this will not carry special meaning for Americans, but it may for the Europeans). How to structure informal contact in SSOC to strengthen intergroup bonds is an open research/design issue for NASA.

Conflict and Communication

The communication system on board the Space Station in SSOC will differ from that in IOC. The size of the SSOC communication network will be larger (i.e., contain more nodes) than that in IOC because the crew will be larger in size. Moreover, the total communication flow (number of messages sent) will be higher in SSOC, although the messages per crew member may remain about the same. Communication flows in SSOC will reflect the clustering of crew members into subgroups; flows will be higher within and lower between subgroups.

From the standpoint of conflict and conflict resolution, however, the most critical difference between IOC and SSOC will be the media of communication used. During IOC, when the Space Station will have a small crew housed in a single module, a significant proportion of communication will doubtless be face-to-face. In SSOC, with a larger crew dispersed in several modules, a smaller proportion of communication will be face-to-face and a larger proportion will occur via other media such as telephone and electronic (computer) mail. This will result naturally because SSOC crew members will have to communicate with others in remote locations in the Space Station.

The shift in communication media between IOC and SSOC may be important because the various media have different properties. Telephones and computers, for example, do not convey some types of information as fully as the face-to-face channel (Mehrabian, 1972). Face-to-face communication transmits linguistic, paralinguistic, kinesic and proxemic cues, while electronic (computer) mail transmits linguistic cues only (Connors et al., 1984; Danford et al., 1983; Hall, 1968). One important consequence is that non-face-to-face media carry less information about personal relations and feelings. Thus, in view of the SSOC system's potential for fractionating conflict, heavy use of non-face-to-face media in SSOC may produce undesirable consequences.

Computer-mediated communication is especially problematic in this regard. The effects of computer-mediated communication are not yet fully understood, but it is increasingly clear that this medium is good for some purposes, poor for others. Computer conferencing tends, for instance, to increase equality of participation more than face-to-face conferencing (Johansen et al., 1979), which may improve the potential for circumspect consideration of issues. Electronic mail is not, however, a good medium by which to conduct bargaining or to resolve interpersonal conflict, because it can foster one-sided proclamations and policy statements couched in concepts not shared by participants. More generally, computer-mediated communication may be less effective than face-to-face communication for reaching consensus on issues where the "correct" answer is not obvious. In addition, research shows that use of computer-mediated communication sometimes leads to polarization and flaming (Kiesler, et al., 1984). Behavior of this type would be especially undesirable in SSOC, given the subgroup differentiation projected for the social system.

The burden placed on computer-mediated communications will increase in SSOC in the sense that failures to communicate adequately may have more serious consequences in SSOC than in IOC. Communication failures will assume higher criticality in SSOC due to the differentiated nature of the social system. To communicate across cultures is difficult enough via face-to-face interaction; to rely heavily on media that filter information in unpredictable ways will make the communication problem even worse.

Thus, a general research/design issue for NASA is how may the SSOC crew best use the communication media on board the Space Station to promote non-polarizing interpersonal contact and to create cross-linkages between members of subgroups.

At the least, NASA may wish to develop an "etiquette" regarding use of the various media on board. This may include not only rules for the use of media, but also rules regarding what types of messages are to be sent over which media. Some theorists have hypothesized a (statistical) interaction effect between media type and message type on communication effectiveness (Geller, 1980; Danford et al., 1983). In view of this, one approach to the SSOC communication problem is to seek a match between media and the type of message being sent (i.e., where "type" refers to message content coded from the standpoint of its functionality for the social system). That is, to achieve high communication effectiveness, send some types of messages by one

channel, other types by other channels. To achieve such regulation, the Space Station will need norms specifying what types of messages are sent via computer mail, what types via telephone, and what types via face-to-face contact. The exact nature of these norms is an open issue.

Conflict Avoidance via Selection and Training

NASA has traditionally placed great emphasis on selection and training of its crews. Selection and training will continue to play an important part in IOC and SSOC Space Station operations. The potential for conflict in SSOC, however, implies that when NASA moves from IOC toward SSOC, it may wish to make some adjustments both in the criteria used to select crew members and in the content of Astronaut training. An important research/design question is what should be the nature of these changes.

Crew Composition and Selection

Certain obvious shifts will occur in NASA crew selection activities from IOC to SSOC. First, the number of persons selected will increase, because NASA will be flying larger crews. Second, the skill-mix of persons selected will shift; compared with IOC, a larger proportion of crew members will be Payload Specialists, a smaller proportion Astronaut Pilots. Third, the nationality of persons on the Space Station will change, to include Japanese and Europeans.

Less self-evident is that, when moving from IOC to SSOC, NASA may find it necessary to change its crew selection criteria. To enhance integration of the SSOC social system, NASA may opt for crew members who, by virtue of their background, can serve as linking-pins across subgroups. For example, in SSOC there may be a premium on crew members who have a background of cross-cultural or international experience, or who are multi-lingual (e.g., NASA Astronauts who speak French, or who have lived in Japan). Alternatively, NASA may choose to "manufacture" persons with such backgrounds by, for example, having its pilots live in Europe or Japan for several years.

Another possible change concerns the personality profile of the ideal Astronaut. In IOC, with small crew size, there will be a premium on persons who are high on interpersonal compatibility and who relate well to others. The concept of interpersonal compatibility, however, is more applicable to small groups of 6-8 than to larger groups of 20-30. Rarely does one find a group of 30 persons, all of whom are interpersonally compatible. Thus, in SSOC, the emphasis on compatibility may fade and give way to other interpersonal skills, such as diplomacy. More generally, a research/design issue for NASA is to discover which personal attributes of crew members best serve to enhance linkages between subgroups in SSOC.

Conflict and Crew Training

Astronauts from different countries and reared in different cultures will hold different expectations regarding patterns of social interaction. Although these may not affect the technical aspects of space flight, some will seriously affect interpersonal sentiments. For instance, respectful interpersonal treatment among the Japanese looks different from that among the Americans or the French. Without adequate preparation, misunderstandings will arise among crew members. NASA may wish to address the implications of this when training Astronauts for SSOC.

Emphasis throughout this section had been on avoidance of conflict. Even with the best preparation, however, some conflict will occur in SSOC. For this reason, NASA may wish to train crew members in conflict resolution techniques. When persons are under stress, some forms of communication and negotiation are more effective than others (Pruitt, 1981; Rubin and Brown, 1975). Useful conflict management skills in American society include: reflective listening, assertion skills, issue control, structured exchange regarding emotional aspects of a controversy, and collaborative problem solving (Bolton, 1979; Walton, 1969). Whether these techniques will work in a cross-cultural context like the SSOC social system is an open issue. If they do work, NASA may wish to include them in its training regimen. Their use could increase crew's effectiveness in dealing with interpersonal disagreements when they arise on board the Space Station. In sum, an important research/design issue is exactly what conflict resolution skills should be taught to crew members.

COPING WITH ENVIRONMENTALLY-INDUCED CRISES

Crisis: A Definition

As used here, the term "crisis" refers to a circumstance in which something threatens to destroy or impair the social system on board the Space Station, and which therefore requires an immediate response from crew members (as well as from Mission Control) to assure the continued functioning of the system. Crises can be precipitated by many different events. For instance, crises might result if: (a) a sudden leak or air-loss occurs, causing the cabin pressure to decline sharply, (b) a sudden loss of power occurs, (c) a crew member becomes seriously ill, (d) some space debris collides with the Space Station, producing serious damage, (e) one of the bio-experiments on board goes awry, releasing pathogens or contaminants that pose a threat to humans, or (f) fire erupts on board the Space Station. This list is illustrative, not exhaustive.

Most of the events listed here are improbable, in the sense that they will occur only infrequently. However, the Space Station will operate in a perilous environment for a planned 25-30 years and, while the probability of a crisis on any given day may be low, the odds of avoiding crises are much less favorable over the full span of 25-30

years. Although not inevitable, one or several crises are probable during the operational lifetime of the Space Station.

Normal Operating Mode vs. Crisis Operating Mode

The structure of many systems in nature is controlled by the manner in which the system might fail (von Neumann, 1966; Weinberg, 1975). In other words, natural systems often incorporate some precautionary measures to prevent failure, or at least to prevent a failure from being lethal. Social systems also display this characteristic, and they often cope with crisis and failure by having several distinct operating modes, such as "normal operating mode" vs. "crisis operating mode." In normal operating mode, when the environment is not disruptive, the social system conducts "business as usual." Human plans drive the action, and the emphasis is on productivity and performance. However, in crisis operating mode, when the social system responds to environmental threats, there is a shift in the social system's objective function. The predominant goal in crisis mode becomes that of assuring the very survival of the system, and activities are reorganized in terms of this goal. Environmental contingencies, not human plans, drive the action; persons in the system become more reactive and less proactive.

Most likely, the IOC and SSOC social systems will use several operating modes. They may even implement several distinct crisis operating modes, contingent upon whatever types of crises occur. Nevertheless, crisis management in SSOC probably will differ from that in IOC, in part because the shift from normal mode to crisis mode will be more difficult to accomplish in SSOC than in IOC.

Crisis Management

Crisis Management in IOC

Crisis-management is never easy, but the characteristics of the IOC social system will equip it well to respond to crises when they arise. The small size and great homogeneity of the crew, the housing of the crew in a single habitat module, and the nature of the supervisory-control structure will enable the IOC system to switch quickly to crisis operating mode from normal operating mode. In IOC, crisis operating mode will (a) establish centralized control of crew activities, (b) assure adequate information flow among members, (c) create the potential for clear, consensual decision-making, (d) rapidly establish coordination among crew members, and (e) apply the greatest expertise available to the problem. In social systems, these are desirable features under emergency conditions.

The IOC's supervisory-control structure, assumed to be patterned after a hierarchical "military command model," will function fairly well during a crisis. Because command is centralized, the system will hold together and coordination of action will be attainable even under

stress. The hierarchical structure will enable the IOC system to focus resources, restrict non-adaptive responses (such as argumentation or countermanding), and achieve an adequate level of communication among crew members. In general, it can provide the high level of interpersonal organization needed to respond to crises.

Crisis Management in SSOC

The SSOC social system will have more resources than the IOC system to deal with crises. For example, its hardware may have better sensors to anticipate crisis-precipitating events before they happen, its expert-system software may provide more accurate diagnoses of problems, and its crew may include a greater mix of skills useful during crises. Nevertheless, crisis-management in SSOC will present its own problems. The incidence of crises may be higher in SSOC than in IOC, because there will be more things to go wrong. There will be more crew members to get sick, more area to get hit by space debris, more bio-experiments to blow up, more on-board hardware to malfunction, etc. Moreover, the organizational form of the SSOC social system will make it more difficult to respond adequately to crises. The SSOC system may have more difficulty switching from normal operating mode to crisis operating mode than the IOC system.

The SSOC social system will be larger, more complex, and more differentiated than IOC. Moreover, as noted above, supervisory-control and decision-making in SSOC will be decentralized in normal operating mode. The presence of different nationality groups and of many Payload Specialists performing diverse tasks will create a heterarchical supervisory-control structure. If a crisis arises, the supervisory-control structure in SSOC must coordinate the response of distinct subgroups living in different physical modules and pursuing divergent goals. This task is not impossible, but it will be more difficult than in IOC.

In all likelihood, a shift from normal operating mode to crisis operating mode in SSOC will entail a quick move from a decentralized heterarchical structure to a centralized hierarchical one. Failure to move back to a hierarchy during a crisis in SSOC will leave the system vulnerable. If the Space Station relied on a decentralized system during crisis, it would risk lack of coordination among crew members, less-than-optimal deployment of resources to deal with the problem, and perhaps even disagreement over the best type of response to the emergency.

Although a shift from heterarchy to hierarchy during crisis seems likely, the exact form of SSOC command during crises is an open research/design issue. Danford et al. (1983) have suggested that it would be appropriate to have control during crisis rest in the hands of a specialized safety officer or "crisis leader." This scheme has some merit, but it may also create excessive complexity because that it requires yet another form of control beyond the heterarchy-plus-Mission Commander structure discussed above. A superior alternative might be

simply to recentralize control during a crisis around the regular leader (Mission Commander).

Recentralization around the Mission Commander will work best if NASA trains crew members in specific skills for dealing with different types of crises. That is, some crew members will be specialists in coping with one type of crisis and other crew members with another type of crisis. Thus, when a crisis occurs, two things will happen. First, crew members will coordinate around the Mission Commander; and second, the Mission Commander, assisted by those persons who are specialists in the particular type of crisis at hand, will direct the efforts of the entire crew to cope with the emergency. This approach brings both special expertise and strengthened command to bear in a crisis.

A related research/design issue concerns the use of AI and computerization to aid decision-making during crises. Expert systems that diagnose the causes of hardware failures will be operational increasingly as the Space Station moves from IOC to SSOC, and these may increase the speed and accuracy of the crew's efforts during crises. To some degree, expert systems will be able to supplement (even supplant) the knowledge and expertise of crew members. On the other hand, use of AI systems in the analysis and diagnosis of life-threatening events raises the issue of trust—to what extent will crew members trust software-based diagnoses. The use of AI may affect not only how the crew is organized to cope with crises, but also what mix of skills is (and is not) placed on board and how crew members are trained. These are matters that can be addressed through research and design efforts.

One final research/design issue concerns the impact of computer-mediated communication during crises. As noted above, computer-mediated communication will be even more important and prevalent in SSOC than in IOC. Whether computer-mediated communication enhances or inhibits satisfactory responses to crises is an open question. It was noted above that computer-mediated communication may be less effective than face-to-face communication for reaching consensus on issues where the "correct" answer is not obvious. Some crises on board the Space Station may have clear-cut diagnoses, but for those that do not, computer-mediated communication may prove more a liability than an asset in achieving adequate response from the crew. The (in)effectiveness of computer-mediated communication during crises is an important research topic.

SUMMARY OF RESEARCH AND DESIGN ISSUES

This paper has discussed issues that arise in the design of the SSOC social system. Attention has been given to three broad problem areas: (a) the characteristics of the SSOC supervisory-control structure, (b) the potential for conflict within the crew, and (c) the capacity of the SSOC system to respond to crises if they arise. Specific research suggestions are summarized below.

Issues Regarding SSOC Supervisory-Control

One important research/design issue for NASA is what type of supervisory-control structure will best serve the SSOC social system, in the sense of providing the greatest efficiency and highest probability of mission success. There are a wide variety of supervisory-control structures that might be deployed on board the Space Station--hierarchical, equalitarian, heterarchical, etc.--and the exact nature of the system to be used is an open issue.

It has been proposed here that the Space Station's supervisory-control structure will take the form of a hierarchy in IOC, and that it may subsequently shift in the direction of a heterarchy in SSOC. This is really no more than a conjecture, however. NASA can make decisions regarding the form of supervisory-control structure to be used in IOC and SSOC on the basis of trial-and-error or past experience with space flight supervision. Alternatively, it might make them on the basis of research findings, such as those obtainable from simulations conducted on the ground.

Specifically, it was suggested above that NASA might conduct simulations to test various outcomes from different supervisory-control structures. These simulations would be done under conditions that closely replicate those found in space--e.g., high stress, high noise, restricted communication, 90-day duration, tasks similar to those done in space, and so on. Major outcome measures include productivity levels, crew satisfaction, lack of conflict, adequacy of response to emergencies, etc. Multiple replications could be run on each of several alternative supervisory-control structures using experimental designs. The results should provide a useful indication of how the alternative supervisory-control structures will perform in space.

One design sub-problem is to determine the appropriate division of control between Space Station crew and Mission Control on Earth. One concrete manifestation of this problem is the issue of who should have control over the crew's day-to-day task assignments. Various suggestions, including the use of AI project planning software to accomplish task assignments, were discussed.

A second design sub-problem is to determine the appropriate division of control within the Space Station's crew. Presumably the Task Specialists in SSOC will be afforded some degree of independence with regard to their particular activities, but the exact range is unclear. The Mission Commander's role during SSOC will likely shift toward coordination of other activities, but the exact definition of the role's prerogatives and powers is problematic.

A related issue is the selection criteria regarding crew officers; this matter is made more complex by the inclusion of crew members from other space agencies (Japan, Europe). Whether the role of Mission Commander will be restricted to NASA Astronauts or open to crew members from other countries is a research/design issue that NASA might address.

Issues Regarding Crew Conflict in SSOC

The risks of interpersonal and intergroup conflict will be greater in SSOC than in IOC. This is true in part because the SSOC system will include many subgroups with distinct identities (Task Specialists/Astronauts; and USA/Japan/Europe). The broad research/design question for NASA is what safeguards to build into the SSOC system to reduce the probability of overt conflict occurring, and to resolve conflict if it occurs.

A wide variety of steps can be taken in the design of the SSOC system to reduce the probability of conflict. Some discussed in this paper include:

(a) Specify objective function(s) for the SSOC crew such that the attainment of goals by one subgroup does not prevent the attainment of goals by other subgroup(s). Approaches to this include the use of superordinate goals and game-theoretic analysis of subgroup interaction. One implementation might involve computer software (project scheduler routines) to optimize not just productivity but also group overlap.

(b) Incorporate boundary-spanning roles in the SSOC social system. An open question is how to interface these roles with the activities of the Space Station's Mission Commander and other officers.

(c) Structure interpersonal contact among crew members to promote cohesive, non-polarizing relations across the subgroups in SSOC. Crew members might be assigned tasks with an eye to creating interdependence and cross-linkages between nationality groups. Likewise, module living and sleeping assignments might be made to promote contact across nationality groups.

(d) Use of the communication media on board the Space Station to promote non-polarizing interpersonal contact and cross-linkages between members of subgroups. Computer-mediated communication is especially problematic in this respect, for it may worsen, not improve, the prospects for intergroup conflict. NASA may wish to develop some rules or "etiquette" regarding use of computers for communication.

(e) When moving from IOC toward SSOC, NASA may need to make some adjustments in the criteria used to select crew members and in the content of Astronaut training. In this regard, a research/design issue for NASA is to discover which personal attributes of crew members best serve to enhance linkages between subgroups in SSOC. Another issue is to determine what conflict resolution skills should be taught to crew members.

Issues Regarding Response to Crises in SSOC

The SSOC social system may have more difficulty than the IOC system in mobilizing to deal with various crises and emergencies on board. This will occur not only because SSOC is a larger system, but also because it is more heterarchical in form with decentralized decision-making. The broad research/design question for NASA is how best to structure the SSOC social system so that it can mobilize adequately for crises.

Some writers have suggested placing control during crises in the hands of a specialized safety officer or "crisis leader." This proposal has some merit, but a better alternative may be to recentralize control around the regular Mission Commander. NASA may wish to investigate this research/design issue more closely.

Moreover, NASA might investigate the use of AI expert systems to help deal with crises--the software system becomes the crisis advisor, assisting or even supplanting human decision-making. Use of expert systems in this context may improve diagnosis of the problem, as well as increase speed and accuracy of response to the emergency.

Finally, NASA may wish to investigate the (in)effectiveness of computer-mediated communication during crises. Whether computer-mediated communication enhances or inhibits responses to crises is an open question. Some crises on board the Space Station may have clear-cut diagnoses, but for those that do not, computer-mediated communication may prevent or diminish an adequate response from the crew. The effects of computer-mediation on communication during crises merits scrutiny.

