Steven L. Mandell

The Management Information System Is Going to Pieces

Consider a primitive logging operation in which we have determined that the increasing weight of the timber has become too much for a single ox to pull. The solution is not to try to grow a bigger ox, but to either make a two-ox team or split the load and make an additional haul.[1] This would be a logical solution constrained by the laws of nature and the environment, whereas the computer is not limited to such restrictions—or is it?

The perceived need to increase data-processing capabilities by either moving to a larger and more powerful version of the present machine or switching to a newer and technologically superior machine is usually considered prima facie evidence that the data-processing department is successfully carrying out its responsibilities. The purpose of this article is not to minimize the miraculous success of computer technology or to reevaluate its accomplishments in providing useful information, but rather to present and defend an alternative approach to the "bigger/better syndrome" that is almost universally accepted as a truism.

Distributed minicomputer systems are designed to bring the computer system into the location where data are collected and information is required through a network of decentralized minicomputers in a communication-based information system. The advantages to be gained from implementing such a system are diseconomies of scale, reduced organizational impact, a greater degree of flexibility and responsiveness, and an increased fault tolerance.

Philosophy

Whenever there is a gathering of computer professionals, the talk always turns to the type of equipment utilized at the installation. Status appears to be measured by the size, reputation, and cost of the machines rather than the criticality and complexity of the information system. This inherent prejudice must be overcome before distributed minicomputer systems will be accepted, regardless of significant support for this new approach. One of the major predispositions that must be challenged is the acceptance of centralized computer systems as natural and effective. The distributed concept requires that both the processing and data bases be segmented or modularized with respect to the functional information requirements of the organization and that a communications network exist among these subsystems.

Although this functional-dispersion approach is gaining industry support, the primary method of implementation has been multiaccess, terminal-oriented computing with the requisite large centralized computer. This method has been defended on the bases of increased economies for resource sharing and optimal equipment utilization. I believe that present research refutes this line of reasoning while providing support for the distributed functional orientation.

Most people involved with minicomputers challenge the belief that there is any difference between one of the smaller computers and a general-purpose, large-scale computer. Advocates of the minicomputer claim that almost anything that can be done on a large computer can be done on one or more small ones in the same manner. This notion, however, negates the significant special benefits that the concept of a *minicomputer network* permits us. The minicomputer must be considered as basically task-oriented! By that I mean that although it could be used in a multiprogrammed environment (many jobs run concurrently), it most often should be used as a single-job processor with communications support. The advantages of small cost-effective machines then become apparent—their ability to handle problems on a segmented or modular basis, interconnected in a communications network. The well-used analogy still holds: there is a fundamental difference between a car and a bus.

The concept of modularity requires that the total set be divided into component subsets. In this way a large problem can be broken down into more manageable smaller problems, which can be solved individually before being recombined. This restructuring requires careful consideration for the areas of overlapping concern, as it is practically impossible to isolate mutually exclusive subsets. The justification for this approach is based upon the belief that the difficulty of a problem is related to its size and complexity not linearly, but exponentially. Therefore, the ability to solve a problem is greatly enhanced by first segmenting out problem areas that are somewhat limited in scope and constructing the necessary linkages among these separate modules. However, two limitations to this approach must be recognized: not all large problems can be segmented; and often the whole is greater than the simple sum of its parts. Within these rare structural limitations (such as complex modeling and simulation), the vast majority of problems facing an information system are divisible and solvable. It is interesting to note the rarity of actual integrated management information systems (MIS) now in operation as compared to the number of personnel, logistics, and financial subsystems.[2]

Steven L. Mandell is Assistant Professor of Quantitative Analysis and Control, Bowling Green State University, Bowling Green, Ohio. A former Associate Director, ADP Management Training Center, U.S. Civil Service Commission, his current research interest is computer-based information systems in multinational corporations.

