A MARKOVIAN STATE-SPACE FRAMEWORK FOR INTEGRATING FLEXIBILITY INTO SPACE SYSTEM DESIGN DECISIONS

A Thesis
Presented to
The Academic Faculty

by

Jarret M. Lafleur

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Aerospace Engineering

Georgia Institute of Technology
May 2012

A MARKOVIAN STATE-SPACE FRAMEWORK FOR INTEGRATING FLEXIBILITY INTO SPACE SYSTEM DESIGN DECISIONS

Approved by:

Dr. Joseph H. Saleh, Advisor
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Robert D. Braun
School of Aerospace Engineering
Georgia Institute of Technology

Dr. John-Paul B. Clarke
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Vitali V. Volovoi
School of Aerospace Engineering
Georgia Institute of Technology

Mr. John F. Connolly
Exploration Missions and Systems Office
NASA Johnson Space Center

Date Approved: December 14, 2011

To Gerard, Sheryl, and Justin Lafleur

ventī sub ālīs meīs

and to Martin and Sharkey Tahakjian
and Alcide and Dolores Lafleur

doctorēs multō antĕquam nātus sum

ACKNOWLEDGEMENTS

My past several years at Georgia Tech have been a dream come true. Without fail, each year has been more fruitful, enriching, and exciting than the preceding one; and the things I have seen, the people I have met, and the activities of which I have had the privilege to be a part have exceeded my greatest expectations many times over. I have spent no small amount of time wondering how this has happened, and have concluded that either I have been consistently lucky in my stochastic experience of life, or I have been surrounded by extraordinary people who have consistently cared to guide me in the right direction no matter what uncertain future events unfolded. All things considered, I think the latter is the more likely explanation. While no list of acknowledgements is ever complete, I know that without the following people, this thesis would have taken a course other than (and quite inferior to) the one you see on the pages that follow.

First and foremost, I could not have asked for a more honest, principled, thoughtful, guiding, curiously poetic, and overall *awesome* advisor than **Dr. Joseph Saleh.** My time with Joe has been truly locupletative (Joe, I hope this is a word for which even *you* need to consult a dictionary!). Since I met Joe in 2007, we have been through thick and thin, and surely seen the best and the worst of each other. No thanks is enough to express my appreciation for his guidance over these years, and I can only hope this and my future accomplishments will come close to repaying my debt of gratitude.

I owe great thanks as well to the members of my thesis advisory committee, **Dr. Robert Braun**, **Dr. John-Paul Clarke**, **Dr. Vitali Volovoi**, and **Mr. John Connolly**. I assembled this committee knowing full well how high its standards would be, and it has been a pleasure to work with and receive the guidance of each one of these gentlemen.

I also owe thanks to my former undergraduate advisor, **Dr. John Olds**, who offered me my very first space systems engineering research experience. I shudder to think where I would be without that first opportunity so long ago. I must also not forget **Dr. Jechiel Jagoda** for his unfailing ability to make things happen and provide objective words of advice during both my undergraduate and graduate years.

Over the past eight years, I have grown as a professional by leaps and bounds through the work experience I have gained at NASA Johnson Space Center, especially in the Flight Mechanics and Trajectory Design Branch. I particularly need to thank **Mr. Christopher Cerimele**, my supervisor and mentor on six occasions since 2004 and someone who has always been open to whatever wild ideas I came up with next (including the idea of applying this thesis to a NASA application). I also need to thank **Mr. Bret Drake, Ms. Vickie Gutierrez, Mr. John Stelly,** and colleagues **Dan Matz, Jeff Gutkowski,** and **John Christian** for their many valuable inputs. In addition, I cannot convey enough thanks to **Mark Orr, Juan Senent,** and **Rick Rohan** for helping me convert my MATLAB MDP computer code, which would have taken several decades to run, into parallelizable Fortran syntax; thanks to these three gentlemen, I am relieved to be earning my terminal degree before my retirement checks.

Furthermore, the past few years would not have been the same without my graduate student colleagues in the Space Systems Design Lab. I especially need to thank **Joy Brathwaite** (who suffered the unspeakable misfortune of having a desk next to mine for four years!), **Jean-François Castet**, **Greg Dubos**, **Mike Grant, Brad Steinfeldt,**

George Burdell, Demyan Lantukh, and **Gregory Lantoine** for the cumulative hours of their lives they willingly (or perhaps reluctantly but politely) sacrificed as sounding boards for my ideas during the development of this thesis. I owe similar thanks to **Dr. Brian German** and **Matt Daskilewicz** in the Aerospace Systems Design Lab.