REFERENCES

- Adams, J. S.
 1976 The structure and dynamics of behavior in organization boundary roles. In M. D. Dunette, ed., Handbook of Industrial and Organizational Psychology. Chicago: Rand-McNally.
- Anir, Y.
 1969 Contact hypothesis in ethnic relations. Psychological Bulletin, 71:319-342.
- 1976 The role intergroup contact in change of prejudice and ethnic relations. In P. A. Katz, ed., Towards the Elimination of Racism. New York: Pergamon.
- Balbaky, E. M. L.
 1980 Strike in Space. Case #1-431-008. Boston, MA: Harvard Business School.
- Blake, R. R., and Mouton, J. S.
 1968 Corporate Excellence Through Grid Organization Development. Houston: Gulf.
- 1976 Diary of an OD Man. Houston: Gulf.
- 1984 Restoring Trust Between Groups in Conflict. San Francisco: Jossey-Bass.

- Bluth, B. J.
 1980 Social and psychological problems of extended space missions. AIAA International Meeting and Technical Display "Global Technology 2000". New York: American Institute of Aeronautics and Astronautics.
- 1981 Soviet space stress. Science 81:2:30-35.
- Bolton R.
 1979 People Skills. Englewood Cliffs, NJ: Prentice-Hall.
- Brewer, M. B.
 1986 The role of ethnocentrism in intergroup conflict. In S. Worchel and W. G. Austin, eds., Psychology of Intergroup Relations. Chicago: Nelson-Hall Publishers.
- Brewer M. B., and Campbell, D.T.
 1976 Ethnocentrism and Intergroup Attitudes. New York: Halsted Press.
- Campbell, D. T.
 1965 Ethnocentrism and other altruistic motives. In D. Levine, ed., Nebraska Symposium on Motivation. Vol. 13. Lincoln: University of Nebraska.
- Casti J. L.
 1979 Chapter 4 in Connectivity, Complexity, and Catastrophe in Large-Scale Systems. New York: Wiley.
- Connors, M. M., Harrison, A. A., and Akins, F. R.
 1984 Living Aloft: Human Requirements for Extended Spaceflight. NASA SP-483. Washington, DC: U. S. Government Printing Office.
- Cooper, H. S. F., Jr.
 1976 A House in Space. New York: Holt, Rinehart, and Winston.
- Cooper, J. and Fazio, R. H.
 1986 The formation and persistence of attitudes that support intergroup conflict. In S. Worchel and W. G. Austin, eds., Psychology of Intergroup Relations. Chicago: Nelson-Hall Publishers.
- Cunningham, W.
 1977 The All-American Boys. New York: Macmillan.

Danford, G. S., Dry, C. M., Ericson, J. E., Freitas, Jr., R. A., Hays, D. G., Lewis, Jr., W. C., Marsolan, Jr., N. F., Mazlack, L. J., Schoonhoven, C. K., Sheskin, T. J., Wallace, R. S., Wellens, A. R., and Wise, J. A.

- 1983 Evolution of space station autonomy. Chapter 7 in Autonomy and the Human Element in Space. Final Report of the 1983 NASA/ASEE Summer Faculty Workshop, NASA.

Danford, G. S., Ericson, J. E., Freeman, H. J., Hays, D. G., Lewis, Jr., W. C., Schoonhoven, C. K., Sims, H. P., Walsh, P. J., Wellens, A. R., and Wise, J. A.

- 1983 Humane space station. Chapter 5 in Autonomy and the Human Element in Space. Final Report of the 1983 NASA/ASEE Summer Faculty Workshop, NASA.

Danford, G. S., Ericson, J. E., Marsolan Jr., N. F., Mazlack, L. J., Schoonhoven, C. K., Wellens, A. R., and Wise J. A.

- 1983 Autonomy in the space station context. Chapter 2 in Autonomy and the Human Element in Space. Final Report of the 1983 NASA/ASEE Summer Faculty Workshop, NASA.

Deutsch, M.

- 1973 The Resolution of Conflict. New Haven: Yale University Press.

Diab, L.

- 1970 A study of intragroup and intergroup relations among experimentally produced small groups. Genetic Psychology Monographs 82:49-82.

Geller, V. J.

- 1980 Toward a Theoretical Framework for Investigating Modality Effects. Bell Laboratories Technical Memo. Murray Hill, NJ.

Hall, E. T.

- 1968 Proxemics. Current Anthropology 9:83-108.

Harrison A. A., and Connors, M. M.

- 1984 Groups in exotic environments. In L. Berkowitz, ed., Advances in Experimental Social Psychology. Vol. 18. New York: Academic Press.

Helmreich, R. L.

- 1983 Applying psychology in outer space: Unfulfilled promises revisited. American Psychologist 38:445-450.

- Helmreich, R. L., Wilhelm, J., Tanner, T. A., Sieber, J. E., and Burgenbach, S.
 1979 A Critical Review of the Life Sciences Management at Ames Research Center for Spacelab Mission Development Test III. NASA Tech. Paper No. 1464.
- Henschy, T., and Glass, D. C.
 1968 Evaluation apprehension and social facilitation of dominant and subordinate responses. Journal of Personality and Social Psychology 10:446-454.
- Holmes, J. G., Ellard, J. H., and Lamm, H.
 1986 Boundary roles and intergroup conflict. In S. Worchel and W. G. Austin, eds., Psychology of Intergroup Relations. Chicago: Nelson-Hall Publishers.
- Johansen, R., Vallee, J., and Spangler K.
 1979 Electronic Meetings. Menlo Park, CA: Addison-Wesley.
- Katz, D., and Kahn, R. L.
 1966 The Social Psychology of Organizations. New York: Wiley.
- Keeney R. L., and Raiffa, H.
 1976 Decisions with Multiple Objectives: Preferences and Value Tradeoffs. New York: Wiley.
- Kiesler, S., Siegel, J., and McGuire T. W.
 1984 Social psychological aspects of computer-mediated communication. American Psychologist 39:1123-1134.
- Kurtzman, C. R., Marsolan Jr., N. F., Mazlack, L. J., Rinalducci, E. J., Topham, W. S., Walsh, P. J., and Wong, D.
 1983 Monitoring and control. Chapter 4 in Autonomy and the Human Element in Space. Final Report of the NASA/ASEE Summer Faculty Workshop, NASA.
- Jeitman G., ed.
 1976 Multicriteria Decision Making and Differential Games. New York: Plenum Press.
- Leonov, A. A., and Lebedev, V. I.
 1975 Psychological Problems of Interplanetary Flight. NASA Tech. Translation. NASA TT F-16536.
- Marcus H.
 1978 The effect of mere presence on social facilitation: An unobtrusive test. Journal of Experimental Social Psychology 14:389-397.
- Mehrabian A.
 1972 Nonverbal Communication. New York: Aldine-Atherton.

- Mesarovic, M. D., Macko, D., and Takahara, Y.
1970 Theory of Hierarchical, Multilevel, Systems. New York: Academic Press.
- Michener, H. A., DeLamater, J. D., and Schwartz S. H.
1986 Chapter 15 in Social Psychology. San Diego: Harcourt Brace Jovanovich.
- Nelson, P. D.
1973 The indirect observation of groups under confinement and/or isolation. Pp. 167-194 in J. E. Rasmussen, ed., Man in Isolation and Confinement. Chicago: Aldine.
- Norvell, N., and Worchel, S.
1981 A reexamination of the relation between equal-status contact and intergroup attraction. Journal of Personality and Social Psychology 41:902-908.
- Oberg, J. E.
1981 Red Star in Orbit. New York: Random House.
- Owen, G.
1982 Game Theory. 2nd edition. New York: Academic Press.
- Pogue, W. R.
1985 How Do You Go to the Bathroom in Space? New York, NY: Tom Doherty Associates.
- Pruitt, D. G.
1981 Negotiation Behavior. New York: Academic Press.
- Pruitt, D. G., and Rubin, J. Z.
1986 Social Conflict. New York: Random House.
- Rubin, J. Z., and Brown, B. R.
1975 The Social Psychology of Bargaining and Negotiation. New York: Academic Press.
- Schoonhoven, C. B.
1986 Sociotechnical considerations for the development of the space station: Autonomy and the human element in space. Journal of Applied Behavioral Science 22:271-286.
- Sells, S. B., and Gunderson, E. K.
1972 A social-systems approach to long duration missions. Pp. 179-208 in Human Factors in Long-Duration Spaceflight. Washington, DC: National Academy of Sciences.
- Sherif, M., Harvey, O. J., White, B. J., Hood, W. R., and Sherif C.
1961 Intergroup Cooperation and Competition: The Robbers Cave Experiment. Norman, OK: University Book Exchange.

- Shubik M.
1982 Game Theory in the Social Sciences: Concepts and Solutions. Cambridge, MA: MIT Press.
- Sloan, A. W.
1979 Man in Extreme Environments. Springfield, IL: Thomas.
- Sutherland, J. W.
1975 Pp. 32-38 and Chapter 2 in Systems: Analysis, Administration and Architecture. New York: Van Nostrand Reinhold.
- Tajfel, H., and Turner, J. C.
1986 The identity theory of intergroup behavior. In S. Worchel and W. G. Austin, eds., Psychology of Intergroup Relations. Chicago: Nelson-Hall Publishers.
- Von Neumann, J.
1966 Theory of Self-Reproducing Automata. Urbana: University of Illinois Press.
- Von Tiesenhausen, G.
1982 An Approach Toward Function Allocation Between Humans and Machines in Space Station Activities. NASA TM-82510. Marshall Space Flight Center, Alabama.
- Vorob'ev, N. N.
1977 Game Theory: Lectures for Economists and System Scientists. New York: Springer-Verlag.
- Wall, J. A.
1974 Some variables affecting a constituent's evaluations of and behavior toward a boundary role occupant. Organizational Behavior and Human Performance 11:390-408.
- Walton, R.
1969 Interpersonal Peacemaking: Confrontations and Third-Party Consultation. Reading, MA: Addison-Wesley.
- Weinberg, G. M.
1975 Chapter 7 in An Introduction to General Systems Thinking. New York: Wiley.
- Wilder, D. A.
1981 Perceiving persons as a group: Categorization and intergroup relations. In D. L. Hamilton, ed., Cognitive Processes in Stereotyping and Intergroup Behavior. Hillsdale, NJ: Erlbaum.

- 1986 Cognitive factors affecting the success of intergroup contact. In S. Worchel and W. G. Austin, eds., Psychology of Intergroup Relations. Chicago: Nelson-Hall Publishers.
- Wilder, D. A., and Cooper, W. E.
1981 Categorization into groups: Consequences for social perception and attributions. In J. Harvey, W. Ickes, and R. Kidd, eds., New Directions in Attribution Research. Vol. 3. Hillsdale, NJ: Erlbaum.
- Worchel, S.
1986 The role of cooperation in reducing intergroup conflict. In S. Worchel and W. G. Austin, eds., Psychology of Intergroup Relations. Chicago: Nelson-Hall Publishers.
- Worchel, S., Andreoli, V. A., and Folger, R.
1977 Intergroup cooperation and intergroup attraction: The effect of previous interaction and outcome of combined effort. Journal of Experimental Social Psychology 13:131-140.
- Worchel, S., Axson, D., Ferris, F., Samaha, C., and Schweizer, S.
1978 Determinants of the effect of intergroup cooperation on intergroup attraction. Journal of Conflict Resolution 22:429-439.
- Zajonc, R. B.
1965 Social facilitation. Science 149:269-274.

DISCUSSION: CONFLICT AND STRESS IN THE SPACE STATION

Oscar Grusky

The primary question both Michener and Cook's papers are concerned with is the impact of social factors on the performance of human groups in manned vehicles in space missions of long duration. My comments first address some issues raised by Michener. Then I turn to selected portions of Cook's paper. Finally, I make some general observations and conclude by arguing that a need exists for a systematic data base on social system processes based on past long duration space flights.

COMMENTS ON MICHENER'S PAPER

The theme of Michener's paper is the impact of social system factors on the management of system conflict. He asserts that as crew size increases, crew composition becomes more varied, system goals become more complex and diverse and onboard artificial intelligence and computerization increases, the likelihood of control and conflict problems will become greater because of social system failures. Such failures come about in part because of the greater complexity, differentiation, and decentralization that is created by changes in crew size and composition, technology, and goals.

As Michener notes, the space station crews confront a perilous environment and one that they must deal with largely on their own (Michener does not mention but is no doubt cognizant of the fact that under the current design there is no way a crew member can return to earth in an emergency, since there is no escape vehicle), relative isolation, and a long period of time in space, i.e. 90 days. The SSOC system will have to deal with very complex supervisory control problems, the risk of intergroup conflict, and the necessity of coping with serious crises.

One contributing cause of conflict in the space station, according to Michener's analysis, is modularity. Modularity refers to a social system composed of multiple and distinct subgroups. Michener claims that modular systems may be particularly vulnerable to conflict, miscoordination, lack of cooperation, and mistrust.

Because space station must function in such a perilous environment it is vulnerable to the risk of exogenous shock, that is, an uncontrollable event. More complex social systems are presumably less capable than less complex social systems of coping with such exogenous shocks.

Naturally, one of the first questions one asks when confronting a theory such as this is how does it square with the available evidence, realizing that most evidence is earth-based and therefore only partially relevant or relevant to an indeterminate degree. There is, however, additional evidence, also of limited or of indeterminate value, that may be obtained from observations on group performance in long duration space flights that have already been undertaken. Unquestionably, there is a need for additional systematic research on the problems of group conflict and performance in the space station.

Effects of Group Size and Complexity

Michener suggests that SSOC is likely to experience greater conflict than earlier missions, in part, because the social system will be larger and more complex. However, the evidence on the effects of group size is not uniform. There is no doubt that as group size increases the potential number of intragroup relations multiplies. But the effects of group size on factors such as member commitment, cooperation, and group performance are unclear. For example, Michener suggests that larger-sized groups have weaker member commitment than smaller-sized ones. However, Doll and Gundersen (1969) studied Antarctic groups and found that perceptions of members of compatibility were more favorable in larger (size 20 to 30 members) than smaller groups (size 8 to 11). More recently, Yamagishi (1986) studied laboratory-created groups of size 2, 6, 11, 51 and 501 in order to study social dilemma or public good situations.¹ Subjects were told they would be given \$5 and would be asked to give any number (0 to 5) of one dollar bills to others in the group (group condition) or to matched participants (give-away condition). Yamagishi found that once group size exceeded ten, additional size increases had no effect on member contribution level. The point is that earth-based studies of the effects of group size on both utilitarian and affiliative type group goals have not produced uniform findings.

Conflict

Michener is not specific about the causes of conflict, but the close interactional situation in the space station provides the potential for seemingly minor events to stimulate interpersonal hostility. For example, there is waiting to gain access to the toilet. It takes longer to use zero-g facilities and space constraints will mean a limited number of toilets. Hence, long waits especially when waking up may be common, and questions of priority may produce conflict.

As Michener points out, since no large social system has been established in space, there are no alternatives but to develop generalizations based on earth-based groups. However, potentially suggestive information may also come from extrapolating from observations on smaller-sized groups that have been in space, such as Skylab. Pogue, (1985) a Skylab astronaut who spent 84 days in space,

has written a book describing some of his experiences. Two others accompanied Pogue on his long mission, the Commander (Carr) and the Scientist Pilot (Gibson). Pogue reported that overall the crew got along well together and that they had so many equipment problems that they "had to help each other often. We had good team spirit." Very little space in his book was devoted to the subject of interpersonal conflict. In response to a question on fights and arguments among the crew, he observed: "We didn't have any fights, and there was only one argument that I can recall. It had to do with a change in procedure, and the instructions were very vague. We resolved this by trying the procedure to see if it worked. We never got truly angry at each other, but we were frequently upset with or had disagreements with some people in Mission control. We were all trying hard to get a job done, so there was probably fault on both sides at one time or another" (Pogue, 1985:67). These comments suggest first, that the crew was reasonably well-integrated and supportive of one another, and second, that group identification was to some extent strengthened as a product of antagonism toward Mission Control.

Pogue recounts only one incident that took place between him and the Scientist Pilot that could be characterized as a disagreement: "I think I upset Ed Gibson one day by putting his ice cream in the food warmer and leaving his steak in the freezer. I really felt badly about it. He couldn't eat the steak because it was still frozen hard, and the ice cream had turned to milk. He had to dig out some contingency food to eat. There wasn't too much conversation at dinner that night. He salvaged the ice cream by refreezing it. In liquid form it had turned into a big hollow ball. The next day, after it refroze, he stuffed it full of freeze-dried strawberries and had the first strawberry sundae in space" (Pogue, 1985:67). Attention to these comments is called mainly to suggest the need for collecting and analyzing systematically social system data on conflict already collected from long duration missions, such as Skylab and Salyut. Analysis of these data may help in identifying potential social system sources of conflict on space station. The Russians have had a small space station, Salyut 7, in orbit for almost five years and have manned that station periodically since then. About a year ago they launched the first element of a modular station which is designed to be permanently manned. Bluth (1984) has reported on Soviet evidence of strong interpersonal hostilities among the crew on the Salyut missions.

Mission Length, Conflict, and Expectations

As Michener notes, the projected length of space station missions is 90 days. Another reason for carefully examining group factors in the Carr-Gibson-Pogue Skylab mission is that its length was 84 days or almost the same as the proposed space station missions. As noted above, the Russians also have completed long duration missions that are of interest. Their experience with such missions exceeds ours. Michener argues that the long length of the space station missions may lower crew tolerance and encourage greater subgroup conflict. However,

Pogue (1985) and the *New Yorker* report (Cooper, 1976)² both note that mission length did not constitute a problem on Skylab in the sense of elevating interpersonal tension. The crew apparently did get very disturbed when Mission Control proposed near the end of the mission the idea of lengthening it. It appeared that two factors contributed to the crew's strong dissatisfaction with this idea. First, the crew was trained and geared from the start for an 84 day mission. They had prepared themselves both mentally and physically with this period of time in mind. Hence, a proposed change in the schedule greatly upset their expectations and was dissonance-arousing. Second, and relatedly, the very fact that it was raised as an issue by Mission Control at the crucial point in the mission may have seriously undermined the crew's sense of personal control over their actions. These were very proud and extremely capable individuals with a strong sense of personal autonomy. The tight daily scheduling of their actions and the close observation and monitoring of even minute aspects of their behavior over a long period of time may have been threatening and stress-arousing to these competent and autonomous individuals. The composition of these groups and their training was designed to mute social system conflict, promote strong group integration, and strengthen identification with the group goal or mission. However, such strong identification with the group goal may inhibit externally-induced changes in the mission. Presumably any such changes, if group resistance is to be avoided, must involve a participative process worked out in advance.

In addition to this issue of whether or not greater size and differentiation actually contribute to greater group conflict, there is the matter of the consequences of such conflict for group functioning, and in particular, productivity. Michener assumes that conflict will increase with greater differentiation and complexity, and furthermore, that conflict in general is detrimental to group performance. There is not a great deal of evidence on this, and what exists, is earth-based. Michener, like most students of conflict, sees conflict as creating disequilibrium in the system. Conflict may cause a "breakdown in decision making" (March and Simon, 1958), that is, it is a malfunction and is negatively valued. However, other social scientists, such as Coser (1966) look at certain kinds of conflict as a source of equilibrium and stability. Coser argues that a multiplicity of small conflicts internal to a group may breed solidarity provided that the conflicts do not divide the group along the same axis, because the conflict coalitions provide a place for exchanging dissenting opinions. In essence, he claims that some conflict or disagreement is inevitable and that it is better to foster minor conflicts of interest and thereby gradually adjust the system, than to allow for the accumulation of many latent deep antagonisms that could completely disrupt it. Coser notes that frequent small conflicts keep antagonists informed of each other's position and strength and hence prevent a serious miscalculation on the part of either party. In a similar vein, Lipset et al. (1956) in a study of the International Typographer's Union showed how institutionally-regulated conflict between the two political parties in the union actually fostered a democratic climate

and organizational stability. Likewise conflict between modules may take the form of healthy competition and this may enhance overall productivity. Thus, conflict and competition are not inherently dysfunctional as Michener suggests. Contrary to Michener's approach, this perspective suggests that a key issue is not merely how much conflict takes place, but the conditions under which conflict occurs, for example, the extent to which it is normatively regulated and controlled.