Diseconomies of Scale

Since the 1950s a widely held belief has related computing power exponentially to computer cost. Grosch's law implies that doubling the outlay for a computer will quadruple the computing power acquired. In recent literature some authors claim—and often incorrectly document with time-series data—that a more accurate description of the existing relationship is the third power, or in our doubling case, increasing the computing ability by a factor of eight.[3,4] Unfortunately, until very recently there have been no published reports of detailed economic analysis regarding this accepted principle. The effect of blind faith in this law has been the bigger/better syndrome accompanied by the centralization psychology. An installation was thought to improve its cost-performance ratio by simply employing a bigger machine. A corollary to this was the belief that 100-percent machine utilization was a key to efficient computer operations. The marketing representatives of computer manufacturers, armed with this accepted principle and artificial restrictions imposed by their designers, continued to sell new and larger computers while technologically and economically reinforcing the obsolescence of existing equipment.

R. F. Littrell of the North Carolina State Board of Education recently conducted an economic study of the concept of "economy of scale" with relationship to IBM System 360 and System 370 computers.[5] The interesting, even revolutionary conclusion of this work is that economies of scale do exist for scientific calculations but not for business applications. "In some cases maximum throughput for minimum cost may be through a system of relatively small computers rather than a single large one," Littrell states.[6] In fact, it appeared that the smaller machines were more economical and efficient for commercial computing, which comprises the vast majority of information systems' applications. Artificial restrictions (such as core memory limitations) placed upon the different machines by manufacturers tend to support the notion that planned obsolescence is the way of life in the computer industry. Significantly, these results were based upon benchmark data supplied by the manufacturer. It is interesting to note that in recent court cases, IBM's confidential marketing reports tend to confirm the findings of Littrell.

It would be impossible to discuss economies of scale without mentioning another feature of cost, systems overhead. The operating systems required by the larger computer systems are orders of magnitude more complex than previously outdated ones. In fact, the concept of single-job batch processing is almost unheard of today in any large installation. There is a tradeoff that must be clearly recognized: the increasing sophistication of these system programs requires greater machine dedication to their support—in plain language, overhead. This reality has been defended by the extended capabilities offered and the ability to make 100-percent utilization of the bigger, more efficient computer. But both of these arguments can be seen as valid only if we accept the traditional bigger/better/centralization philosophy. Software development by the manufacturers is following this trend as data-base management systems, security systems, virtual storage systems, and others are made available in more complex versions. Eventually the modern computer may spend all its time in the systems state, with the scoreboard reading: "Direct Labor 0—Overhead 100."

Technology has made great advances in the area of computer science, with resulting increased cost-performance ratios for central processors. These developments have been viewed differently by the minicomputer manufacturers than by conventional general-purpose computer manufacturers.[7] Minicomputers were traditionally applications limited; hence the cost-performance improvement was passed along through cost reductions. Conventional systems, however, were cost-limited; their cost-performance improvement has been translated into performance enhancement. Additionally, the number of minicomputers sold has increased dramatically faster than other types, leading to manufacturing economies of scale and thereby further increasing the economic advantage and decreasing the cost of minicomputer systems. Unfortunately, similar advances cannot be found in the areas of memory and peripherals; thus, the cost-sensitive component is shifting away from the processor.

Much has been said about the economies associated with centralization of the data-processing organization in terms of manpower. We apparently forget that the increasing complexity of large-scale centralized systems requires that many people concern themselves with the system programs. Many authors concur in the belief that centralizing the data-processing organization leads to insolvable problems and increased game playing (negative work).[8,9] Historically, computer centers have grown into empire-building enterprises where little senior-level management involvement can be found. As a result, they are more often than not autonomous units stuck in an organization chart, regardless of reporting responsibility. Computers must be considered resources and as such must be controlled by sound management principles. Similar arguments, both pro and con, should be made for centralization of data processing and for centralized marketing operations.

A distributed network of minicomputers will entail additional costs in certain areas. If multiaccess remote computing is presently not being employed, then telecommunication costs can become substantial. Although the necessity for large-scale operating systems is avoided, a network-management structure must be created. The relatively high costs associated with memory and peripherals discourage extensive duplication of files and reporting facilities that would be advantageous in a network design. However, it is my firm belief that a network of distributed minicomputer systems is often more cost effective than a large centralized computer system and that the future direction of technical development will increasingly give the advantage to the network concept.