In addition, several friends have been exceptionally supportive over the past several years, among them **Gaurav Nagle, Masa Yano, Jonathan Sharma, Charity Duke, Brittany Williamson,** and of course, **Jason LeMaire**. I also must thank **Pat Biltgen** for his sage advice and support ever since my first day as a Georgia Tech student.

I also must thank **Mrs. Susan Mosher**, my calculus teacher from years ago for (as I recall) first planting Georgia Tech in my mind. I owe thanks, as well, to many other teachers in my home town of Burrillville – many more than I have the latitude to list here – for preparing me for the academic journey ahead. I never imagined how far it would go! (I must, however, specifically thank **Mr. Henry Coutu**, or as his Latin students know him, “Magister”, for checking the accuracy of my Latin for this thesis’ dedication.)

Several organizations have also generously funded my several years of education, and I would like to particularly acknowledge the National Science Foundation, Air Force Office of Scientific Research, Astronaut Scholarship Foundation, Ocean State Power / Burrillville Scholarship Foundation, Levy Memorial Scholarship Foundation, and the Georgia Institute of Technology for their support.

Finally, my greatest thanks goes to my family, **Gerard, Sheryl,** and **Justin Lafleur**. Without their unconditional love and support, neither this thesis nor I would have flown very far at all.

TABLE OF CONTENTS

Page

ACKNOWLEDGEMENTS iv

LIST OF TABLES xii

LIST OF SYMBOLS xix

LIST OF ABBREVIATIONS xxii

SUMMARY xxv

CHAPTER 1: Introduction: Flexibility in Space Systems 1

- 1.1. Uncertainty in Modern Space Missions 4
- 1.2. Flexibility Defined 5
  - 1.2.1. Flexibility in the Context of Optimization and Robustness 5
  - 1.2.2. Observed and Observable Flexibility 7
- 1.3. Recent Examples from Industry and Government 7
  - 1.3.1. Exploration Systems Architecture Study (NASA) 7
  - 1.3.2. System F6 (DARPA) 9
  - 1.3.3. Flexible-Path Human Space Exploration (NASA) 10
- 1.4. Flexibility in Space System Design Decisions 12

CHAPTER 2: Review of Literature on Flexibility 13

- 2.1. Early Economic Notions of Flexibility 13
- 2.2. The Two-Period State-Centric Notion of Flexibility 17
- 2.3. Flexible Manufacturing Systems 19
- 2.4. Examples from Aerospace Engineering Academia 21
- 2.5. Gaps in the Flexibility Literature 28

CHAPTER 3: Objectives and Contributions 31

CHAPTER 4: Theoretical Basis for a Markovian State-Space Flexibility Framework 35

- 4.1. Step 1: Define Configuration Options and the Cost Transition Matrix 36
  - 4.1.1. Defining the Configuration Space 36

4.1.2. Defining the Cost Transition Matrix 39
4.1.3. Analyzing the Cost Transition Matrices 44
4.1.4. Limitations of Cost-Only Considerations 59
4.2. Step 2: Define Markovian Demand Environment Evolution Probabilities 61
4.2.1. Definition of the Demand Environment 62
4.2.2. The Markovian Stochastic Model 63
4.3. Step 3: Define State-Dependent Performance Matrix 67
4.4. Step 4: Decision Support Analysis 69
4.4.1. Find Pareto-Optimal “Open-Loop” Paths 72
4.4.2. Find Pareto-Optimal “Closed-Loop” Policies 79
4.5. Step 5: Implications for Initial System Selection 95
4.5.1. Implications based on the Expected-Value Pareto Frontier 95
4.5.2. Accounting for Non-Expected-Value Objectives 98
4.5.3. Flexibility, Entropy, and Policy 103
4.6. Summary 111
CHAPTER 5: Application to Distributed- or Multi-Payload Satellite Design Decisions 113
5.1. Step 1: Define Configuration Options and the Cost Transition Matrix 114
5.1.1. Defining the Configuration Space 115
5.1.2. Defining the Cost Transition Matrix 117
5.1.3. Analyzing the Cost Transition Matrices 121
5.2. Step 2: Define Markovian Demand Environment Evolution 125
5.3. Step 3: Define State-Dependent Performance Matrix 128
5.4. Step 4: Decision Support Analysis 129
5.4.1. Find Pareto-Optimal “Open-Loop” Paths 130
5.4.2. Find Pareto-Optimal “Closed-Loop” Policies 136
5.5. Step 5: Implications for Initial System Selection 145
5.5.1. Implications based on the Expected-Value Pareto Frontier 145
5.5.2. Accounting for Non-Expected-Value Objectives 147