Crew Rotation

Michener mentions only briefly that rotating crews under extended duration space flight may effect their functioning. The effects of rotation, succession, or turnover, merits more detailed treatment. There is a substantial literature on this topic concerning the effects of rate of succession on group and managerial effectiveness (e.g. See Grusky, 1963, 1964; Brown, 1982). Practical research questions include: optimal mission length, optimal method of crew rotation (replace individuals, subgroups, or total crews), optimal method of leader rotation, etc.

COMMENTS ON COOK'S PAPER

Cook's paper is concerned first of all with stress and the relationship between stress and productivity. In addition, she examines the issue of mediated communication, particularly computer-mediated communication and its effects on productivity.

Stress

Cook proposes that space be used as a site for basic research on stress. She points out that reliable and valid non-physical health-related measures of stress are lacking. Space station is a good site for stress research, she claims, because there are so many stressors in space, such as crowding, noise, workload, and life-threatening crises. She describes a model of stress produced by interpersonal factors such as inequitable assignment of rewards, task or role ambiguity, arbitrary exercise of authority, and others. Cook wishes to complement physiological and psychological stress research by investigating social system properties of stress, an approach that has not been heavily utilized in the past. She also wants to explore adaptive group strategies for coping with stress. She proposes the intriguing idea of developing a computer-aided system to rectify cognitive processing deficiencies that appear under high stress levels. However, one of the problems with stress measurement is that so many factors can be stressful that objective quantitative measurement is difficult.

Gender and Stress

Although up to the present space missions have been male-dominated, it is evident that future missions will involve more female astronauts and mission specialists. Research in psychiatric epidemiology has consistently documented an association between gender and psychological distress. Women are more than twice as likely as men to report affective disorders and extreme levels of distress (Al-Issa, 1982; Kessler and McRae, 1981). Although male prevalence of some psychiatric disorders is greater than females and for some disorders there is no reported association with gender, the best available evidence indicates that the psychological well-being of women is different than that of men.

The major sociological interpretation of this evidence is that women's roles expose them to greater stress than men's (Gove, 1978). Gove (1972) has claimed that female role stress is especially pronounced in traditional role situations.

A number of investigators have shown that women are more vulnerable than men to a range of what have been called network events, that is life crises that are significant to the lives of persons important to the respondent (Kessler, 1979; Radloff and Rae, 1981). Kessler has proposed that women care more about people, and because this is the case, they are more vulnerable to crises that take place "at the edges of their caring networks." (Kessler, 1985). Men are emotionally affected by crises that occur within their nuclear family, but women are more deeply affected by both crises among members of their nuclear family and among persons who may be classified as friends and associates.

There are a number of major limitations in the analysis presented above:

1. The findings showing a relationship between gender and psychiatric distress and subclinical distress can be explained by selection factors.
2. Most of the evidence on role-related stresses has been based on scales using subjective evaluations.
3. Evidence on the differences suggested between men and women claiming that the latter are more vulnerable to crises in their networks is sparse.

Despite these limitations of which Cook is well aware, this information and the speculations described above raise some potentially important issues regarding long duration space missions. Specifically, one issue is whether or not male and female crew members will take on different roles and respond differently to crises that may take place in the space station. Kanter (1977) has studied the lone woman in male-dominated work organizations as part of her study of what she calls "skewed sex ratios." She has distinguished between dominants and tokens in these organizations and suggests that (1) tokens are more

visible than dominants (2) differences between dominants and tokens tend to be polarized and (3) tokens' attributes tend to be "distorted to fit pre-existing generalizations about their social type." At issue is the effect (if any) of the gender distribution in the group on command, control, and communication processes. It may very well be the case that selection factors that have up to now worked well in identifying crew members capable of handling stress will continue to work effectively in the future. It is also reasonable to anticipate that besides selection effects, situational effects will be overpowering and hence the gender differences suggested above will be masked. Alternatively, it may be that the larger-sized groups in space station 1990s combined with the existence of a "skewed sex ratio" (Kanter, 1977) in work groups will have problematic impacts on group functioning. Research is needed to explore these and related issues.

Computer-Mediated Communication

As Cook has observed, the social consequences for systems of long duration where the primary communications are computer-mediated are simply unknown. As Cook notes, the recent finding by Siegel et al. (1986) that computer-mediated communication facilitates the upward flow of negative communications or information that challenges those in high status positions merits replication. This problem also should be studied developmentally to see if changes occur as groups exist over long periods of time. Another related problem that merits study is the potential impact of cultural differences on computer mediated communication. People of different cultural backgrounds may respond in radically different ways. Such differences if found could be consequential to communication between the various space station modules, the Japanese, European, and that of the United States.

Cook cites Connors (1985:32) research as justification for the proposition that "computer mediation may mitigate the inhibiting effects of face-to-face communication when "subordinates" have access to critical information and may need to challenge authority." However, Connors' groups bear little resemblance to the environment experienced by past long duration space flights or presumably will be faced by future flight crews, such as weightlessness, continuous peril and public exposure, continuous high task-load, small amounts of space per person, etc. Moreover, the idea of challenging authority and attitudes toward work are culture-bound. Hence, even if the findings were applicable to the United States' space module, they would not necessarily be as applicable to the Japanese or European modules.

CONCLUDING COMMENTS

Most of the comments in the two papers focus on negative effects such as conflict, stress and miscommunication. Michener stresses the perilous environment, the possibility of conflict between modules, human error possibilities, and breakdown possibilities that stem from

the sophistication of the technology. Michener neglects the potential positive contributions of small conflicts and competition to group functioning providing that such conflict and competition is institutionalized and is expressed in legitimate ways. Cook focuses on the problems of decisional and interpersonal stress. Yet, in contrast, what was highlighted in the narrative reports, both by Pogue (1985) and The New Yorker (Cooper, 1976) accounts, was the relatively smoothness of interpersonal relations among the Skylab crews, their high motivation, high productivity, high goal identification, and group commitment. Perhaps these reports have been "sanitized". In any case, it is clear that the social system impacts with respect to conflict and stress are unknown, although we do know that these factors can have consequential effects, and as Michener suggests, increasing social system complexity may enhance the likelihood of social system problems.

Both the Cook and the Michener papers stress the importance of social organizational factors on productivity or performance. Cook calls attention to the work of Foushee (1984) who has used flight simulators to study group process. Foushee cites a study by Ruffell Smith (1979) who had B-747 crews fly a simulated flight from New York to London. A failed engine, hydraulic system failure, poor weather, and other problems created an emergency situation. Foushee observes significantly that "Perhaps the most salient aspects of this flight simulation study was the finding that the majority of problems were related to breakdowns in crew coordination, not to a lack of technical knowledge and skill." Research on social factors affecting group conflict, stress and other related issues as both Michener and Cook have observed, is essential.

In summary, four major observations were made on Michener's paper, as follows:

1. Findings from earth-based laboratory and field research on the effects of group size and complexity on task performance have produced inconsistent results.
2. Michener's social system theory suggests considerable potential for group conflict on the space station. However, narrative accounts describing a Skylab mission do not conform to this theory.
3. Michener's theory assumes that conflict and competition (seen as a form of conflict) has only deleterious consequences for social systems and this may not be the case when conflict is institutionally regulated.
4. Michener omits extensive discussion of problems associated with crew rotation.

The following observations were made on Cook's paper:

1. The development of new methods of measuring stress and coping techniques are needed. Existing data on Skylab crew behavior should

be examined in order to identify effective interpersonal coping strategies, that is, techniques that crew members have used that reduced, controlled, or made stress more tolerable.

2. Research is needed to explore systematically the relationship between gender, group structure, and stress.
3. Research is needed on the social impacts of computer-mediated communication. We need to know the positive and negative consequences of computer-mediated communication for individual and group decision processes. Cook has identified a set of hypotheses that merit intensive study.

Some of these problems can best be studied by means of human simulations where the space station situation is simulated in the laboratory by means of a mock-up and human crews of eight to ten or even twenty volunteer subjects are studied continuously in the laboratory for long periods of time. The crews would be given carefully assigned tasks as similar as possible to those to be performed by space station crews. The noise level is manipulated as are living conditions to approximate as closely as possible the real situation. Ideally, one would develop a set of experimental studies using the simulation method which would enable the close study of the effects of key independent variables such as authority structure on crew productivity, performance, and satisfaction. The same technique could be employed to examine the effects of various methods of crew and leader rotation.

A fundamental research recommendation should be added to those noted by Michener and Cook, namely the need for development of a systematic data base in the area of group performance of past (and future) astronauts in long duration space missions. Such a data base is especially needed because the space station is a unique environment due to the interaction of a very unusual set of characteristics such as weightlessness, constant danger, restricted or computer-mediated communications, high stress due to noise, and other environmental hazards. Valuable although limited information can be obtained from studies of social systems facing quite different but presumably comparable situations such as polar environments and long duration submarine missions. Hence, there is a great need for data on this particular type of social system that is unique to long duration space missions. The types of data that should be included in such a data base are demographic information on the astronauts, performance data, and perhaps most important of all, audio and videotapes of missions, such as the three-person Skylab missions discussed above. Research access to these tapes would facilitate development of new measures of stress and conflict and their relationship to decision processes and would permit study of microgroup processes such as initiation of interaction, rates of interaction, and measures of power (such as interruptions, talkovers, and overlaps, etc.). The hope is that NASA might be convinced that a data base of this kind would be a valuable research resource for them and that such a data base could be assembled

and the data analyzed in such a manner as to conceal appropriately as necessary the identities of particular astronauts and their specific missions.

NOTES

1. Obviously Yamagishi did not create actual groups with 501 members in the laboratory. Instead, he allowed no communication or contact between subjects, who were isolated from one another, and told them the number of persons in their "group". No data were presented on the validity of this manipulation.
2. The New Yorker account also suggested that the three-person Skylab crews varied substantially in their productivity. One major determinant of this variation was how much was demanded of them by Mission Control. When a point was reached that seemed to the members of the crew to overtax their capacity, they complained and Mission Control reduced the workload.

REFERENCES

- Al-Issa, I.
1982 Gender and adult psychopathology. Pp. 54-84 in I. Al-Issa, ed., Gender and Psychopathology. New York: Academic Press
- Bluth, B. J.
1984 Pilots of outer space. Society 1/2:31-36
- Brown, M. C.
1982 Administrative succession and organizational performance: the succession effect. Administrative Science Quarterly 27:1-16.
- Connors, M. M., Harrison, A. A., and Atkins, F. R., eds.
1985 Living aloft: Human Requirements for Extended Spaceflight. NASA
- Cooper, H. F. S.
1976 A Reporter at large. New Yorker 9/10:34-70
- Coser, L.
1966 The Functions of Social Conflict. New York: The Free Press
- Doll, R. E. and Gunderson, E. K. E.
1969 Hobby interests and leisure activity behavior among station members in antarctica. U.S. MNMRU. Rep. No. 69-81

- Foushee, C.
1984 Dyads and triads at 35,000 feet: factors affecting group process and aircrew performance. American Psychologist 39:885-893.
- Gove, W. R.
1972 The relationship between the sex roles, marital status, and mental illness. Social Forces 51:34-44
1978 Sex differences in mental illness among adult men and women: an evaluation of four questions raised regarding evidence on the higher rates of women. Social Science and Medicine 12B:187-98
- Grusky, C.
1963 Managerial succession and organizational effectiveness. The American Journal of Sociology 59:21-31.
1964 The effects of succession: a comparative study of military and business organization, in M. Janowitz, ed., The New Military. New York: Russell Sage Foundation.
- Kanter, R.
1977 Men and Women of the Corporation. New York: Basic Books.
- Kessler, R. C.
1979 Stress, social status, and psychological distress. Journal of Health and Social Behavior 20:259-72.
1985 Dyadic Responses to Stress: A Study of Married Couples. Unpublished manuscript.
- Kessler, R. C. and McRae, J. A. Jr.
1981 Trends in the relationship between sex and psychological distress: 1957-1976. American Sociological Review 46:443-452.
- Lipset, S. M., Trow, M. A., Coleman, J. S.
1956 Union Democracy. New York: The Free Press.
- March, J. G., Simon, Herbert, A.
1958 Organizations. New York: John Wiley.
- Pogue, W. R.
1985 How Do You Go to the Bathroom in Space? New York: Tom Doherty Associates.
- Radloff, L. S., Rae, D. S.
1981 Components of the sex difference in depression. Pp. 77-110 in Roberta G. Simmons, ed., Research in Community and Mental Health. Greenwich, CT: JAI Press.

- Siegel, J., Dubrovsky, V., Kiesler, S., and McGuire, T. W.
1986 Group processes in computer-mediated communication.
Organizational Behavior and Human Performance 37:157-187.
- Yamagishi, T.
1986 Limits of generalization of small group social dilemma
experiments. Pp. 1-17 in American Sociological
Association. New York.

SYNOPSIS OF GENERAL AUDIENCE DISCUSSION

Due to time limitations only a brief discussion of the papers in this session was possible. Two main comments on these papers were made, as follows:

1. It was pointed out that neither of the papers considered the relationship between the airborne or space station crew and the larger community that participates in the operation of the station. The role of mission control, for example, was not mentioned and merits careful examination. The airborne crew does not exist in isolation and reflects the objectives of the larger organization and of the nation (or nations) as a whole. Mission control is in constant communication contact with the airborne crew and serves important functions with regard to its safe, effective, and efficient operation.
2. It was suggested that the extant literature on social system behavior in a number of other analogous "hostile" environments such as undersea or in Antarctica be reviewed carefully for information that might be relevant to the situation of long-duration space missions.

SESSION VII
THE HUMAN ROLE IN SPACE SYSTEMS

THE ROLES OF HUMANS AND MACHINES IN SPACE

David L. Akin

INTRODUCTION

Throughout the history of the space program, there has been a dichotomy of opinions on the relative importance of manned and unmanned (i.e., robotic) applications. Until the arrival of the shuttle, manned and unmanned operations occupied different sections of NASA Headquarters, involved different groups of NASA field centers, and were generally viewed as competing for the limited funds available. There were (and still are) areas, such as planetary exploration, where there were no viable options to the use of unmanned systems. The arguments, rather, tended to the utility of humans in space, and the cost of replacing each of their functions with robotic alternatives.

Any self-contained device performing a useful function in space, whether a human or a robot, must contain the same set of basic functions to adequately perform the mission. In many cases, of course, the mission is actually constrained to work around the limitation of the state-of-the-art in one or more of these areas. These basic functions for autonomy include:

- Sensation In order to operate on the local environment, a system requires sensors for detecting objects. These typically break down into remote sensors (such as vision or other ranging systems) and proximal (such as tactile and force sensors).

- Computation Having the capability to detect objects does not translate directly into the capability for manipulation. Understanding the spatial relationships, having a knowledge base of both general activities (tools, forces and motions) as well as specific knowledge (specific satellite design details) are necessary for effecting a complete system.

- Manipulation This area has trailed the others considerably, as many of the original space objectives did not involve manipulative activities. Manipulation to date has been

performed by the sampling arms of the Surveyor and Viking landing spacecraft in small scale, and by the Remote Manipulator System of the shuttle in larger scale. None of these systems has involved any appreciable dexterity in either the arms or the end effectors. Nonetheless, this area is pivotal for future space activities, as it relates to the capability of the spacecraft system to interact with, and to alter, its local environment.

Locomotion This is a necessary function, often relegated to a supporting role. The capability to maneuver around in space, either on an existing structure or in free space, is required for any robotic system to be generally useful. It might be anticipated that space systems will evolve a wider range of locomotive capabilities than humans have evolved in a gravity field. For example, legs on a human provide both locomotion and anchoring functions. In the microgravity environment of space, locomotion might well be relegated to the equivalent of arms, which have the finer dexterity and force control required in the absence of damping, and anchoring left to sets of specialized manipulators with strength, but little other capability. Thrusters for free-flight propulsion will also be common, at least for those situations not constrained to minimize use of consumables.

Support This category includes all the other functions necessary for the system to exist. This would include power, cooling, structural integration, navigation, and communications.

It is interesting to examine a known autonomous system (a human) in the context of these functions. The head is the sensor platform, located in the optimal location for bipedal locomotion. The computational system (brain) is co-located with the sensors in the head, to minimize the length (and vulnerability) of the high-bandwidth data paths, particularly vision. The arms form a dexterous manipulative system, and the legs similarly perform locomotion tasks. The torso thus encompasses most of the support functions, as well as tying all of the other systems together in a self-contained unit. The human body is thus a wonderful example of a possible design for a robot. However, the human paradigm should not be extended too far, as many of the optimal choices for a system which stands erect in a gravity field may have little logical application in a system optimized for weightlessness.

The task, therefore, is to come to an understanding of the past and present roles of humans and machines in space activities, and extrapolate to the future to come to a meaningful understanding of the capabilities and limitations of each. In fact, it is worth emphasizing at this point an essential conclusion of this paper: it is not an

"either-or" choice between humans and machines. There are necessary and sufficient roles for both in the foreseeable future in space.

HISTORICAL PERSPECTIVES FROM SPACE FLIGHT

With the limited payload capability of early launch systems, there was no viable alternative to the use of unmanned satellites. These early payloads were composed of sensor packages, communications gear, and support systems, and were required to do nothing more than observe/measure and report their findings. Even today, many of the satellites being launched to orbit are still limited to these functions; for the purposes of this paper, these systems may be considered to be subrobotic systems.

It seems clear that the original intention of the Mercury program was to use the humans as an experimental subject, in order to study the effects of spaceflight on humans. The choice of experienced military test pilots for Mercury astronauts led to some predictable dissatisfaction with this role, and the desire for incorporation of manual control capability in the vehicle. This led to the use of the human as a controller, albeit primarily in the backup mode. A case in point was the Mercury attitude control system. The primary system was an automatic one, which maintained the capsule in retrofire attitude during orbital operations. A second selection was "fly-by-wire", in which the astronaut could command attitude maneuvers by use of a side-stick controller, which would then be performed by the attitude control system. The final mode, however, was purely manual, as the astronaut actuated push-pull rods which mechanically opened and closed thruster valves.

This issue of humans as the final back-up is a pivotal one. For example, Mercury was a simple spacecraft, designed primarily for a single, sequentially organized mission. It carried no on-board computer, but instead relied on activating systems at set times on a mission clock. Contingencies, such as the decision to enter without jettisoning the retropack on John Glenn's orbital mission, relied on manual activation of retrofire commands to prevent the sequencer from automatically separating the retropack following retrofire. Thus, throughout the Mercury program, the human represented the adaptable (reconfigurable) element of the Mercury control system.