Organizational Impact

An information system should support the overall goals and objectives of an organization, as should any other staff service. Very often the data-processing group has gone far beyond this limited scope of providing support and has directly affected the goals, structures, policies, and personality of the organization in which it resides. This has occurred (as with no other staff operation) because of a singular lack of top-management involvement in its operations. Today, many organizations successfully operate in spite of, rather than because of, the information support obtained from the computer department.

The organizational impact of installing an information system must be minimal; hence it should be designed to merge into and match the existing organization. But computer systems have been designed by computer people for the large central computer, with the net result that most information-system designers successfully impose their will upon organizational operations.[10] Intrusion by the computer has created a great deal of the resentment felt by line-operations people, resulting in poor user-machine interfaces and leading to ineffective management information systems.

Organizations are normally structured to permit optimal operating effectiveness. People are assigned tasks that contribute to the overall accomplishment of group objectives. It would appear, then, that an infonnation system should be constructed along similar task-module lines. Despite this logical approach (which is at the heart of computer programming), system designers have continued to effect operating change rather than to encompass and support existing operations. The systems designed continually place artificial restrictions on local operations because of the need for centralized rules and regulations. Communication channels, both formal and informal, often are adversely affected and abused by the implementation of a so-called information system which requires a central point of focus.

Existing information systems that depend upon a central computer system have placed operational management in a difficult position. Although the computer people call for user involvement, there is no resultant cooperation, consideration, or responsibility. Management information must flow up the organization through a filtering, stripping, and summarizing process.[11] The centralized computer system does not support this, however, for the massive data base collects all available information. If nothing more, the concept of an information system should permit unbiased selective information to be placed in the hands of the person responsible for each specific management function.

The decentralized network approach of distributed minicomputer systems with each module matching the "people module" within the organization readily provides "a natural set of windows for monitoring the operation at all levels of management."[12] The user responsibility associated with such a network generates the missing ingredient in many modern systems—user involvement. Communication links can be structured to match organizational flows, thus requiring no change in existing operations. By matching the distributed computer modules with their associated functional elements, the system is more likely to be successful and useful. The degree to which decisions are based on the information supplied by a computer describes its usefulness, while the outcome of such decisions will quickly give an indication of its success. When computing power and data bases are distributed to the areas functionally responsible, users will be far more likely to trust the system and, at the same time, the input information probably will prove more reliable.

The distributed approach does require far more user involvement and responsibility than any of the existing systems approaches. The only hope is that functional managers have not been totally brainwashed as to the complexity of computers. The mystique surrounding information systems is unwarranted except for the quest by computer people for professional status. Despite previous experiences, most managers should learn to work with computers and to regard them as information tools such as typewriters or automated filing systems.

Flexibility and Responsiveness

When we are dealing with a network of distributed minicomputers, the flexibility of reconfiguration is unmatched by a centralized system. If an organization were to grow rapidly, the addition of minicomputers to the network should prove no problem; conversely, if volume were to decrease, the ability to shrink the system is available. Another important advantage is the capability to shift the workload among the many components within the system. This would become extremely useful when one part of the organization is placed under peak load conditions.

This distributed concept places strong emphasis on having functional management involved in and responsible for the component modules of the system. These modules must be adapted to the needs and requirements of both the organization and the functional unit. Isolated problems regarding an information system can be corrected rapidly without the need to interfere with the other sectors as long as the information required by the network is provided. Rigid adherence to standard rules found in centralized systems is not required, and operations can be modified according to local conditions. This is not to say that there are no systemwide rules, but that the adaptability of such a distributed system makes it easily responsive to the individual working with the computer.

The design and implementation of a computer information system is usually traumatic, time-consuming, and expensive. By being broken down into modules, however, the project becomes manageable. Installation and implementation can be made on a phased approach far more easily than with present centrally located systems.[13] The ultimate effect is that the emphasis is shifted from the information system to the persons responsible for it. This will permit the system to be structured to the individual rather than attempting to restructure the individual. Responsiveness should breed acceptance, which leads to involvement, a necessary prerequisite for a management information system.