5.6. Summary 151

CHAPTER 6: Advanced Application: NASA Human Exploration Architecture Decision-Making....... 154

6.1. Step 1: Define Configuration Options and the Cost Transition Matrix.. 159

6.1.1. Defining the Configuration Space 159

6.1.2. Defining the Cost Transition Matrix.. 167

6.1.3. Analyzing the Cost Transition Matrix.. 173

6.2. Step 2: Define Markovian Demand Environment Evolution 190

6.3. Step 3: Define State-Dependent Performance Matrix.. 195

6.3.1. Selecting the Performance Metric... 196

6.3.2. Populating the Performance Matrix 197

6.3.3. Populating the Boolean Demand Fulfillment Matrix 201

6.4. Step 4: Decision Support Analysis 202

6.4.1. Preclusion of Open-Loop Path Analysis 204

6.4.2. Timeline Assumptions and Expanded Probability Definitions.. 204

6.4.3. Computational Resources and Implementation 208

6.4.4. Primary Results: The Potential of an Unconstrained Per-Period Budget 210

6.4.5. Implications of a Per-Period Budget Constraint 222

6.5. Step 5: Implications for Initial System Selection 226

6.5.1. Implications based on the Expected-Value Pareto Frontier.. 226

6.5.2. Accounting for Non-Expected-Value Objectives 234

6.6. Summary 239

CHAPTER 7: Conclusion and Avenues for Future Work.. 243

7.1. Summary 243

7.2. Contributions 248

7.3. Avenues for Future Work.. 249

7.3.1. Multi-Period Expansion of the Φ Transition Metric 250

7.3.2. Decision-Makers with Authority Limitations.. 251

7.3.3. Robustness to Changing Decision-Makers... 252

7.3.4. Additional Theory and Algorithm Investigations 253
7.3.5. Additional Application-Specific Questions 253
7.4. Closing Remarks 255
APPENDIX A: A Heuristic Method for Identifying Concave Pareto Frontiers in Multi-Objective Dynamic Programming Problems 258
A.1. Introduction 258
A.2. Application to Dynamic Programming 261
A.3. A Heuristic Method for Identifying Pareto-Optimal Solutions 264
A.4. Example Application 267
A.4.1. Full-Factorial Pareto Frontier 268
A.4.2. Heuristically-Generated Pareto Frontier 270
A.5. Summary 273
APPENDIX B: Transition Cost Model for Human Space Exploration Configurations 276
B.1. Development and First-Period Production Costs 277
B.1.1. Total Development and First-Period Production Cost Estimation 278
B.1.2. Development and First-Period Production Cost Estimation Logic 282
B.2. Program Management and Systems Engineering Costs 283
B.3. Recurring Production Costs 284
B.4. Ground and Mission Operations Costs 284
B.5. Shutdown Costs 285
B.5.1. Retirement Costs 285
B.5.2. Termination Liability Costs 286
B.6. Validation 287
B.7. Summary 290
APPENDIX C: Expert Inputs for Human Space Exploration Stochastic System Decision Modeling 292
C.1. Survey Description 292
C.2. Results 302
C.2.1. Participant Experience 302

C.2.2. Figures of Merit 305

C.2.3. Markov Chain Estimates 309

C.3. Summary 329

REFERENCES 331

VĪTA 346

LIST OF TABLES

Page

Table 1. Summary of Flexibility-Related Literature in Aerospace Academia. 27

Table 2. Example Form of a Morphological Matrix. 39

Table 3. Example Format for a Cost Transition Matrix *C*. 40

Table 4. Configuration Definitions for Satellite Example. 45

Table 5. C_{dev} for Satellite Example. Costs are in millions of dollars. 45

Table 6. C_{rec} for Satellite Example. Costs are in millions of dollars. 46

Table 7. *C* for Satellite Example. Costs are in millions of dollars. 46

Table 8. Development Cost Savings Matrix. 48

Table 9. Demand Environment Definitions for Satellite Example. 63

Table 10. Sample Markov Chain Transition Matrix for the Satellite Example Application. 64

Table 11. *U* for Satellite Example. 68

Table 12. Transition Probabilities for Notional MDP Example. 85

Table 13. Current-Period Reward Function for Notional MDP Example. 85

Table 14. Backward Induction Calculations for Notional MDP Example. 86

Table 15. Optimal Actions and Expected Rewards-to-Go for Notional MDP Example. 87

Table 16. Pareto-Optimal Policy #14 for Satellite Example. 93

Table 17. Anticipatory Policy for Satellite Example. 95

Table 18. Low-Entropy-Rate Demand Markov Chain Transition Matrix (H′=0.67 bits). 104