The Gemini program was an interesting "backwater" of space flight development. Originally conceived as a Mark II version of the Mercury capsule, Gemini was developed as an interim program to increase space flight experience while waiting for the development of the Apollo system. Since it represented to some an evolutionary dead-end in manned space flight, the manned elements were permitted to have unusual sway in the systems development. Thus, where Mercury was largely automatic, Gemini was almost entirely manual. It might indeed be argued that, more so than any other space program before or since (including Shuttle), Gemini was a pilot's spacecraft. There were no automatic abort modes: the crew had to decide the appropriate action based on the reports of the instruments. For the first time, a space

vehicle could be accurately described as a spacecraft, since Gemini had the capability to change orbits and achieve rendezvous. The crew had windows which faced forward, and hatches which could be opened and closed again in flight. Even in landing, the vehicle was positioned to allow the crew to sit upright, and much development effort took place towards a Rogallo-wing recovery system which would have allowed Gemini to maneuver to a landing on the dry lake bed at Edwards Air Force Base.

Even in the midst of this manual spacecraft, additional elements of automation had to be incorporated. The Gemini was the first spacecraft to fly with an on-board computer, used for calculating rendezvous maneuvers and for control of the lifting reentry. Although many of the procedures used for rendezvous and docking were manual in nature, the complexities of orbital mechanics required the use of ground or on-board computer calculations; the crew were primarily used as interpreters of visual and radar data.

The presence of humans on board Apollo may be considered as entirely a political decision, as the entire objective of the Apollo program was to place a man on the moon and safely return him to earth. The greater complexities of the spacecraft and mission led to a return to automated systems, after the largely manual nature of the Gemini spacecraft. Thus, for example, many of the abort modes were automatically initiated, although the crew did agitate for manual control of launch vehicle trajectory as a backup for the Saturn flight control system. The manual docking techniques developed during Gemini were utilized by Apollo in lunar orbit.

Apollo again showed the utility of humans as a robust backup system. It was not possible to do a survey of landing sites down to the level of all possible hazards to the Lunar Module; it was therefore planned that the pilot would take over and steer the lunar lander to a safe landing site. This system worked well in every instance: the initial aim point for Apollo 11, for example, turned out to be right in the middle of a boulder field. Manual control of the landing vehicle allowed the targeting of landings next to an unmanned Surveyor spacecraft, adjacent to a deep lunar rille, and in the lunar highlands. This greatly augmented the data return, as later flights were targeted into areas of greater geological interest, with fewer options for safe landing sites.

The presence of humans to pilot the landers into safe locations may be compared to the Viking landings on Mars a few years later: since the unmanned vehicles did not have the image processing and decision making capabilities of humans, both of the landing craft had to be targeted to the flattest, smoothest, and therefore least interesting landing sites available. Similarly, the Soviet Union performed lunar exploration with unmanned vehicles. However, the quantity of samples returned differed from Apollo by 3-4 orders of magnitude; since the samples were selected randomly from the immediate location of the landing vehicle, it may be assumed that the quality of samples varied widely from Apollo as well.

Skylab, as the first American space station, involved the long-term habitation of space by humans. Indeed, one of the major objectives of Skylab was to study the effects of long-term space flight on human

physiology; however, to use this objective as a justification for manned space flight constitutes circular logic. Much more may instead be said of the other science objectives of Skylab, such as earth resources, solar physics, and space operations. In all of these, the Skylab crews played an essential role in the success of the mission.

Since Skylab was constructed of surplus Apollo components, there was little significant difference between the two programs in the automation levels of the vehicle system themselves. The only significant difference was in the experiment packages, which in Skylab represented a later generation of technology from the spacecraft hardware. For example, the solar observing instruments in the Apollo Telescope Mount could be (and were) operated remotely from the ground. However, the onboard crewmen could provide more immediate decisions when faced with fast-breaking phenomena, and in fact managed to record solar flares from their inception. Modifications to the onboard control panel of these instruments during the course of the Skylab mission were primarily to increase the ability of the crew to make immediate data records for use onboard, by the addition of an instant-print scope camera.

Of greatest significance, perhaps, was the role played by the crew in the repair of the workshop and salvation of the mission. Extensive extravehicular activities (EVAs) were performed to free the jammed solar array, and to deploy a sunshade to reduce temperatures in the workshop to habitable levels. The three Skylab crews regularly repaired failed equipment, both inside and outside of the space station, and clearly made possible the success of the program: had Skylab been an unmanned station with the state-of-the-art robotics of its time, it clearly would have had little or no recourse beyond those capabilities left by the launch accident.

The greater complexity of the Space Shuttle has led to the greatest amount of automation yet. Flight crews have referred to the Orbiter as the "electric airplane", since almost all functions are controlled through the four general-purpose computers (GPCs). The atmospheric flight characteristics of the Orbiter are such as to be practically unflyable without stability augmentation. Although a manual direct mode does exist, few of the flight crew have much success in this mode in training simulations, and even this mode relies on the GPCs to interpret hand controller data and command motions of the flight control surfaces. Although the flight control system is capable of flying the vehicle all the way through landing ("autoland"), it is interesting to note that no crew has yet allowed this to be tested on their mission: the commander always takes over in control stick steering mode (i.e., stability augmented) at subsonic transition, or certainly by the pre-flare maneuver at 2000 feet altitude. This is representative of many of the lessons learned from shuttle operations: the flight crew have now been cast in the role of systems managers, but still demand active involvement in all safety-critical aspects of the mission. It would be unwise to assume that this trend will not continue into the era of the space station.

CAPABILITIES AND LIMITATIONS

It has been said that humans are the only self-programming, highly dexterous autonomous devices capable of being mass-produced by unskilled labor. Be that as it may, there are significant limitations on both humans and machines in the space environment. Having evolved in the environment of the earth's surface, it is necessary to (in some degree) take the conditions of earth along with humans in space. Constraints to be considered include atmosphere, consumables, volume, work cycles, and gravity.

Humans need oxygen above a partial pressure of approximately 3 psi in order to survive. Through the Apollo program, spacecraft were supplied with a pure oxygen atmosphere at 4 psi. This simplified several operational problems: the structures could be simpler, as the internal pressures were less; only a single gas had to be stored and delivered; and there was no requirement for denitrogenification prior to an extravehicular activity. However, the Apollo 1 fire showed graphically the primary disadvantage of a single-gas system.

In Skylab, the atmosphere was kept at 5 psi, with nitrogen forming the additional partial pressure beyond that required for oxygen. While this reduced the flame propagation problem, the crew was less than satisfied with the atmosphere, as it was difficult to carry on conversations beyond their immediate vicinity. Current plans for the Space Station assume a sea-level pressure of 14.7 psi, as used on the Orbiter. This decision is coupled into the choice of avionics: the sea-level pressure of the Orbiter was partially chosen to allow the use of "off-the-shelf" air-cooled avionics. This had an effect on habitability, as the number of cooling fans on the Orbiter creates an appreciable amount of noise, thus limiting conversations to the immediate vicinity of the individuals. The Orbiter has been operated extensively at 10.2 psi during pre-breathe cycles prior to an EVA, but this requires a significant power-down of avionics to prevent overheating.

A biological organism, such as a human, is powered by a series of chemical reactions, and must be replenished regularly. In a totally open-loop system (that is, no attempt at recycling anything), humans will require approximately 5 kg/day of food, water, and oxygen. Recycling water and air will reduce this to 1 kg/person-day: this is equivalent to 540 kg of consumables for a six-person crew over a 90 day resupply cycle. Even without recycling, then, consumables are not a pacing item for a space station if the crew sizes are kept small. These figures also do not take into account such operational factors as air loss, inefficiencies in recycling, or food carried for reasons beyond base-level nutrition, and therefore the actual figures planned for consumables in space stations will be higher than these academic minimums. Many of the techniques for effective recycling are currently highly experimental, and will require a great deal of development prior to operational use.

Studies have shown a direct relationship between habitable volume and crew performance; the minimum volume is also a function of mission duration. In addition to the working volume, humans need to have

shared facilities for eating, exercising, and personal hygiene, and are usually best provided with some private locations for recreation and sleep. Deciding on these issues are some of the most difficult choices in interior station design, as there is often no clear relationship between productivity and volume; indeed, there is often no generally agreed-upon metric for productivity itself. Other desirable modifications to a spacecraft designed for long-term human occupancy include windows (as many and as large as the structural designers can be forced to incorporate), airlocks, and redundant escape paths in case of contingencies such as hull penetration or fire.

Humans are not capable of working "around the clock": some amount of recreation is required, along with natural housekeeping and other support functions and a sufficient amount of sleep. A normal 40 hour week represents a 24% duty cycle for a human. Assuming five hours per day for meals, housekeeping, and exercise represents a further 21% of the time, leaving 55% of the day for sleep, recreation, and general off-duty activities. This may be compared to the averages for Skylab: 25.6% experiment operation (work), 33.9% meals, housekeeping, and exercise, and 40.5% for sleep, rest, and other. It is interesting that the net percentage of time spent on experiments is so close to that of a typical 40 hour week; the exhaustive pace reported by the Skylab crews clearly demonstrates the increased overhead associated with living in space. Evidence indicates that the work pace established in Skylab would be difficult to maintain over indefinite periods on a space station: therefore, planners must either accept lower than normal duty cycles on experiments and other output-oriented activities, or plan ways of automating the housekeeping functions to bring these back in line (from a perspective of time) with comparable activities associated with living on earth.

One of the origins of the increased housekeeping times is the necessity of adapting to routine living in the weightless environment. Although it can certainly be maintained that if sufficient experience has yet been obtained to provide definitive conclusions in this area, clearly it will be difficult to overcome the millions of years of evolution in a gravity field in a brief time, and some performance degradation in weightlessness is to be expected in the foreseeable future. Physiological reactions to extended microgravity include a number of hormonal and fluid shifts: the only long-term effect which seems to be both serious and progressively degenerative is a decalcification of bone material. This effect can be retarded to some degree by strenuous exercise, particularly involving compression of the large bones of the leg: this has led to the development of treadmills with elastic cords replacing some of the force of gravity, allowing aerobic running exercises.

Some effort has gone into examining the options for providing appreciable gravity on a space station, by rotating the components to provide a centripetal acceleration. This effect can be quantified as

$$g = \omega^2 r$$

where w is the angular velocity, and g is the effective acceleration at a radius of r . Early plans (prior to Skylab) indicated that an angular velocity of 4 rpm would be acceptable, producing a required radius of 55.8 m for earth-normal gravity. Some research has suggested that 3 rpm (99.3 m) might be a better rotational velocity for human adaptation, even with a select crew population. If selection standards are relaxed to most of the general population, that implies a rotation speed of 1 rpm, with a resultant radius of 894 m required.

Obviously, it would be extremely complex and expensive to provide stations of this size. One method of easing this requirement would be to provide partial gravity: an early space station proposed with a radius of 25 m at a spin rate of 4 rpm would have produced an apparent gravity of .45 g. However, nothing is known of the effects of partial gravity on bone decalcification or other microgravity effects; this is clearly an important research issue to be addressed by a space station. Short of this information, the logical approach is probably that being considered: do not provide artificial gravity, and rotate the crews at intervals known to be safe, such as three months.

It would be unwise, however, to overly emphasize the limitations of humans, without some equal attention to their assets. The capabilities of humans have been demonstrated repeatedly throughout the history of manned space flight. The list of experiments repaired, satellites retrieved, and missions saved would be too long to go into in this paper. Of greater importance than reviewing the individual performances is to summarize the individual capabilities which made them possible.

Manual dexterity is obviously highly critical for those tasks requiring physical manipulations. No manipulator has yet been developed with anything remotely approaching the dexterity of the human hand. Some experimental efforts in this direction (the Utah/MIT hand and the Salisbury hand) have produced impressive manipulator arm at the current time. The approach taken in the nuclear and the undersea communities (the other two areas for application of general-purpose robotics) have tended towards the use of simple and effectors, and the alteration of tasks to allow for limited dexterity. To some extent, the same is true of space systems designed for EVA involvement: current pressure suit gloves are still far more dexterous than manipulator and effectors, and are likely to continue to evolve in the future.

Strength is (perhaps surprisingly) still an important issue in microgravity. The Remote Manipulator System of the Orbiter is capable of manipulating payloads up to the Orbiter limit of 65,000 lbs., but is severely strength-limited, and therefore handling time goes up as mass goes down. The most capable system for retrieval has been shown to be an EVA astronaut in the Manipulator Foot Restraints, attached to an RMS with its joints locked. This configuration was used for grappling the two HS-376 satellites retrieved on shuttle mission STS 51-A, as well as the Leasat HS-393 satellite captured, repaired, and re-released on STS 51-I. This last procedure especially, with the requirement to despin and capture, and later respin and deploy a massive satellite, could not have been effected without the strength and dexterity of a human.

This raises an interesting side point: in most robotic systems available today, manipulators are specialized for either strength or dexterity, but not both. Those arms used for positioning large masses generally do not have the positioning accuracy of arms used for exact pointing or positioning tasks with lightweight payloads. To some extent, the microgravity environment of space may tend to help this problem, as no appreciable strength of the arm will go to maintaining its position in the absence of external forces. At the same time, mass limitations tend to produce lightweight space manipulator designs, requiring either tasks adapted to their flexibility, or sophisticated compensatory control systems to actively reduce the structural modes.

In general, humans are excellent adaptive control systems. Humans routinely change gains and algorithms based on the physical parameters of the system being controlled, and are capable of adapting and changing to a continuously varying system, within limits. Humans improve with practice, and can transfer learned responses to new control tasks of a similar nature.

Humans are especially suited for rapid processing and integration of visual data. From the first manned orbital flights, crews have reported being able to see features on the ground indistinguishable from the best photographic records. Nuances of color, shading, and pattern may be instantly apparent to a human, yet be below the resolution of an electronic imaging system. Humans have the capability to receive and derive spacial information from both static and dynamic scenes, and continuously update their world model based on visual data.

The human capacity for judgement is certainly well-discussed, but it might be maintained that there is a greater utility for low-level reasoning than for intellectual decision-making capability. For example, neutral buoyancy tests of EVA show a human capacity for instinctive maneuvering in the simulated weightless environment, resulting in improvement in task performance without the need for restraints, and without conscious consideration of body actions. This sort of maneuvering, which is computationally complex for a robot, can be performed by a human in "background" mode while concentrating on task planning. While expert system shells will be important for error diagnosis and strategic planning, it is the robotic equivalent of reflexes, instincts, and common sense which will provide the greatest challenge for the artificial intelligence community.

FUTURE RESEARCH NEEDS

Quantitization

Many of the important decisions on the applications of humans and machines in space have been (and are currently being) based on prejudices from limited prior experience, a priori arguments, and large, costly system analyses which have no meaningful underlying data base. Certainly, the path of following past experience will probably result in an operable space station. However, much could and should be done to formulate and follow a logical plan for ground-based analyses

and simulations, and flight experiments, which would produce a meaningful data base on human and machine capabilities and limitations in each of the operational categories needed for a successful space station program. There are two caveats for such a program: first, of course, the research must be performed. But equally important, the program managers must be willing to listen and act on the outcomes of the research, and not revert to "tried and true" solutions for the sake of engineering conservatism.

Appropriate Roles

One of the outgrowths of the data base development described above would be a greater quantitative understanding of the appropriate roles of humans and machines in space operations, and the most favorable combinations of each to accomplish any particular task. This may imply the altering of traditional roles. For example, as discussed earlier, the flight crew has insisted on maintaining an active, controlling role in those areas critical to safety of flight, or of mission success. However, the (appropriate) risk adversity of mission planners prohibits intuitive solutions to any problem which can be foreseen prior to flight. This has led to the plethora of checklists which describe the appropriate actions of both the flight crew and the ground controllers in any contingency. But, it might be argued, this algorithmic approach obviates the need for most of those capabilities currently unique to humans, such as insight and judgement. Shouldn't this argue for automated systems to implement corrective action in the event of critical malfunctions?

In response to this question, an interesting parallel may be drawn from current findings in aeronautical human factors. With the increased autonomy of transport flight control systems, the airline flight crew are assuming to greater extents than ever the role of system managers. Flight control systems have become capable of completely controlling the aircraft from liftoff through touchdown and rollout. However, serious accidents have already occurred in airline service, due to a flight crew which is neither fully aware of the intricacies of the flight control system, nor highly practiced in manual control of the aircraft. It seems clear that, short of removing the flight deck crew and automating airliners, too much automation breeds overconfidence and inattentiveness in the cockpit; the same will probably be found in space flight.

The conclusion of this argument is to show that it is not enough to fully understand the limitations and capabilities of each of the component technologies: the interactions of the pieces may be far more important to safety and mission success than the pieces themselves. Since the possible number of interactions is a combinatorial problem, it is hopeless to postulate a rigorous or analytical solution to this problem. It is clear, however, that it must be approached in a logical and methodical way if programs as complex as space station are to be successful.

Improved Metrics

A problem which is at once conceptually simple and, in implementation, difficult is that of appropriate metrics for human and machine performance in space. Performance indices based on task performance tend to be unique, or specialized to a small subset of tasks. Indices based on more generic factors, such as motions or subtasks, must take into account the fact that humans and machines may be able to perform the same tasks, but will likely use different techniques in accomplishing them. Even among limited communities, such as EVA, there has yet to form any consensus on the appropriate measurements to produce meaningful comparisons between tasks or experiments. This will be true in larger measure as the field expands to include a wider range of human and robotic activities.

An Assessment of Anthropocentrism

Almost all of the designs currently proposed for telerobotic systems are highly anthropocentric: that is, they tend towards a robotic duplication of the human form. Artist's concepts show a head (sensor platform), with two arms mounted on a torso, and with one or two "legs" used for grappling. This approach is understandable for a system which is designed to incorporate (or at least allow) teleoperation, but its assumption for a fully robotic system can only be attributed to engineering conservatism ("stick with a known configuration"). Some recent results from simulation indicate that a number of manipulators with limited degrees of freedom, designed to perform limited or dedicated tasks, may offer performance increased over two anthropomorphic general-purpose manipulators. The human form, evolved in a gravity field for effective protection from predators, is not necessarily the best adaptation for space activities, and alternate forms and technologies should be encouraged and studied carefully.

CONCLUSION THE (FAR?) FUTURE

Given sufficient time, support, and determination, human beings have demonstrated that they are capable of doing almost any physical or intellectual task. They have shown over the last quarter-century that they are fully capable of living and working in space, performing a wide variety of tasks, from the routine and mundane to innovative, immediate actions needed to save a mission or a life. One may postulate a new unit of measurement: the "human-equivalent", or a system in space with the same effectiveness as a single human. Such a system might be composed of a full-time human, living and working in space; of a human in space working part-time with a robotic system; of a teleoperated system controlled by a human on the ground; or even of a fully autonomous robot with learning and reasoning capabilities.

It is clear that the "human-equivalent" presence in space is on a monotonically-increasing curve. As the societies on earth start to gain advantages from space, the need for capabilities in space will continue to grow. This implies a parallel growth in the requirement to operate routinely in space.

As a thought experiment, let us pick that point in the future at which machine systems have become as capable as a human. It may even be maintained that this point is not in the far distant future: manipulative capabilities are already approaching that of a human in a pressure suit, and human decisions on-orbit have been constricted to algorithmic logic trees easily implemented on modern computers. It is clear that, at some point in time, machines will be capable of performing everything currently done by humans in space. At that point, will we (as a nation, or a civilization) pull all the people out of space, and rely totally on robotic systems to continue the exploration and exploitation of this last, infinite frontier?

At this philosophical question, the author has reached the limits of his original charter. History indicates that humans are capable of performing important, complex tasks in the space environment. As adaptive mechanisms, humans have only begun to learn how to operate in this new environment.