Fault Tolerance

Many managers are reluctant to depend fully on an information system because of the possibility of system crashes, or periods in which the computer is not functioning. Data-processing managers speak with pride of 99 percent "up time" and tend to ignore the effect of the 1 percent "down time." Reliability and dependability are of utmost concern to the functional manager who must speak in terms of 100 percent up time, regardless of the air-conditioning or other factors. It is interesting to realize that in computer operations, the standard method for providing backup support is additional hardware, if the enormous expense is justified. Otherwise, the goal is to have the system back in operation as soon as possible by quickly isolating and correcting the problem.

Von Neuman noted that nature deals with malfunctions by making their effect as unimportant as possible and then applying correctives, if needed, at leisure.[14] The distributed minicomputer network permits this ability to "fail soft" rather than to crash. In this way performance will be degraded, but operations can continue while the problem is corrected. Through wise architecture, the system will have the ability to bypass the faulty component and permit the workload to be shared by the remaining members of the network. In addition, the cost of maintaining back-up hardware is well below that of a dual central system. The old adage, "Don't put all your eggs in one basket," is appropriate when looking at the increased potential for total disruption when computer operations are centralized rather than distributed among diverse locations.

Maintenance itself is far easier in a distributed minicomputer system. The problem is automatically isolated into a definitive sector, regardless of hardware or software culpability. Modularizing is becoming universally accepted as preplanned maintenance engineering: witness the television manufacturers' "works in a drawer" and snap-in integrated circuitry. The same concept is applicable to a distributed information system that has been presegmented. Finally, the absolute cost associated with a minicomputer permits an economic supply of spare processors to be kept and plugged into the network as needed (similar to an auto repair shop's inventory of spare parts).

Security has become an overwhelming and insolvable problem in the central computer system. It is further argued that multiaccess computing, especially through the use of remote terminals, has added to the exposure. If the same logic were then to be employed, a distributed network would appear to be extremely vulnerable. However, on the contrary, it appears that there are significant security advantages to the distributed network. All information is not in one location requiring multiple break points, either in the form of multiple physical incursions or a single intrusion accompanied by a breaking of the communication system followed by subsequent remote entries. The independence and capability of each distributed processing unit permits better continued monitoring of sensitive files and poses greater difficulty in obtaining access.

Architecture

Two basic types of structures are being used to implement the distributed network concept: a communications ring or a central controller. Employing a central controller requires that all communication first must be made to a central minicomputer (or set of minicomputers) before messages can be routed to the appropriate network member. The effect is to create a central decision point, facilitating workload distribution and resource sharing but exposing the system to single-point vulnerability. The alternative system approach utilizes a number of minicomputers connected to a single transmission line in a ring configuration. The minicomputers are hooked into this ring through an interface which itself is a minicomputer responsible for preparing messages for the communication line and being able to recognize and pull correctly addressed messages off the line. This type of system is capable of bypassing a malfunctioning unit without disrupting operations within the rest of the network. In addition, it is possible to distribute the important software programs responsible for organizational operations in a pre-engineered manner to preclude a system crash.[15] The advantages of the ring approach must be balanced against its increased cost in implementation.

The data bases also must be organized around a functionally distributed concept. Only the information that is required at a particular level of management should be stored or made available at that level. Because of the need for communications between minicomputers and segmented data bases, however, some common form of data description and data manipulation must be employed on a systemwide basis. Similar problems are being encountered in the multiaccess environment, and a great deal of research is being conducted in this area. The two concepts applicable to segmented data bases are vertical distribution and horizontal distribution. As their names imply, they are fundamentally different in emphasis, with the vertical concept based upon building blocks and the horizontal based on functional equalizations and sharing of responsibilities.

Richard Sprague and others have proposed an interesting hybrid type of information system, the hierarchical network.[16,17] In this concept an organization's needs are divided into multiple levels, which correspond to differing levels of computer support. The lowest level is the user level and only the appropriate computing power is supplied. But this level is connected through a network to the next higher level and its associated information system. This continues until the top level is connected to the system, thereby creating a totally integrated network. As one moves up the levels, the machine size increases while the need for distribution decreases. While the hierarchical network must be transparent to the users, functional responsibilities must match the computing power supplied at the appropriate level. In operation, this information net would basically consist of a network of minicomputers tied into a large central computing complex.

Conclusions

Although this article has been highly positive on the capabilities of a distributed network of minicomputers, it is obvious that it has certain limitations. There are times when massive computing power is needed, and the network is simply unable to provide such support.[18] However, I feel that this is the exception rather than the rule with regard to business information systems. As technology continues to improve the cost-performance ratios of memory and peripheral devices to the extent presently available in the processor field, the advantages associated with the new approach will increase significantly. At present, it is debatable whether the distributed minicomputer network is superior technically, operationally, and economically to the multiaccess centralized information system. In the future, there will be no debate!

The on-line integrated management information systems proposed by many authors are nothing more than a dream. In fact, the sophistication of today's computer systems is highly suspect and their usefulness to the organization is questionable. A study of more than fifty large corporations found that most of their computing systems were deliberately not both complex and critical.[19] The integration of data files is normally nothing more than the sharing of a particular storage device by more than one file. Very few companies classify their information system as critical, indicating that failure would not disrupt operations.

It is obvious that the computer world has a long way to go before it will achieve its mission of providing a reliable and efficient total information system to the organization. I believe that this failure can be attributed to the attitudes of systems architects or analysts who feel that in order to design and implement an information system, the structure and method of present operations must be changed. The centralization psychology probably has contributed more to this situation than any other single element. Distributed minicomputer systems provide a concept that permits an information system to fit into an organization rather than thrust itself upon it. And seduction is always more enjoyable and rewarding than rape!

REFERENCES

1. Captain Grace Hopper, private interview (Washington, D.C., September 1973).

2. John Dearden, "MIS is a Mirage," *Harvard Business Review* (January-February 1972), p. 93.

3. Bernhard Schwab, "The Economics of Sharing Computers," *Harvard Business Review* (September-October 1968), pp. 62-63.

4. Robert M. McClure, "Trade-off Between Software and Machine Size," Sixteenth Annual EDP Conference, American Management Association (February 24, 1970), p. 5.

5. R. F. Littrell, "The Costs of Computer Systems: Part I: Economies of Scale in Hardware Systems for Commercial and Scientific Computing," (Raleigh, North Carolina, unpublished).

6. Ibid., p. 2.

7. C. Gordon Bell; Robert Chen, and Satish Rege, "Effect of Technology on Near Term Computer Structures," *Computer* (March-April 1972), pp. 31-32.

8. Dearden, op. cit., pp. 95-96.

9. Henry Oswald, "Maxi-empire, No! Mini-empire, Maybe?" *Infosystems* (June 1972).

10. Fred Gruenberger and David Babcock, "Speaking of Minis," *Datamation* (July 1973), p. 58.

11. Nathan Grier Parke III, "Increasing Popularity of Modular Design," American Management Association, Briefing Session Number 6377-01.

12. Parke, op. cit., p. 4.

13. Christopher B. Newport, "Maturing Mini-Computers," *Honeywell Computer Journal*, pp. 30-35.

14. Richard G. Canning, "In Your Future: Distributed Systems," *EDP Analyzer* (August 1973), p. 9.

15. Roy Brun, "Researchers Pioneer Application Advances with Minis," *Infosystems* (October 1973), pp. 45-46.

16. Canning, op. cit., p. 2.

17. James C. Emery, "Problems and Promises of Regional Computing," *Datamation* (August 1973), pp. 56-57.

18. Wallace B. Riley, "Minicomputer Networks—A Challenge to Maxicomputers?" *Electronics Magazine* (March 29, 1971), p. 60.

19. "Systems Fragility in Complex Data Bases," (private working paper, McCaffrey, Seligman, and Von Simson, Inc.; used with permission).