However, much of manned space flight to date has been involved with overcoming the limitations of biological organisms. The evolution of robotic systems has been orders of magnitude more rapid than that of biological systems; there is no reason to assume that this new evolution will stop short of full human capacities, particularly if measured against the currently limited capabilities of humans in space. It is clear that both systems have strengths and weaknesses; that the best mixture of each is a time-dependent solution; and that, for the foreseeable future, the presence of each in space is an absolute necessity for the effective use of the other. If continued development of robotic systems renders humans in space obsolete, that must be a rational, conscious decision made by society as a whole, based on factors beyond those appropriate to an engineering overview paper.

BIBLIOGRAPHY

- Atkinson, J. D., Jr., and Shafritz, J. M.
1985 The Real Stuff: A History of NASA's Astronaut
 Recruitment Program. Praeger Scientific.
- Bilstein, R. E.
Stages to Saturn: A Technological History of the Apollo/Saturn
Launch Vehicles. National Aeronautics and Space
Administration. NASA SP-4206

- Brooks, C. G., Grimwood, J. M., and Swenson, L. S., Jr.
1979 Chariots for Apollo: A History of Manned Lunar Spacecraft. National Aeronautics and Space Administration. NASA SP-4205.
- Collins, M.
1974 Carrying the Fire. Ballantine Books.
- Compton, W. D., and Benson, C. D.
1983 Living and Working in Space: A History of Skylab. National Aeronautics and Space Administration. NASA SP-4208.
- Connors, M. M., Harrison, A. A., and Akins, F. R.
1985 Living Aloft: Human Requirements for Extended Spaceflight. National Aeronautics and Space Administration. NASA SP-483.
- Furness, T.
1983 Manned Spaceflight Log. Van Nostrand Reinhold Co., Inc.
- Hacker B. C., and Grimwood J. M.
1977 On the Shoulders of Titans: A History of Project Gemini. National Aeronautics and Space Administration. NASA SP-4203.
- National Academy of Sciences National Research Council
1972 Human Factors in Long-Duration Space Flight. Space Science Board.
- National Aeronautics and Space Administration
1969 Proceedings of the Winter Study on Uses of Manned Space Flight. NASA Science and Technology Advisory Committee for Manned Space Flight, NASA SP-196.
- Pitts, J. A.
1985 The Human Factor: Biomedicine in the Manned Space Program to 1980. National Aeronautics and Space Administration. NASA SP-4213.
- Smith, D. B. S., ed.
1976 A Systems Design for a Prototype Space Colony. Department of Aeronautics and Astronautics, Massachusetts Institute of Technology.
- Swenson, L. S., Jr., Grimwood, J. M., and Alexander, C. C.
1966 This New Ocean: A History of Project Mercury. National Aeronautics and Space Administration. NASA SP-4210.

SHARING COGNITIVE TASKS BETWEEN PEOPLE AND COMPUTERS IN SPACE SYSTEMS

William H. Starbuck

WHAT ARE THE RELATIVE ADVANTAGES OF PEOPLE AND COMPUTERS?

Mankind's capabilities change very slowly, whereas computers' capabilities have been fast-changing. The cost of a memory component has dropped forty percent per annum for over thirty years, and memory sizes have grown even more rapidly than that (Albus, 1981; Toong and Gupta, 1982). Computation speeds have been accelerating nearly 25 percent yearly, the cost of logic hardware has been dropping equally rapidly, and the computation work done with each unit of energy has been rising thirty percent per annum. Computing hardware has become much more reliable and very much smaller. User interfaces and programming languages have improved considerably, especially over the last decade. If human beings had evolved as rapidly as computers since the mid 1950s, the best runners would now finish a 26-mile marathon in 2.3 seconds, a bright student would complete all schooling from kindergarten through a Ph.D. in a bit over two days, normal eaters would consume one calorie per month, and half of America's families would be earning more than \$141,000,000 annually.

The improvements in computing costs, sizes, and speeds have generally exceeded the most optimistic forecasts of yesteryear, as has the proliferation of computers. Unfulfilled, however, have been the forecasts predicting that computers would shortly be able to imitate human beings. For example, in 1960 Simon optimistically speculated that "Duplicating the problem-solving and information-handling capabilities of the brain is not far off; it would be surprising if it were not accomplished within the next decade" (Simon, 1960:32).

Computers have not, in fact, developed an ability to reason very much like people, and computer simulation of human thought has had little success (Albus, 1981). When computers look most effective solving problems, the computers use quite different techniques than people apply (Weizenbaum, 1965; Winograd and Flores, 1986). For example, Newell et al. (1957) studied students' efforts to prove theorems in mathematical logic, and inferred that the students search for proofs, using heuristics that generally lead toward proofs but do not guarantee them. Challenged by such work, Wang (1963) devised a computer program that efficiently proved all 200 theorems in the first five chapters of Principia Mathematica. Job-shop scheduling affords

another example: Scientific-management studies of human production schedulers led to the development of Gantt charts to portray graphically the activities of various machines, and thus to help human schedulers visualize the cascading implications of alternative assignments. Computers generate job-shop schedules by solving integer-programming problems that no human could solve correctly without machine assistance.

The differences between people and computers have an illusory quality, insofar as people tend to take prevalent human abilities for granted and to notice rare or inhuman abilities. If computers did operate exactly like people do—working at the same speeds, making the same mistakes, showing the same fatigue, complaining about unpleasant tasks, and so on—people would regard computers merely as inhuman labor. Computers most impress people when they augment human abilities significantly—by working silently and tirelessly, by calculating with dazzling speed, or by displaying total consistency.

But the quite real differences between people and computers are persistent and profound. Rather than regard computers as potential imitators of human beings, it makes better sense to look upon them as a distinct species—a species that prefers different languages, reasons with somewhat different logic, finds comfort in different habitats, and consumes different foods.

Computers are much better symbol manipulators and much stricter logicians than people; and computers are much more decisive, literal, precise, obedient, reliable, consistent, and transparent. Computers can act both much more quickly and much more slowly than people. If so instructed, computers will carry out utterly absurd instructions or they will remain completely calm in the face of impending disaster. Computers easily simulate what-if conditions; and they can extrapolate even the most farfetched implications of theories or conjectures.

People, on the other hand, possess brains that are so much more complex than the largest computers that comparisons make no sense. These brains carry on numerous simultaneous and interacting processes, some of which operate entirely automatically. Without even trying, people process visual and auditory data of great complexity. People can shift levels of abstraction from detail to generality and back, they separate foreground images from background images, they distinguish patterns while remaining aware of contexts, and they attend to important or unusual stimuli while ignoring unimportant or routine stimuli. People have quite extensive memories that possess meaningful structures; and if they have relevant information in their memories, people usually know it and they can usually find it. People can operate with imprecise and somewhat incomplete plans, and they can extrapolate their past experiences to novel situations while recognizing that they are indeed operating outside the limits of their direct experience (Allen, 1982; Dreyfus and Dreyfus, 1986; Moray, 1986; Reason, 1986; Winograd and Flores, 1986).

Perhaps most importantly, people are more playful than computers and better at making mistakes. Whereas computers obey instructions literally, people often ignore or forget instructions, or interpret them loosely. Not only do people tend to deviate from plans and to

test the limits of assumptions, but many human perceptual skills and response modes depend on observing deviations from expectations or goals that may be evolving. Sometimes, people begin to doubt even their most basic beliefs. Thus, people generally expect to make mistakes and to learn from them, and creative people may be very good at learning from mistakes. If they have sufficient time, people can learn to correct their mistakes and they can reprogram themselves to take advantage of unexpected situations. Although computers also observe and react to deviations, computers have not yet exhibited much capability to devise goals for themselves, to reprogram themselves, or to question their own basic premises (Valiant, 1984). Computers must be told to learn from their experiences, and efforts to enable them to learn have, so far, been restricted to very narrow domains of activity. Also, computers are good at not making mistakes in the first place, so they have less need to learn from mistakes.

People are, however, pretty diverse and flexible. Some people can learn skills and perform tasks that other people find impossible; and since NASA can choose from a large pool of applicants, the extreme capabilities of exceptional people are more important in space systems than the average capabilities of typical people. The people who operate space systems first receive thorough training, so their deficits of inexperience should be small; but this training itself may impose serious liabilities, such as a tendency to rely on well-practiced habits in novel situations.

Because people are flexible and complex, they often surprise scientists and systems designers: People may change their behaviors significantly in response to ostensibly small environmental changes, or people may change their behaviors hardly at all in response to apparently large environmental changes. How people react to a situation may depend quite strongly on the sequence of events leading up to that situation, including the degree to which the people see themselves as having helped to create the situation. Accurate statements about microscopic details of human behavior rarely prove accurate as statements about general, macroscopic behavioral patterns, or vice versa. For example, experimental studies of people who are being paid low hourly wages for making repeated choices between two clearly defined, abstract symbols that have no implications for later events probably say little about human behavior in real-life settings where actions may have persistent and personally significant consequences and where actors may not even perceive themselves as having choices. Conversely, broad generalizations about the behaviors of most people in diverse situations probably say little about the behaviors of carefully selected people who are performing unusual tasks in which they have great experience.

The research issues that are important for designing human-computer systems seem to be ones concerning the proper balances among opposing advantages and disadvantages, rather than ones demanding new concepts; and the best resolutions of these issues are certain to shift as computers acquire greater capabilities. Consequently, I will not attempt to state any generalizations about the proper dividing lines between human and computer responsibilities in space systems, and I am

not advocating any research aimed at describing human capabilities in general. The designers of space systems should not depend on general theories, but should test fairly realistic mock-ups of interfaces, hardware, and software, with people who are as well trained and as able as real astronauts and controllers. The designers should also investigate the sensitivity of performance measures to small variations in their designs (Gruenenfelder and Whitten, 1985): Do small design changes produce large changes in performance? Both to improve the quality of designs and to improve users' acceptance of designs, experienced astronauts and controllers should participate in the designing of interfaces and systems; and because early decisions often constrain later modifications, astronauts and controllers should participate from the beginning of any new project (Grudin, 1986).

PEOPLE INTERACTING WITH COMPUTERS

Today's computers cannot imitate people very closely, but the differences between people and computers imply that combinations of the two can achieve results beyond the capabilities of each alone. For that reason, NASA should devote research effort to improving the interactions and synergies between people and computers.

Five research topics seem especially interesting and important because (a) I can see how to pursue them and (b) I can foresee some research findings that would translate directly into improved performances by space systems.

1. Fostering Trust Between People and Expert Systems
2. Creating Useful Workloads
3. Anticipating Human Errors
4. Developing Effective Interface Languages
5. Using Meaningful Interface Metaphors

Fostering Trust Between People and Expert Systems

Decision-support systems are computer programs and data bases that are intended to help people solve problems. Some decision-support systems merely afford their users easy access to data; other decision-support systems actually propose solutions, possibly basing these proposals on data supplied by their users (Woods, 1986b).

Expert systems are decision-support systems that attempt to embody the specialized knowledge of human experts. Their proponents argue that expert systems can, in principle, make specialists' knowledge available to nonspecialists: every CPA might be able to draw upon the combined expertise of several tax specialists; every general practitioner might be able to make subtle diagnoses that reflect advanced training in many specialties. Expert systems might perform even better than human experts: Computers may be able to obtain data that would be unavailable to people (Burke and Normand, 1987). Computers' huge memories and high speeds might enable them to

investigate more alternatives or to take account of more contingencies than people consider. Computers may also avoid some of the logical errors to which people typically fall prey, and thus may draw some inferences that people would miss (Bobrow et al., 1986). Advocates of statistical decision theory value computers' ability to adhere quite strictly to such formulae. Some proposals would have computers formulating recommendations and people then screening these recommendations and deciding whether to accept them (Burke and Normand, 1987; Dreyfus and Dreyfus, 1986; Woods, 1986a, 1986b).

Not everyone holds an optimistic view of expert systems' potential. Stanfill and Waltz (1986:1216) remarked: "Rule-based expert systems ... tend to fail badly for problems even slightly outside their area of expertise and in unforeseen situations." Dreyfus and Dreyfus (1986:108) have argued that human experts do not follow decision rules but instead they remember "the actual outcomes of tens of thousands of situations", and that "If one asks the experts for rules one will, in effect, force the expert to regress to the level of a beginner and state the rules he still remembers but no longer uses." Consequently, Dreyfus and Dreyfus (1986:109) predicted "that in any domain in which people exhibit holistic understanding, no systems based upon heuristics will consistently do as well as experienced experts, even if those experts were the informants who provided the heuristic rules."

Dreyfus and Dreyfus' critique may be valid. Dutton and I (1971) spent six years studying an expert production scheduler named Charlie, including one full year investigating his procedure for estimating how much production time any schedule represented. Charlie estimated time by using the relation:

$$\text{Production Time} = \text{Schedule Length} / \text{Speed}$$

"We gradually were disabused of the idea that Charlie has a computation procedure for speed and were convinced that he obtains his speed estimates by a table look-up. That is, Charlie has memorized the associations between speed and schedule characteristics, and he looks up speeds in his memory in somewhat the way one looks up telephone numbers in a directory. In our interviews, Charlie talked as if the existence of a computation procedure was a novel idea, intriguing to contemplate but difficult to conceive of. He thinks of the speeds in his table as discrete numbers distilled from a long series of unique experiences. Although he can interpolate and extrapolate these numbers--implying that the stored speeds must be specific examples from a systematic family of numbers--he distrusts the interpolated values and speaks of them as hypotheses to be tested in application. The stored values are so much more reliable that they might be a different kind of information altogether. In fact, Charlie can recount, for a large proportion of his table entries, specific remembered situations in which the circumstance was encountered and the speed observed. The only speeds that he does not so document, apparently, are those appropriate to situations arising almost daily" (Dutton and Starbuck, 1971:230).

We calculated that Charlie had memorized approximately 5000 production speeds corresponding to various situations. But we also discovered that Charlie's production-time estimates could be predicted quite accurately by a simple linear equation that had a meaningful and generalizable interpretation in terms of the physics of the production process. Rather than thousands of machine speeds, this linear equation required only a few hundred parameters. Thus, we could state a procedure that was simpler than the one Charlie used; and because this artificial procedure had a physical interpretation, a user could more confidently extrapolate it to novel production situations.

One of the best-known expert-system projects not only produced a heuristic program, DENDRAL, but also led to the development of an efficient algorithm for generating molecular structures (Bennett et al., 1981). Evidently, the heuristic program has received little practical use whereas the algorithm has had much (Dreyfus and Dreyfus, 1986).

One obvious question is: why must expert systems closely resemble human experts? The proponents of expert systems typically equate expertise with human beings, so they see imitating human expertise as essential to creating expert systems; and their critics focus on the differences between computers and people. Yet, computers possess different abilities than people. Computer programming efforts that have begun by imitating human behavior have often ended up using techniques that made no pretense of imitating human behaviors; and engineers and scientists have devised, without imitating human expertise, many techniques that enable computers to exceed the best of human capabilities.

Other questions arise concerning people's willingness to depend upon computer-based expertise. Collins (1986) interviewed actual and potential users of several widely known expert systems for accounting, chemical analysis, mathematics, medical diagnosis, and computer-components ordering. She found only one of these expert systems that has active users: the one for ordering computer components (R1). It has straight-forward logical processes and it draws no subtle inferences; it mainly helps sales personnel forget no details when they fill in orders, and the sales personnel said they appreciated not having to waste their time worrying about details or waiting for access to a human expert. It may be relevant that the users of this system sold computing equipment. Concerning the other expert systems, potential users expressed considerable distrust, of other human experts as well as computers; and the potential users may view these systems as threatening their own expertise. However, the people who actually participated in creating these systems said they do trust them and would, but do not, use them. Collins inferred that trust in an expert system comes either from participating in the design process or from being able to change the system to reflect one's own expertise. This inference meshes with the general pattern of psychological research, but neither of these options was available to the computing-equipment sales personnel, who were the users voicing the greatest trust in an expert system.

Complex issues surround the idea that a user should screen an expert system's recommendations and decide whether to accept them. If an expert system draws the same inferences that its user would draw and if it recommends the same actions that the user would choose, that user will easily learn to trust the system. Such seems to be the case with the expert system for computer-components ordering. Such a system may relieve people from having to perform boring or easy work, but it adds very little to a user's intellectual capabilities, whereas in principle, computers' precise logic and extensive computation capabilities and the incorporation of exceptionally high-quality expertise might enable expert systems to draw substantially better inferences than their users and to choose distinctly better actions. Yet a user is quite likely to distrust an expert system that draws significantly different inferences and that chooses significantly different actions than the user would do. If the expert system also uses a computational procedure that diverges quite dramatically from human reasoning, the system may be unable to explain, in a way that satisfies users, why it draws certain conclusions and not others. Distrustful users may never discover whether an expert system is making good recommendations or bad ones.

This calls to mind the experience of a manufacturing firm that installed one of the first computer-based systems for job-shop scheduling. The system's creators promised that computer-generated schedules would produce considerable savings in comparison to human-generated schedules. The factory's managers, however, were not entirely sure of the goodness of computer-generated schedules, and they wanted to minimize the implied insult to their human production schedulers, so the managers told the schedulers to follow the computer's recommendations as long as they agreed with them, but to substitute their own judgement when they thought the computer had made bad recommendations. An evaluation conducted after one year showed that the computer-based system had yielded no improvements whatever.

But research may be able to suggest some answers to these issues, at least in part; and good design may be able to resolve them: Expert systems, even the ones that cannot meaningfully explain the reasoning that leads them to make certain recommendations, should be able to explain why they believe their recommendations to be good. People who cannot formulate a good recommendation may be able to recognize a good recommendation or a bad one, and people do sometimes recognize their own limitations. At least some of the people who manage factories have learned to trust computer programs for production scheduling or inventory control even though these people could not themselves generate the computers' solutions.

The foregoing observations highlight the practical significance of research about the factors that influence people's trust in computers' expertise. In what ways should a decision-support system's knowledge and logical rules fit each user individually? Given opportunities to tailor interfaces to their personal preferences, inexperienced users may design interfaces poorly (Dumais and Landauer, 1982): Do users trust systems more or less when tailoring is postponed until the users gain considerable experience? How do task characteristics affect a

user's willingness to trust a decision-support system? In what circumstances does a user decide to trust a computer system that captures the knowledge of experts whom the user does not know personally? What kinds of experiences lead a user to trust a decision-support system that the user regards, at least partly, as a blackbox? What kinds of experiences encourage a user to see a decision-support system's limitations and to override bad recommendations?

Creating Useful Workloads

Automation tends to make computers responsible for routine, easy tasks and to leave the nonroutine, difficult tasks for people. One reason for this may be the perception that nonroutine tasks are interesting and challenging, and thus worthy of human attention, whereas routine tasks appear easy and uninteresting, and so demeaning to people. But a more important reason may be the practicality that designers can figure out how to automate routinized activities whereas they cannot effectively automate activities that vary.

This division of labor produces the consequence that, as automation progresses, people's work becomes more and more diverse and unpredictable and it takes on more and more of an emergency fire-fighting character. At the same time, cutting people out of routine tasks isolates them from on-going information about what is happening and forces them to acquire this information while they are trying to perform nonroutine, difficult tasks. The human controllers in a system may not even be warned of gradually developing problems until the system exceeds critical limits and alarms go off (Weiner, 1985). Thus, people's work grows less do-able and more stressful (Senders, 1980); and extreme stress and extreme time pressure may cause people to do poorer work and less of it.

In many tasks, automation also increases the short-term stability of the variables used to monitor performance; as Weiner (1985:83) put it, "automation tunes out small errors and creates opportunities for large ones." De Keyser (1986) has suggested that this short-term stabilization causes the human operators to shift from an anticipation logic to a recovery logic: instead of keeping track of events and trying to manage them, the operators wait for significant undesirable events to occur. Furthermore, "At the highest automation stage, the production operator has only very sketchy operating images of process and installation.... He will not make a huge investment in observation, checking, judging, establishing relationships, gathering of data without being certain of its usefulness. The operator does not invest psychologically in a role which escapes him" (De Keyser, 1986:234-235). Hence, De Keyser et al., (1986:135) have advocated that "the person still play an active part in the ongoing activity, not because this presence is required, but because it automatically keeps the person up to date on the current status of the system, the better to respond if an emergency situation develops." This seems a plausible hypothesis, but an equally plausible hypothesis would be that operators

tend to work mechanistically when they are performing the kinds of activities that could be automated.

De Keyser also, however, pointed out that serious emergencies call for as much automation as possible because they produce extreme time pressures, extremely complex problems, and extreme dangers--all of which greatly degrade the capabilities of human operators. Of course, people are utterly unable to respond as quickly as some emergencies demand. This poses a Catch-22. As long as the designers of a system have sufficient understanding to be able to prescribe how the system should respond to a serious emergency, they should incorporate this understanding in the system's automatic responses. But such complete understanding should imply that the automatic system works so well that a planned-for serious emergency never occurs. Consequently, when a serious emergency does arise, is not design error one prominent hypothesis about its cause, and does that hypothesis not render suspect the diagnostic information being produced by the system? Any system-design process establishes a frame of reference that identifies some events as relevant and important, and other events as irrelevant or unimportant; and a cost-effective system monitors the relevant and important events and ignores the irrelevant and unimportant ones. But this is likely to mean that the system lacks information about some of the events that produce a serious emergency, and the incomplete information that the system does have available may well lead human diagnosticians astray. Moreover, human operators who participate continuously in a system might grow so familiar with the system and its current status that they overlook anomalies and lack the objectivity to respond effectively to a serious emergency.

Trying to diagnose the causes of an unexpected emergency and to develop remedies, human operators must understand computers and other machines extremely well, which implies that they are quite comfortable with computers and with the causal models they incorporate; but on the other hand, human operators must distrust their computers and computer-based models sufficiently to be able to sift computer-generated information with skeptical eyes. Similarly, confidence in their training can help people remain calm in an emergency, but confidence in their training also blinds people to its shortcomings. It thus seems likely that the people who do the most good in emergencies have an ability to discard their preconceptions and to look at situations from new points of view (Luchins and Luchins, 1959; Watzlawick et al., 1974). NASA should investigate the degrees to which such an ability varies among people and can be predicted or taught.

Workloads vary in duration as well as intensity. People can cope with very intense workloads for short periods, yet they experience stress from moderate workloads that persist for long periods (Turner and Karasek, 1984). Some physiological reactions to stress, such as ulcers and vulnerability to infection, take time to develop. Thus, the short-duration shuttle flights do not afford a good basis for forecasting the workloads to be experienced on long-duration tours in a space station. NASA should continue to investigate the workload

experiences gained from long stays in confined spaces such as Antarctica, Sealab, and nuclear submarines (Bluth, 1984).

Anticipating Human Errors

Overloading causes people to make errors, but so do boredom, inattention, and indifference. Human errors are both prevalent and inevitable (Senders, 1980), and many human errors are desirable despite their costs. People experiment, and some of their experiments turn out badly. People deviate from their instructions, and some of these deviations have bad consequences.

Norman (1983, 1986) and Reason (1979, 1986) have initiated research into the causes of errors and ways to prevent or correct them. Norman, for instance, distinguished errors in intention, which he called mistakes, from errors in carrying out intentions, which he called slips. He classified slips according to their sources, and then sought to prescribe remedies for various slips. Table 1 lists some of Norman's categories and prescriptions.

Recognizing errors' importance, NASA's Human Factors Research Division is currently conducting some well-thought-out research on error-detection and on error-tolerant systems. Error-detection systems would warn people when they appear to have omitted actions, to have acted out-of-order, or to have taken harmful actions. Error-tolerant systems would first detect human errors through unobtrusive monitoring and then try to remedy them.

This research has much to recommend it. But some errors are very costly to tolerate, and some errors are very costly or impossible to correct. So human-computer systems should also try to predict human errors in order to make serious errors unlikely in advance (Schneider et al., 1980; Shneiderman, 1986). That is, prevention may be cheaper and more effective than cure, and research on error prevention might usefully complement the current projects.

Of course, all human-computer systems express some assumptions about their human participants. These assumptions have nearly always been implicit; and they have nearly always been static, insofar as the assumptions have not changed in response to people's actual behaviors (Rouse, 1981; Turner and Karasek, 1984). For many tasks, it would be feasible to explicate fairly accurate models of people. In fact, models need not be very accurate in order to make useful predictions or to suggest where adaptability to people's actual behaviors might pay off. Computers might, for example, predict that people who respond to stimuli quickly are more alert than people who respond slowly; or they might predict that experienced people would respond more quickly than inexperienced ones; or they might predict that people would be more likely to behave in habitual ways than in unusual ways; or they might predict that people would be less concerned about small discrepancies when much activity is occurring. Based on a review of human-factors research, Simes and Sirsky (1985) hypothesized that:

TABLE 1 Some Error Categories and Prescriptions

Forming the Wrong Intentions

Mode errors:

misclassifications of systems' modes

Eliminate modes.
 Give better indications
 of modes.
 Use different commands
 in different modes.

Description errors:

ambiguous statements of
 intentions

Arrange controls
 meaningfully.
 Give controls
 distinctive shapes.
 Make it difficult or
 impossible to take
 actions that have
 serious, irreversible
 consequences.

Misdiagnoses:

Suggest alternative
 explanations.
 Point out discrepancies
 that might be
 overlooked.

Activating the Wrong Behaviors or Triggering Behaviors at the Wrong Times

Omissions:

Remind people of
 uncompleted actions.

Capture errors:

very familiar behaviors replace
 less familiar behaviors

Minimize overlapping
 behaviors.

Monitor actual behaviors where
 similar behavior sequences diverge.

SOURCE: Norman (1983, 1986)

- experience or frequent use of a computer system decreases people's need for immediate feedback (closure),
- experience or frequent use decreases the importance of human limitations in information processing,
- experience or frequent use decreases the impact of sensory overstimulation,

- task complexity increases inexperienced people's need for immediate feedback,
- task complexity increases the importance of human limitations in the information processing by inexperienced people, and
- task complexity increases the impact of sensory overstimulation.

As NASA's human-factors scientists well understand, computers that predict, detect, and remedy human errors raise issues about who is actually in control. When should people have the right to experiment or to deviate from their instructions?

Developing Effective Interface Languages

Communication between people and computers may resemble communication between people who come from very different backgrounds, say a tribesman from the Kalahari desert and a whiz-kid mathematician from Brooklyn. Because computers do differ from people, the people who interact with computers need to remain aware of these differences, and the interfaces for human-computer interaction should remind users of these differences. This need became clear during the 1960s, when Weizenbaum created a program, ELIZA, that conversed in English. ELIZA had almost no understanding of the topics about which it conversed. Instead, it imitated blindly the vocabularies of the people with whom it conversed; in effect, ELIZA merely repeated people's words back to them. Yet Weizenbaum (1976:6) observed: "I was startled to see how quickly and how very deeply people conversing with [ELIZA] became emotionally involved with the computer and how unequivocally they anthropomorphized it."

Weizenbaum's more colorful examples concerned people who did not have close acquaintance with computers. Nearly all of the research on human-computer interaction has focused on people who lacked thorough training and who had little experience with computers. Although such research findings can benefit the design of training programs, design characteristics that have strong effects on novices may have negligible effects on expert users, so most of these findings may not extrapolate to the well-trained and experienced operators of space systems. There is need for studies of well-trained and experienced users.

Sheppard, Bailey, and their colleagues (Sheppard et al., 1980, 1984) have run experiments with professional programmers having several years of experience. The first three experiments involved programs or program specifications that were stated either in flowchart symbols, or in a constrained program-design language, or in carefully phrased, normal English. These experiments asked experienced programmers to answer questions about program specifications, to write and debug programs, or to correct faulty programs. The fourth experiment omitted flowchart symbols and substituted an abbreviated English in which variables' names replaced their English descriptions; and the programmers were asked to add instructions to programs. Table 2 summarizes the results: Normal English turned out to be consistently inferior, and the program-design language proved consistently superior.

TABLE 2 How Experienced Programmers' Performances Vary with Different Languages

<u>First experiment: answer questions about program specifications</u>			
	Normal English	Flowchart Symbols	Program-design Language
Time needed to answer:			
Forward-tracing questions	45.9	37.6	35.1
Backward-tracing questions	46.8	37.6	35.8
Input-output questions	42.9	39.4	41.0
Percent of programmers preferring	14	33	53
<u>Second experiment: write and debug programs</u>			
	Normal English	Flowchart Symbols	Program-design Language
Time needed to write and debug programs	29.7	23.9	20.5
Editor transactions before solution	37	39	32
Attempts before solution	3.0	2.7	2.2
Semantic errors	2.4	1.4	.8
% of programmers preferring	6	35	59
<u>Third experiment: correct faulty programs</u>			
	Normal English	Flowchart Symbols	Program-design Language
Time needed to correct faulty programs	18.7	14.2	14.5
Attempts before solution	1.9	2.2	1.9
Percent of programmers preferring	33	34	33
<u>Fourth experiment: modify and debug programs</u>			
	Normal English	Abbreviated English	Program-design Language
Time needed to modify and debug	28.1	26.6	25.0
Semantic errors	.9	1.3	1.0
Percent of programmers preferring	18	32	50

SOURCE: Sheppard et al. (1980, 1984)

One liability of a natural language such as English is its generality: Because vocabularies are large and linguistic structures are flexible, much ambiguity surrounds each word, phrase, and sentence. Speakers can make statements that mean almost anything, or nothing. Even a restricted natural language, probably because it resembles unrestricted natural language, may make users uncertain what commands are legitimate and meaningful to the computer system (Jarke et al., 1985; Shneiderman, 1986). Ambiguity and unused complexity create noise.

Both people and computers absorb information faster and more accurately when their interactions make good use of themes, chunking, and sequences (Badre, 1982; Simes and Sirsky, 1985). Overall themes can help people or computers to predict what information to expect and what information is important. Effective chunking aggregates information into batches that have meaning within the context of specific tasks. Effective sequencing presents information in a familiar, predictable order. Themes, chunking, and sequences can improve communication in any language, but they may become more important when a language has more generality.

A second liability is that natural language evokes the habits of thinking and problem solving that people use in everyday life. Green et al. (1980:900-901) remarked, for example:

"The fundamental strategies of parsing used by people seem, in fact, to be aimed first and foremost at avoiding parsing altogether

- (i) if the end of the sentence can be guessed, stop listening;
- (ii) if semantic cues or perceptual cues (boldface, indenting, pitch and stress in speech) are enough to show what the sentence means, stop parsing;
- (iii) if syntactic signals (and, -s, -ly, etc.) are available, use them to make a guess at the sentence structure;
- (iv) if there is no help for it, make a first shot at parsing by cementing together the closest acceptable pairings--noun to the nearest verb, if to the next then, etc.;
-) only if that first shot fails, try to figure out the structure by matching up constituents properly.

Not until Step (v) does the human start to parse in a manner anything like the computer scientists' idea of parsing; and the phrase 'figure out' has been used advisedly, for by the time that step is reached people are doing something more like problem solving than routine reading or listening."

Information displays can improve comprehension by offering symbolic and, especially, perceptual cues that help people to interpret messages. However, designing good displays is made complicated by the potentially large effects of overtly small cues. In a study of a command language, for instance, Payne et al. (1984) found that users' errors dropped 77 percent when the operator words were displayed in upper case and the operands were displayed in lower case, thus providing visual distinction between the two categories. Further, changes that improve performance in one context often degrade

performance in another context, and changes that improve one dimension of performance often degrade another dimension of performance. A flowchart, for example, may help users to trace forward to the consequences of some initial conditions but it may impede their backward inferences about the antecedents of some terminal conditions (Green, 1982).

A third liability may be that natural languages lead users to assume that computers' reasoning resembles human reasoning, whereas artificial programming or query languages remind users that computers' reasoning differs from human reasoning. This suggests that languages resembling natural ones might be more effective media for communication between people and computers in contexts where the computers closely simulate human reasoning and understanding, even though artificial languages might be more effective communication media in applications where computers deviate from human reasoning.

Unstudied so far are the interactions between social contexts and interface languages; virtually all studies of interface languages have involved people working on tasks that they could perform alone. Yet space systems create strong social contexts. The operators talk with each other while they are interacting with computers: Queries between people instigate queries to computers, and messages from computers become oral statements to other people. De Bachtin (1985) found that sales personnel who were interacting with a computer and customers simultaneously greatly preferred an interface that allowed them to pose queries in rather free sequence and phrasing. Thus, interface languages that approximate natural languages might turn out to be more valuable in space systems than in the situations that have been studied.

Using Meaningful Interface Metaphors

One very significant contribution to human-computer interaction was Xerox's Star interface, which derived from many years of research by many researchers. The Star interface embodies a number of design principles that evolved from experiments with prototypes. According to Canfield Smith et al. (1982:248-252), "Some types of concepts are inherently difficult for people to grasp. Without being too formal about it, our experience before and during the Star design led us to the following classification:

<u>Easy</u>	<u>Hard</u>
concrete	abstract
copying	creating
choosing	filling in
recognizing	generating
editing	programming
interactive	batch

The characteristics on the left were incorporated into the Star user's conceptual model. The characteristics on the right we attempted to avoid.

"The following main goals were pursued in designing the Star user interface:

familiar user's conceptual model
 seeing and pointing versus remembering and typing
 what you see is what you get
 universal commands
 consistency
 simplicity
 modeless interaction
 user tailorability

"...We decided to create electronic counterparts to the physical objects in an office: paper, folders, file cabinets, mail boxes, and so on--an electronic metaphor for the office. We hoped this would make the electronic 'world' seem more familiar, less alien, and require less training.... We further decided to make the electronic analogues be concrete objects. Documents would be more than file names on a disk; they would be represented by pictures on the display screen. They would be selected by pointing to them.... To file a document, you would move it to a picture of a file drawer, just as you take a physical piece of paper to a physical file cabinet."

NASA's Virtual Environment Workstation illustrates a much more avant-garde metaphor (Fisher et al., 1986). This project would give a robot's operator the sensations and perspective of the robot: Screens in the operator's helmet would show views taken by cameras on the robot; sensors would pick up the operator's arm and finger movements and translate them into movements of the robot's arms; and the operator's gloves would let the operator feel pressures that the robot's fingers feel. The operator would have the sensation of being inside the robot, and the robot would become an extension of the operator's arm and hand movements, even though the robot might be many miles from the operator.

Although metaphors constitute a fairly new frame of reference for the designers of interfaces, a designer or user can look upon every interface as a metaphor of something, and thus the design issue is not whether to adopt a metaphor but what metaphor to adopt. Each metaphor has both advantages and disadvantages. As Star's designers noted, an effective metaphor can both reduce the amount of learning that inexperienced users must do and accelerate that learning. An effective metaphor can also tap into users' well-developed habits and thereby reduce errors and speed responses; and experienced users as well as inexperienced users show such improvements. For instance, Ledgard et al. (1980) slightly modified a text editor so that its commands resembled short English sentences: The original, notational command RS:/KO/,/OK/* became CHANGE ALL "KO" TO "OK", and the notational command FIND:/TOOTH/ became FORWARD TO "TOOTH". As Table 3 shows, such

changes improved the performances of fairly experienced users as well as inexperienced users.

TABLE 3 Text Editing With Different Command Languages

	English-like Commands	Notational Commands
<u>Users with less than 6 hours of experience:</u>		
Percentage of tasks completed correctly	42	28
Percentage of erroneous commands	11	19
<u>Users with more than 100 hours of experience:</u>		
Percentage of tasks completed correctly	84	74
Percentage of erroneous commands	5.6	9.9

SOURCE: Ledgard et al. (1980)

But every interface metaphor breaks down at some point, both because a metaphor differs from the situation it simulates and because an interface differs from the computer it represents. People in real offices can take actions that users cannot simulate in Star's electronic office, and Star's electronic office allows actions that would be impossible in a real office. Similarly, a robot might be unable to reproduce some of its operator's instinctive finger movements, and an operator in a shuttle or space station would lack the mobility of an unconfined robot. Yet, users are likely to draw strong inferences about a computer's capabilities from the human-computer interface. Ledgard et al. (1980:561) noticed that "the users made no distinction between syntax and semantics... To them, the actual commands embodied the editor to such an extent that many were surprised when told after the experiment that the two editors were functionally identical."

One implication is that an interface metaphor, like an interface language, should maintain some intentional artificiality in order to warn users of its limitations. Are some of the intuitive expectations that users bring to metaphors especially important to fulfill? For example, in designing the Virtual Environment Workstation, might it be essential to use cameras that closely approximate the spacing and movements of human eyes in order to avoid having to retrain the operator's stereoscopic vision? Under stress, people tend to revert from specific, learned, complex models back to generic, common-sense, simple models: Which of the expectations that users have unlearned through training does stress reawaken? Does stress, for instance,

increase users' responsiveness to concrete, visible stimuli and decrease their responsiveness to abstract, invisible stimuli?

A second implication is that designers should carefully explore the limitations of an interface metaphor before they adopt it, and they should look upon a metaphor as one choice from a set of alternatives, each of which has advantages and disadvantages. However, the existing interface metaphors have been developed separately, with considerable emphasis being given to their uniqueness; and the processes that developed them have been poorly documented. So, interface designers need to be able to generate alternative metaphors, they need conceptual frameworks that highlight the significant properties of different metaphors, and they need systematic research to document these properties.

* * *

All of the foregoing topics imply that a computer should adapt both its appearance and the rules in programs to its user--to take account, for example, of its user's technical expertise, experience, frequency of use, or manual dexterity. This calls for development of sophisticated interface software (a so-called User Interface Management System) that will recognize the needs of different users, allow different users to express their personal preferences, and protect users' individuality. Thus, the computer needs to be able to identify a user quickly and unequivocally, and if possible, without imposing an identification procedure that would irritate people or delay their access in an emergency.

PEOPLE ADD IMAGINATION AND POETRY

Efforts to justify space systems in economic terms will keep pressing for higher and higher levels of measurable productivity, and so planners will tend to program the operators' activities in detail. But very heavy workloads raise the probabilities of human error, and computers will always be better than people at working tirelessly and obediently adhering to plans. People contribute to space systems their ability to deal with the unexpected, and in fact, to create the unexpected by experimenting and innovating. They can make these contributions better if they are allowed some slack.

Space systems' tasks are not all located in space. Space systems inevitably make educational contributions that transcend any of their immediate operational goals. One of the major contributions of the space program to date has been a photograph--a photograph of a cloud-bedecked ball of water and dirt isolated in a black void. Before they saw that photograph, people's understanding that mankind shares a common fate had to be abstract and intellectual; the photograph has made this understanding more tangible and visceral.

People play central roles in educational activities because they serve as identifiable points of reference in settings that would otherwise seem mechanistic, remote, and alien. Another of the space

program's major contributions, because it put space exploration into words that caught the human imagination, was Neil A. Armstrong's unforgettable observation: "That's one small step for a man, one giant leap for mankind" (July 20, 1969).

SUMMARY OF RECOMMENDATIONS AND QUESTIONS FOR RESEARCH

Fostering Trust Between People and Expert Systems

What ways should a decision-support system's knowledge and logical rules fit each user individually? Do users trust systems more or less when tailoring is postponed until the users gain considerable experience?

How do task characteristics affect a user's willingness to trust a decision-support system?

In what circumstances does a user decide to trust a computer system that captures the knowledge of experts whom the user does not know personally?

What kinds of experiences lead a user to trust a decision-support system that the user regards, at least partly, as a black-box?

What kinds of experiences encourage a user to see a decision-support system's limitations and to override bad recommendations?

Creating Useful Workloads

Does performing activities that could be automated actually keep human operators up to date on the status of a system, or do operators tend to work mechanistically when they are performing routine activities? Do human operators who perform activities that could be automated respond more effectively to a serious emergency because their participation updates them on the current status of the system, or does continuous participation make operators so familiar with the system and its current status that they overlook anomalies and lack the objectivity to respond effectively to a serious emergency?

NASA should investigate the degrees to which an ability to discard preconceptions varies among people and can be predicted or taught.

What have been the workload of experiences during long stays in confined spaces such as Sealab, Antarctica, and nuclear submarines?

Anticipating Human Errors

Research on error prevention might usefully complement the current projects on error detection and error tolerance. For many tasks, it would be feasible to explicate fairly accurate models of people that would enable human-computer systems to predict and adapt to human errors. In fact, models need not be very accurate in order to make

useful predictions or to suggest where adaptability to people's actual behaviors might pay off.

Developing Effective Interface Languages

Virtually all studies of interface languages have involved individual people working on tasks that they could perform alone. Because space systems create strong social contexts, interface languages that approximate natural languages may turn out to be much more valuable in space systems.

Using Meaningful Interface Metaphors

Are some of the intuitive expectations that users bring to metaphors especially important to fulfill?

Under stress, people tend to revert from specific, learned, complex models back to generic, common-sense, simple models: Which of the expectations that users have unlearned through training does stress reawaken?

Interface designers need to be able to generate alternative metaphors, they need conceptual frameworks that highlight the significant properties of different metaphors, and they need systematic research to document these properties.

General

NASA should develop a sophisticated User Interface Management System that will recognize the needs of different users, allow different users to express their personal preferences, and protect users' individuality.

Is there a way for a computer to identify its user quickly and unequivocally, without imposing an identification procedure that would irritate people or delay their access in an emergency?

Since NASA can choose from a large pool of applicants, the extreme capabilities of exceptional people are more important than the average capabilities of typical people.

The people who operate space systems first receive thorough training, so their deficits of inexperience should be small. Nearly all of the research on human-computer interaction has focused on people who lacked thorough training and who had little experience with computers, so most of these findings may not extrapolate to the well-trained and experienced operators of space systems. There is need for studies of well-trained and experienced users.

Avoid research aimed at describing human capabilities in general. Instead, test fairly realistic mock-ups of interfaces and systems, with people who are as well trained and as able as real astronauts and controllers.

Investigate the sensitivity of performance measures to small variations in designs: Do small design changes produce large changes in performance?

Both to improve the quality of designs and to improve users' acceptance of designs, experienced astronauts and controllers should participate in the designing of interfaces and systems. Because early decisions often constrain later modifications, astronauts and controllers should participate from the beginning of any new project.

ACKNOWLEDGEMENTS

This report has been improved by constructive suggestions from Michael Burke, Janet Dukerich, Kenneth Laudon, Henry Lucas, Frances Milliken, Jon Turner, Jane Webster, Keith Weigelt, and Harry Wolbers.

REFERENCES

- Albus, J. S.
1981 Pp. 97-99 in Brains, Behavior, and Robotics. New York: McGraw-Hill.
- Allen, R. B.
1982 Cognitive factors in human interaction with computers. Pp. 1-26 in A. N. Badre and B. Shneiderman, eds., Directions in Human/Computer Interaction. Norwood, N. J.: Ablex.
- Badre, A. N.
1982 Designing chunks for sequentially displayed information. Pp. 9-193 in A. N. Badre and B. Shneiderman, eds., Directions in Human/Computer Interaction. Norwood, N. J.: Ablex.
- Bennett, J., Buchanan, B., Cohen, P., Fisher, F., and others
1981 The DENDRAL programs. Pp. 106-123 in A. Barr and E. A. Feigenbaum, eds., The Handbook of Artificial Intelligence, Volume II. Los Altos, Calif.: William Kaufmann.
- Bluth, B. J.
1984 The psychology and safety of weightlessness. Pp. 63-79 in G. W. Heath, ed., Space Safety and Rescue, 1982-1983. San Diego, Calif.: American Astronautical Society.
- Bobrow, D. G., Mittal, S., and Stefik, M. J.
1986 Expert systems: perils and promise. Communications of the Association for Computing Machinery 29(9):880-894.

- Burke, M. J., and Normand, J.
 1987 Computerized psychological testing: an overview and critique. Professional Psychology: Research and Practice 18(1).
- Canfield Smith, D., Irby, C., Kimball, R., Verplank, B., and Harslem, E.
 1982 Designing the Star user interface. BYTE 7(4):242-282.
- Collins, J. S.
 1986 The Usefulness of Expert Systems: A Mirage? Manuscript, School of Business Administration, University of Wisconsin-Milwaukee.
- De Bachtin, O.
 1985 It is what it's used for--job perception and system evaluation. Pp. 689-692 in B. Schackel, ed., Human-Computer Interaction--Interact '84. Amsterdam: Elsevier.
- De Keyser, V.
 1986 Technical assistance to the operator in case of incident: some lines of thought. Pp. 229-253 in E. Hollnagel, G. Mancini, and D. D. Woods, eds., Intelligent Decision Support in Process Environments. Berlin: Springer.
- Dreyfus, H. L., and Dreyfus, S. E.
 1986 Pp. 110-118 in Mind Over Machine. New York: Free Press.
- Dumais, S. T., and Landauer, T. K.
 1982 Psychological investigations of natural terminology for command & query languages. Pp. 95-109 in A. N. Badre and B. Shneiderman, eds., Directions in Human/Computer Interaction. Norwood, N. J.: Ablex.
- Dutton, J. M., and Starbuck, W. H.
 1971 Finding Charlie's run-time estimator. Pp. 218-242 in J. M. Dutton and W. H. Starbuck, eds., Computer Simulation of Human Behavior. New York: Wiley.
- Fisher, S. S., McGreevy, M., Humphries, J., and Robinett, W.
 1986 Virtual Environment Display System. Manuscript, Aerospace Human Factors Research Division, NASA Ames Research Center, Moffett Field, Calif.
- Green, T. R. G.
 1982 Pictures of programs and other processes, or how to do things with lines. Behaviour and Information Technology 1:3-36.

- Green, T. R. G., Sime, M. E., and Fitter, M. J.
 1980 The problems the programmer faces. Ergonomics 23(9):893-907.
- Grudin, J.
 1986 Designing in the dark: logics that compete with the user. Pp. 281-284 in M. Mantei and P. Orbeton, eds., Human Factors in Computing Systems, CHI '86 Conference Proceedings. New York: Association for Computing Machinery.
- Gruenenfelder, T. M., and Whitten, W. B., II
 1985 Augmenting generic research with prototype evaluation: experience in applying generic research to specific products. Pp. 913-917 in B. Schackel, ed., Human-Computer Interaction--Interact '84. Amsterdam: Elsevier.
- Jarke, M., Turner, J. A., Stohr, E. A., Vassiliou, Y., White, N. A., and Michielsen, K.
 1985 A field evaluation of natural language for data retrieval. IEEE Transactions on Software Engineering SE-11(1):97-114.
- Ledgard, H. F., Whiteside, J. A., Singer, A., and Seymour, W.
 1980 The natural language of interactive systems. Communications of the Association for Computing Machinery 23(10):556-563.
- Luchins, A. S., and Luchins, E. H.
 1959 Rigidity of Behavior. Eugene: University of Oregon Books.
- Moray, N.
 1986 Modelling cognitive activities: human limitations in relation to computer aids. Pp. 273-291 in E. Hollnagel, G. Mancini, and D. D. Woods, eds., Intelligent Decision Support in Process Environments. Berlin: Springer.
- Newell, A., Shaw, J. C., and Simon, H. A.
 1957 Empirical explorations of the logic theory machine: a case study in heuristic. Joint Computer Conference, Western Proceedings 11:218-230.
- Nozman, D. A.
 1983 Design rules based on analyses of human error. Communications of the Association for Computing Machinery 26(4):254-258.
- 1986 New views of information processing: implications for intelligent decision support systems. Pp. 123-136 in E. Hollnagel, G. Mancini, and D. D. Woods, eds., Intelligent Decision Support in Process Environments. Berlin: Springer.

- Payne, S. J., Sime, M. E., and Green, T. R. G.
 1984 Perceptual structure cueing in a simple command language. International Journal of Man-Machine Studies 21:19-29.
- Reason, J.
 1979 Actions not as planned: the price of automatization. Pp. 67-89 in G. Underwood and R. Stevens, eds., Aspects of Consciousness, Volume I. London: Academic Press.
- 1986 Recurrent errors in process environments: some implications for the design of intelligent decision support systems. Pp. 255-270 in E. Hollnagel, G. Mancini, and D. D. Woods, eds., Intelligent Decision Support in Process Environments. Berlin: Springer.
- Rouse, W. B.
 1981 Human-computer interaction in the control of dynamic systems. Computing Surveys 13(1):71-99.
- Schneider, M. L., Wexelblat, R. L., and Jende, M. S.
 1980 Designing control languages from the user's perspective. Pp. 181-198 in D. Beech, ed., Command Language Directions. New York: North-Holland.
- Senders, J. W.
 1980 Is there a cure for human error? Psychology Today 14(4):52-62.
- Sheppard, S. B., Bailey, J. W., and Bailey, [Kruesi], E.
 1984 An empirical evaluation of software documentation formats. Pp. 135-164 in J. C. Thomas and M. L. Schneider, eds., Human Factors in Computer Systems. Norwood, N. J.: Ablex.
- Sheppard, S. B., Kruesi [Bailey], E., and Curtis, B.
 1980 The Effects of Symbology and Spatial Arrangement on the Comprehension of Software Specifications. Arlington, Va.: General Electric Company, TR-80-388200-2.
- Shneiderman, Ben
 1986 Seven plus or minus two: central issues in human-computer interaction. Pp. 343-349 in M. Mantei and P. Orbeton, eds., Human Factors in Computing Systems, CHI '86 Conference Proceedings. New York: Association for Computing Machinery.
- Simes, D. K., and Sirsky, P. A.
 1985 Human factors: an exploration of the psychology of human-computer dialogues. Pp. 49-103 in H. R. Hartson, ed., Advances in Human-Computer Interaction, Volume 1. Norwood, N. J.: Ablex.

- Simon, H. A.
1960 The corporation: will it be managed by machines? Pp. 17-55 in M. Anshen and G. L. Bach, eds., Management and Corporations 1985. Evanston, Ill.: Harper & Row.
- Stanfill, C., and Waltz, D.
1986 Toward memory-based reasoning. Communications of the Association for Computing Machinery 29(12):1213-1228.
- Toong, H. D., and Gupta, A.
1982 Personal computers. Scientific American 247(6):86-107.
- Turner, J. A., and Karasek, R. A.
1984 Software ergonomics: effects of computer application design parameters on operator task performance and health. Ergonomics 27(5):663-690.
- Valiant, L. G.
1984 A theory of the learnable. Communications of the Association for Computing Machinery 27(11):1134-1142.
- Wang, H.
1963 Toward mechanical mathematics. Pp. 91-120 in K. M. Sayre and F. J. Crosson, eds., The Modeling of Mind. Notre Dame, Ind.: University of Notre Dame Press.
- Watzlawick, P., Weakland, J., and Fisch, R.
1974 Change: Principles of Problem Formation and Problem Resolution. New York: Norton.
- Weiner, E. L.
1985 Beyond the sterile cockpit. Human Factors 27(1):75-90.
- Weizenbaum, J.
1965 ELIZA--a computer program for the study of natural language communication between man and machine. Communications of the Association for Computing Machinery 9(1):36-45.
1976 Computer Power and Human Reason. San Francisco: Freeman.
- Winograd, T., and Flores, F.
1986 Understanding Computers and Cognition. Norwood, N. J.: Ablex.
- Woods, D. D.
1986a Cognitive technologies: the design of joint human-machine cognitive systems. AI Magazine 6(4):86-92.

- 1986b Paradigms for intelligent decision support. Pp. 153-173 in E. Hollnagel, G. Mancini, and D. D. Woods, eds., Intelligent Decision Support in Process Environments. Berlin: Springer.

DISCUSSION: COMMENTS ON THE HUMAN ROLE IN SPACE SYSTEMS

Harry L. Wolbers

The theme of this symposium has been to delineate key research areas that need to be addressed in order to establish effective and reliable interaction of humans with automated and robotic systems in future manned space systems. Topics addressed in the earlier sessions included System Productivity, Expert Systems, Language and Display for Human-Computer Communication, Computer-Aided Monitoring and Decision Making, Telepresence and Supervisory Control, and Social Factors in Productivity and Performance. In this final session the speakers have addressed some of the broader issues related to the human role in future space systems.

Professor Starbuck has examined the sharing of cognitive tasks between people and computers and Professor Akin has examined the roles of humans and machines in previous space missions and has considered how these roles may change in the future.

In his paper, David Akin points out that any self contained device performing a useful function in space, whether human or robot, must rely on the same set of basic functions to adequately perform its mission. These include: sensory, computational, manipulative and locomotive capabilities and the environmental support functions necessary for the device to exist. Humans evolved in the environment of Earth's surface and are dependent upon a similar atmosphere and gravitational reference along with food, water and periodic rest/sleep periods. The space support systems for extended-duration manned missions must accommodate these human needs, perhaps even including a form of artificial gravity if it should prove necessary. On the other hand, machines can be designed to operate under a wide range of environmental conditions. The task which we face is to understand the capabilities and limitations of humans and machines as determined from their past and present roles in space and to extrapolate to the future. Akin presents the thesis that it is not an either/or choice because there are necessary and sufficient roles for both humans and machines and there are significant limitations on both.

Recent space missions have shown that the human operator offers combined advantages of manual dexterity and strength whereas most robotic systems available today are designed to provide either strength (e.g., for positioning large masses) or dexterity, but not both. On the other hand, humans can offer both capabilities. Humans represent

excellent adaptive control systems, especially well suited for rapid processing and integration of visual data. They have demonstrated their capabilities in space to move large masses along with the capability for precise psychomotor coordination in delicate mechanical adjustments.

Akin suggests that future research should be planned to produce a meaningful data base on human and machine capabilities and limitations in each of the functional categories. This will lead to a better quantitative understanding of the appropriate roles of humans and machines and will allow system planners to know which tasks are worth automating and which ones will best be done by humans for the foreseeable future. He points out that it is not enough to understand limitations and capabilities of each of the component technologies, but we must also understand the subtle interactions between the human and the machines to define the appropriate roles of each.

Recognizing that humans and machines may be able to perform the same tasks but may use different techniques in accomplishing them, Akin suggests that we also need to develop appropriate metrics in order to be able to produce meaningful comparisons.

He further points out that almost all of the designs currently proposed for telerobotic systems are anthropocentric tending to duplicate the human form. He suggests that since the human form evolved in a gravity field it may not be the best model for space activities and alternate forms and technologies should be studied.

Akin concludes that: (1) robotic systems are evolving rapidly, (2) both human and robotic systems have strengths and weaknesses; (3) for any future systems the best mixture of each is a time dependent solution; and (4) for the immediate future, the presence of each in space is an absolute necessity for the effective use of the other.

From my personal perspective, the criteria of performance, cost and missions success probability (program confidence based on schedule risk and technological risk) are the principal factors that program managers and system engineers use in selecting the optimum design approach for meeting mission objectives. Much as we may wish it to be otherwise, cost and cost effectiveness will continue to be important factors in designing future systems. I would urge, in addition to the metric comparisons of performance suggested by Akin, that where possible, indices of relative cost also be provided in order that design engineers may have a basis for ensuring the most cost effective utilization of the human operator in the space system of the future.

William Starbuck, in his paper, reminds us that people are flexible and complex. On one hand, they can change their behavior significantly in response to small environmental changes and on the other hand, they change hardly at all in response to apparently large environmental changes.

Starbuck has very eloquently highlighted the behavioral differences between people and computers and suggests that these differences can also mean that combinations of the two can achieve results beyond the capabilities of either alone. He stresses that in defining important research issues in human-computer systems we should be concerned with achieving the proper balance among the opposing advantages and

disadvantages and we must recognize that the dividing lines are fluid and depend heavily upon the evolving state-of-the-art in computer design. Accordingly, Starbuck suggests that space system designers should not depend on general theories but rather test specific implementation concepts with the actual users as subjects.

Starbuck suggests that future research efforts can profitably be directed toward improving the interactions and synergies between people and computers. He suggests five research topics as being especially interesting. These are:

- (1) Fostering Trust between People and Expert Systems: e.g., exploring questions regarding the degree a decision-support systems' knowledge and logical rules should be tailored to each user, and the factors that impact the users trust and acceptance of the computer system.
- (2) Avoiding Overload of Human Controllers: e.g., exploring the delicate balance between information overload, yet keeping the human in the loop by providing sufficient information for the human to respond appropriately when emergencies do arise.
- (3) Anticipating Human Errors: e.g., exploring the basic questions of people monitoring machines or machines monitoring people. Computers that predict, detect and remedy human errors raise issues about who is actually in control. Starbuck asks "When should people have the right to experiment or deviate from their instructions?"
- 4) Developing Effective Interface Languages: e.g., exploring the interactions between social contexts and interface languages. Starbuck points out that for experts, working alone, program design languages may be superior to natural language interfaces. On the other hand in space systems, operators with different cultural and scientific backgrounds may need to talk to each other while interfacing with computers and natural language interfaces may prove more effective.
- 5) Using Meaningful Interface Metaphors: e.g., exploring and establishing the conceptual frameworks that highlight the significant properties of different metaphors and their applications. (Every interface is a metaphor of something.)

Starbuck believes that NASA should develop a sophisticated User Interface Management System that will recognize the needs of different users, allow different users to express their personal preferences, and protect the user's individuality. He points out that in the foreseeable future, space crews will continue to represent an exceptional class of people in abilities, training and experience. This suggests to Starbuck a more immediate need for studies of well trained experienced users, rather than research aimed at describing human capabilities in general.

In providing a frame of reference for commenting upon the human factors research areas identified by William Starbuck and David Alin, we might note that NASA's current Space Station mission model covers a broad range of scientific and technical objectives. This model suggests that as the sophistication of future payloads increases, there will be an accompanying shift in crew support skills and requirements. A transition can be anticipated with the progression of time, from the more physical tasks of orbital assembly and installation to more intellectually oriented work activities.

To more effectively use human intelligence, a better match is required with machine intelligence and with "expert" systems. Work stations must (1) communicate fluently with humans (speaking, writing, drawing, etc.), (2) assist in interactive problem solving and inference functions, and (3) provide knowledge base functions (information storage, retrieval, and "expert" systems) for support.

Based upon the observations of the preceding speakers it would appear that the research issues related to work-station design would logically fall into three categories. These are: (1) Research on Information Seeking Processes, (2) Research on Information/Data Handling Processes, and (3) Research on Operation Enhancement Processes.

Research programs dealing with Information Seeking Processes should include sensory/perceptual research dealing with all sense modalities as well as continuing visual display development. (Continuing effort is required in the development of visual display formats, inasmuch as it is anticipated that, just as today, 80% of the information required by future space crews will be obtained through the sense of sight.)

I would group Starbuck's five research topics under the subject of Information/Data Handling Processes. In expanding his recommendations for establishing Meaningful Interface Metaphors I would also include, as a related topic, research and development of a Universal User Interface Management System (UIMS). This concept for a software system that handles all direct interaction with the user, potentially for a wide variety of underlying applications, began to emerge in the human-computer interface literature several years ago. The concept involves two main components: (1) a set of tools for developers to use in specifying visual and logical aspects of the user interface; and (2) a set of run-time programs and data bases for actually controlling interaction with the users. Some of the potential advantages of a UIMS would be:

- Independence of the user interface software and the application software.
- More intelligent user interfaces.
- Rapid prototyping capability for use in development.
- Easier involvement of manual systems and flight crew personnel in user interface design and evaluation.

- Consistency across applications.
- Multiple user interfaces to the same application if desired (e.g., novice vs. expert modes of interaction)
- Device independence (i.e., application software does not have to know anything about what type of input device a request came from or what type of output device the results will be displayed on.)

Although Starbuck does not advocate research aimed at describing human capabilities in general, I can't help but believe that continuing research on the nature of human cognition can provide insights that will lead to the development of work stations permitting more effective use of human cognitive capabilities. Conversely, studying the best and brightest representatives of the user community as they interact with the evolving concepts of expert systems, may in turn provide insights toward defining a structure of human intellect for mankind in general.

Research on Operation Enhancement should include those research areas identified by Akin such as intelligent robotics, and the mechanization of effector/actuator systems.

In addition to research dealing with Information Seeking, Information Handling, and Operational Enhancement Processes continuing attention also should be directed to the development of assessment techniques. These might include such areas as:

- Measurement of Human Productivity; i.e., continuing effort to develop valid measures of human performance and productivity in order to have meaningful criteria for evaluating performance and productivity adjustments caused by changes in operational procedures and system design concepts.
- Critical Incident Analyses of Human Performance; i.e., continuing effort to investigate and understand the cause of "human error" in space system operations, as well as incidents of exceptional performance, in order to identify and classify the causal factors of exceptional performance, in order to identify and classify the causal factors and establish guidelines for the designing of future space systems.

In closing this session on the human role in space, we can perhaps gain some perspective on the future research needs by looking at the lessons learned in previous manned space missions. We have learned from the US and Soviet¹ space programs to date that (1) systems can have indefinite operational lifetimes in space if they are designed to permit the contingency of in-flight repair and maintenance; (2) structures too large to be launched intact can be constructed and assembled on orbit, using man's unique capabilities; and (3) the flexibility and creative insights provided by the crew in situ significantly enhance the probability of successfully achieving mission objectives.

Reflecting upon their experiences as crew members of the Spacelab-1 mission, Garriott et al. (1984) succinctly described their activities in space by describing three levels of crew participation in accomplishing the mission objectives. At one level, the space crew found themselves highly involved in research activities and working together with principal investigators on the ground in the performance and real-time interpretation of research results. This was the case in areas such as space plasma physics, life sciences, and some materials-science and fluid-physics experiments. At another level, the crew found themselves performing other technical tasks with very little ground interaction. This was the case in the installation of cameras on a high-quality window or scientific airlock table and in the verification of their proper performance. At a third level, the specific experiments were largely controlled from the ground with the space crew participating only when needed to verify experiment performance or to assist in malfunction analysis and correction.

It can be anticipated that future space missions are likely to continue to require human support at each of these levels.

The ability of the crew to manually assemble delicate instruments and components and to remove protective devices, such as covers, lens caps, etc., means that less-rugged instruments can be used as compared to those formerly required to survive the high launch-acceleration loads of unmanned launch vehicles. As a result, complex mechanisms secondary to the main purpose of the instrument will no longer need to be installed for removing peripheral protective devices or activating and calibrating instruments remotely. With the crew members available to load film, for example, complex film transport systems are not needed, and malfunctions such as film jams can be easily corrected manually. The time required to calibrate and align instruments directly can be as little as 1/40th of that required to do the same job by telemetry from a remote location. Even for pure manipulative tasks, experienced operators are found to take as much as eight times longer using dexterous electronic-force-reflecting servomanipulators as compared to performing the same tasks by direct contact.

In future space missions specific experiments and operations no longer will need to be rigidly planned in advance, but can change as requirements dictate. One of the greatest contributions of crews in scientific space missions can be in reducing the quantity of data to be transmitted to Earth. One second of data gathered on SEASAT, for example, required 1 hour of ground-based computer time for processing before it could be used or examined, or a value assessment made. Before recording and transmitting data, scientist-astronauts in situ could determine in real-time whether cloud cover or other factors are within acceptable ranges.

The astronaut can abstract data from various sources and can combine multiple sensory inputs (e.g., visual, auditory, tactile) to interpret, understand, and take appropriate action, when required. In some cases the human perceptual abilities permit signals below noise levels to be detected. Man can react selectively to a large number of possible variables and can respond to dynamically changing situations. He can operate in the absence of complete information. He can perform a broad

spectrum of manual movement patterns, from gross positioning actions to highly refined adjustments. In this sense, he is a variable-gain servo system.

Thus, with the advent of manned platforms in space, there are alternatives to the expensive deployment of remotely manned systems, with their operational complexity and high cost of system failure. Long-term repetitive functions, routine computations or operations, and large-scale data processing functions, however, can be expected to be performed by computers capable of being programmed and serviced by crews in orbit, just as they are now serviced in ground installations. In addition, the normal functions of the terrestrial shop, laboratory, and production staff will find corollary activities in the work done by the crews manning the space platforms of the coming generation.

The human being represents a remarkably flexible and adaptable system. In terms of his basic capabilities and limitations, however, we must also remember that man is essentially invariant. In terms of basic abilities, people will not be much different in the year 2050 than they are today. Recognizing this constancy in sensory, perceptual, cognitive, and psychomotor abilities, the objective of the proposed research programs should be to improve system productivity through (1) hardware, software, and other system improvements that can enhance human performance, and (2) procedure and operational changes that will allow more effective use of the human element in the man-machine systems of the future.

NOTE

1. The Soviets have been reported to rely heavily on manned involvement in order to repair equipment and subsystems with serious shortcomings in reliable and trouble-free service life.

REFERENCES

- Garriott, O. W., Parker, R. A. R., Lichtenberg, B. K. and Merbold, U.
1984 Payload crewmembers view of spacelab operations. Science
225(July):165-167.

SYNOPSIS OF GENERAL AUDIENCE DISCUSSION

Following the presentations by the invited symposium speakers, the proceedings were opened to discussion and comment from the floor. A synopsis of the remarks made during this period of open discussion is presented below.

Stephen Hall, NASA Marshall Space Flight Center, referring to the apparent lack of acceptance of expert systems by many potential users (mentioned by Starbuck) asked, "Is this a fundamental limitation of expert systems or if not, what can be done to increase potential user acceptance?" In reply William Starbuck of New York University suggested that there are ways to teach people to trust expert systems. Starbuck pointed out that there are factory scheduling programs, for example, that people now trust. Many factory schedulers use such programs but have no idea how they work and couldn't replicate them if they wanted to. After using them for a period of time they learn to accept them. One key to acceptance is that the users learn that even if the computer may not be able to explain how it derived the answers to a problem, it can present the solution and provide an indication of how good it thinks the answer or solution is. Over time, the correlation of predicted and observed results instills confidence in the user.

Giulio Varsi, NASA Headquarters, suggested that not enough attention has been paid to the impact which the degree of media exposure can have on the acceptance and performance of space missions, and raised the question of the degree to which such exposure is appropriate. He cited the heroic image of the astronauts created to date. He wondered whether they are likely to receive this same degree of exposure in the future and how this exposure or lack of it may influence future performance. Varsi also commented on the issue of mission safety, pointing out that in addition to the criteria of performance and cost, safety--especially as related to human safety--should be of continuing concern. Varsi asked the question "As we move from the heroic to the routine, what is the real level of risk we are prepared to sustain?" As a final point, Varsi commented that many interesting research issues and questions for investigation were raised during the symposium and he suggested that an ordering of these research issues should be provided, highlighting their urgency not so much from the standpoint of priority but rather from the sequencing or logic to be followed in attacking

these problems. He asked, "Is there any one research program sequence that offers a more effective path to addressing the critical issues than any other one?"

In reply to Varsi's comments on risk adversity and safety, David Akin of MIT pointed out that, in his experience, NASA is already orders of magnitude more risk adverse than the undersea community, and if anything, NASA is becoming even more so in light of the Challenger accident. Akin suggested that if anything is going to drive people out of space entirely it is being absolutely risk free. The ultimate in risk adversity is for humans not to go into space at all. While robotic devices may appear to expand the options, in reality the considerations of risk adversity apply to equipment as well. Akin pointed out that in deciding to risk a one-of-a-kind \$100 million telerobotic servicer to service a satellite with an unfired solid rocket motor, the same issues of risk adversity must be raised for the hardware as would be raised for the crew in a manned mission. To put the issue in proper perspective it is necessary to consider risks and risk adversity in space in relation to potential risks and risk adversity in other fields.

Allen Newell of Carnegie-Mellon observed that no matter how dangerous it is, people believe it to be important and still want to go into space. One of the realities which must be faced is that by being so careful for the first 25 years, levels of National and World expectations of safety in space operations are very high and as a Nation we will suffer from that high level of expectation in the future.

Joseph Loftus of NASA Johnson Space Center observed that an airplane that is safe in peacetime is too dangerous to go to war. He pointed out that in an adversary relationship an airplane is needed that is at the peak edge of performance in order to succeed in its mission. Loftus commented that this is an important point when thinking of space operations because space operation is not a venture in isolation--it is a competition. It is an exploration at a frontier and safety standards cannot be set so high that the frontier is forfeited. At this point Session 7 of the Symposium was concluded.

CONCLUSION

CONCLUDING REMARKS

Allen Newell

In my view, three major issues emerge from this symposium:

1. The merging of AI and robotics.
2. The need to consider the human aspects of these AI-Robotic systems.
3. The potential benefits of incorporating the social sciences into the AI-robotic research effort.

Merging AI and robotics appears to be something that NASA has already identified as an important issue. It is, in fact, one of the great intellectual tasks in this part of the scientific world. With the merging of AI and robotics, AI will finally come to deal, not just with the symbolic world, but with interactions with space (the space of three local dimensions, not NASA's outer space): physical devices, movement, real time, compliance, etc.. This will radically change the field of AI. It is a big step, and its success will depend upon developing a real understanding of the nature of intelligence.

Once AI and robotics are welded together, the concerns for the human aspects of these systems must be addressed along with the concerns for the AI-robotic aspects. There are three distinct reasons for combining behavioral/cognitive science and AI/computer science in a single research program. First, the field of cognitive science--including physiological and motor behavior, not just cognitive behavior--provide major clues about developing effective AI-robotic systems. Second, the combination will allow researchers to address the concerns about human-computer interaction from several perspectives. Third, in order to evaluate the performance of automatic devices, much more needs to be known about human functioning in the tasks-to-be-automated. Human performance can be used as a metric of AI-robotic performance.

Finally, a move by NASA towards the social sciences, to incorporate them into an AI-robotic-cognitive science research program, would be very important in the long run. An area that could benefit from such a combination is communication--how people use the technology to communicate and interact with that technology and with each other. In

this regard, the human-computer interaction field is currently taking tentative steps to become much more socially and communication oriented.

Let me end with a remark about university research efforts. The universities, at this moment, are in an extremely pliant state with respect to developing cooperative efforts with external agencies. "Pliant", in this context, means that they are exploring, in a historic way, how to live with much deeper involvement with the industrial, commercial and government sectors. The ideal of the ivy tower seems far away indeed, although the concepts of independence and objectivity remain solidly in place. There are real opportunities for NASA to build its research programs in ways that will benefit both NASA and the universities involved.

Immense benefits can be garnered from long range, cooperative research programs established in conjunction with places like universities. A ten to fifteen year research relationship between NASA and a university might be expected to yield important dividends beyond the actual research accomplished. The university researchers come automatically to think in terms of NASA and its problems when developing (or expanding) their own research programs. Graduate students, raised in the NASA-oriented research environment, will have an ingrained concern for NASA problems--and are likely to make a career of dealing with those types of problems. These aspects, though not the stuff out of which research contracts can be made, can be of the highest importance to efforts such as inhabiting space that stretch out into the far future.

CONCLUDING REMARKS

Thomas Sheridan

The first thing I want to do is thank the speakers. We really appreciate the efforts you have put in. I also thank the organizers. A lot of effort went into getting this together. And I thank the participants--many useful and interesting comments have come from the floor. Our job, now, is to put together a report that makes sense, is not self-serving, in terms of 'please, Ma, send more money', but says, in effect, 'look, there are some really important research issues out there that are not receiving proper attention'.

I was taking notes, and some of my notes have little stars to indicate important points, for example:

- The idea of monitoring physiological state of the operator, as well as monitoring the computer and the mechanical state of the equipment was suggested. It seems to me that continually assessing the health of both is something that we don't still quite know how to do.
- There were a number of issues related to the difficulties of defining, and measuring, system productivity. At the very beginning, Ray Nickerson, addressed these issues. Bob Williges insisted that performance is a relative measure.
- Bruce Buchanan and Thomas Mitchell talked about the reality of non-numerical constraints. AI people have known this all along, but some of us other engineering types haven't particularly appreciated the importance of coping with those non-numerical, or qualitative, aspects of time, space, and resources. They also pointed out the problems of maintaining expert systems as situations change and new knowledge becomes available.
- Allen Newell characterized the trade-off between knowing versus searching, a priori knowledge versus getting new knowledge (somewhat related to the problem of optimal stopping in operations research).
- Robustness was mentioned many times, but we are not always clear what robustness implies.

- We heard about the difficulties of eliciting (and the need for a better "bedside manner" for eliciting) knowledge for the construction of expert systems.
- We also heard some questions raised about trust. I've looked in the literature on trust and there "ain't much there". We need to understand trust and transparency and that kind of thing vis-a-vis the relationship between intelligent systems and their users.
- Phil Hayes emphasized the graphical interface and how basic that is to the way people see, think, and make decisions. Peter Polson mentioned the fact that we are now able to, as the pilots say, "kill ourselves with kindness"--that is, provide graphic displays and "aids" that are so complicated that nobody understands them. This certainly could happen with expert systems. Randy Davis picked up the same point when he talked about designing to make understanding easier.
- Natural language was mentioned time and again, but it was also pointed out that it's no panacea. That there may be languages which are not "natural", but which are better for certain applications.
- Baruch Fischhoff talked about the need for shared models and the fact that people are not very well calibrated with respect to other people's questions and models of reality.
- We talked about the mechanical work, "manipulation". It also was pointed out that we need better models of (and notation for) characterizing the process of manipulation.
- Allen Newell suggested that we need a theory of presence. We know a little bit about the effects of fidelity in simulators from this point of view, but we need a much better understanding of what it means to feel "present".
- Karen Cook talked about computer-mediated communication, which we are going to have one hell of a lot more of than we have had in the past. We are not going to have situations where people are holding hands; they are going to be separated, and their communication is going to be mediated by computers. Questions of social stress and contending objectives are going to be aggravated or, at least, changed by computer mediated communication--and by all this "non-human expertise" that's floating around.
- In the last session, Dave Akin raised questions about the paucity of our human performance database, and what people can do relative to what machines can do. Harry Wolbers picked up on the same point.

- And, finally, a lovely notion, I think, made by Bill Starbuck is the importance of being playful and deviant.

Guilio Varsi asked about prioritizing these ideas. That takes a great deal of wisdom--but we will try¹.

There is a further comment that should be made. NASA has been extremely cautious about avoiding the risk of errors in space, especially when human life is concerned. This caution is very laudatory. Where human safety is not an issue, however, there can be more risk taking with respect to such areas as budgetary considerations, testing of equipment, and studies on the allocation of functions between people and automation/robotics to derive the best mix based on empirical evidence.

We have seen the evidence of this symposium that the computer scientists and the behavioral and human factors scientists can arrive at a common ground. We believe that this interface is obvious and extremely important for mission success based on the best of both worlds that is superior to either automation or humans used alone. In fact, we don't believe that either one can be used alone successfully at this time or in the future.

In conclusion, I thank you all for trudging through the snow and sleet and for your worth while contributions. I'm sure that it has been useful for all of us.

NOTE

1. The ideas presented by Professors Newell and Sheridan in their closing remarks were categorized and summarized (along with the "issues and research recommendations" presented within each paper) in the section titled "Assumptions, Conclusions and Recommendations".

APPENDIX: SYMPOSIUM PROGRAM

HUMAN FACTORS IN AUTOMATED AND ROBOTIC SPACE SYSTEMS

National Academy of Sciences
Lecture Hall
2101 Constitution Avenue, N.W.
Washington, D.C

January 29-30, 1987

PROGRAM

Thursday, January 29, 1987

- 8:00 Registration
- 9:00 Welcome and Introduction
Thomas Sheridan, (Chair, CoHF) MIT
Ray Colladay, Associate Administrator, Office of Aeronautics
and Space Technology, NASA Headquarters
Washington, D.C.
David Goslin, Executive Director, CBASSE, NRC
- 9:15 Keynote Address
Allen Newell, Carnegie-Mellon
- 9:35 Break
- 9:45 Session 1
System Productivity: People & Machines
Paper: Raymond Nickerson, Bolt Beranek and Newman Labs
Discussant: Robert Williges, VPI&SU
- 10:45 Break
- 11:00 Session 2
Expert Systems and Their Use
Paper: Thomas Mitchell, Rutgers
Paper: Bruce Buchanan, Stanford
Discussant: Allen Newell, Carnegie-Mellon
- 12:30 Break for Lunch
- 1:30 Session 3.
Language and Displays for Human Computer Communication
Paper: Phillip Hayes, Carnegie-Mellon
Paper: Peter Polson, U of Colorado
Discussant: Judith Reitman Olson, U of Michigan
- 3:00 Break

- 3:15 Session 4
Computer Aided Monitoring & Decision Making
Paper: Randall Davis, MIT
Paper: Baruch Fischhoff, Decision Research
Discussant: William Howell, Rice
- 4:45 Open Discussion
- 5:30 Reception in Great Hall

Friday, January 30, 1987

- 8:30 Session 5
Telepresence & Supervisory Control
Paper: Thomas Sheridan, MIT
Paper: Lawrence Stark, U of California
Discussant: Antal Beczky, JPL
- 10:00 Session 6
Social Factors in Productivity & Performance
Paper: Karen Cook, U of Washington
Paper: H. Andrew Michener, U of Wisconsin
Discussant: Oscar Grusky, U of California
- 11:30 Break for Lunch
- 12:30 Session 7
The Human Role in Space Systems
Paper: David Akin, MIT
Paper: William Starbuck, New York U
Discussant: Harry Wolbers, McDonnell Douglas
- 2:00 Concluding Remarks and Open Discussion
Allen Newell, Carnegie-Mellon
Thomas Sheridan, MIT (Chair)
- 2:30 Adjourn