Table 19. Medium-Entropy-Rate Demand Markov Chain Transition Matrix (H′=1.26 bits). 105

Table 20. Special Deterministic Markov Chain Transition Matrix (H′=0 bits). 107

Table 21. Special Uniform Random Markov Chain Transition Matrix (H′=2.58 bits). 108

Table 22. Assumed payload characteristics for example design. [98]-[102] 115

Table 23. Development cost transition matrix, C_{dev} (data in $FY08M). 118

Table 24. Recurring cost transition matrix, C_{rec} (data in $FY08M). 120

Table 25. Total cost transition matrix, *C* (data in $FY08M). 120

Table 26. Assumed demand environment transition probability matrix. 127

Table 27. Performance matrix quantifying number of demanded services performed in a time period. 129

Table 28. Pareto-optimal policy #3. 141

Table 29. Anticipatory Policy.. 144

Table 30. Morphological Matrix for the Human Spaceflight Architecture Application. 160

Table 31. Architecture Definitions for the Human Space Exploration Application. 162

Table 32. Architectural Projection Two-Dimensional Coordinates. 164

Table 33. Time-Independent Architecture Transition Costs (in $FY11B). 178

Table 34. Markov chain for median inputs and Δt = 2 years, for condition that demand is fulfilled. 191

Table 35. Markov chain for median inputs and Δt = 2 years, for condition that demand is *not* fulfilled. 192

Table 36. Demand Environment Requirement Definitions for the Space Exploration Application. 199

Table 37. Performance matrix quantifying number of missions flown to the demanded destination. 201

Table 38. Boolean demand fulfillment matrix. 202

Table 39. Mean relative schedule slippage as a function of TRL, from the model of Ref. [110]. 206

Table 40. Target Coordinates for Example Application. 267

Table 41. Target Values for Example Application. 268

Table 42. Comparison of Pareto Frontier Search Methods in the Example Application. 272

Table 43. Coefficients for Eqs. (B1) and (B2) as a function of Vehicle/Hardware Class. 279

Table 44. Architectural Component Weight Assumptions and DDT&E and TFU Cost Outputs. 281

Table 45. Operations cost model inputs and outputs for each operations architecture. 285

Table 46. Discrepancies for transition cost model with a best-fit single scaling factor. 288

Table 47. Discrepancies for transition cost model with two best-fit scaling factors. 289

Table 48. Survey Invitee Affiliations. 293

Table 49. Survey Activity Timeline. 294

Table 50. Markov chain for median non-dominated experts, for condition that demand is fulfilled. 324

Table 51. Markov chain for median non-dominated expert, for condition that demand is *not* fulfilled. 325

LIST OF FIGURES

Page

Figure 1. An early concept for Skylab (left) [9], and the Skylab configuration as launched (right) [10]. 2

Figure 2. Optimization, Robustness, and Flexibility Notionally Compared. 6

Figure 3. Sample summary of figure of merit ratings for concepts in the ESAS report. 9

Figure 4. Possible "Flexible Path" Mission Sequence. [34] 11

Figure 5. Possible Flexible Path Destination Sequences. [34] 11

Figure 6. Marginal Cost and Average Cost Curves for an Inflexible Plant and Flexible Plant [37] 15

Figure 7. Visualization of Koopmans' Partitioning of Opportunities. [38] 16

Figure 8. Graphical interpretation of Jones and Ostroy's interpretation [51] of flexibility of positions. 18

Figure 9. Example setup of a small flexible manufacturing system. [59] 19

Figure 10. The deterministic flexibility example of Tempelmeier and Tetzlaff [61]. 21

Figure 11. Critical components for decision frameworks addressing gaps in present flexibility literature. 32

Figure 12. Five major steps of this thesis' framework. 36

Figure 13. Example configuration state space in which five systems are defined by two design variables. 38

Figure 14. The M-N space. 42

Figure 15. Switching Cost vs. Initial Cost from Config. 1 (S_1) for the satellite example. 47

Figure 16. Configuration state space for the satellite example. 49

Figure 17. Available configuration transitions for $400, 550, 700, and 850 million budgets. 51

Figure 18. Number of available transitions (Φ) for $400, 550, 700, and 850 million budgets 52

Figure 19. Available configuration transitions as a function of available per-period budget. 54

Figure 20. Available transitions for Configs. 3 and 4 vs. budget, illustrating "flexibility reversal". 54

Figure 21. Number of transitions gained when budget on the x-axis is raised by $250 million 56

Figure 22. Transition Elasticity for a forward difference of $250 million 57

Figure 23. Probabilistic comparison of the transitions available to Configs. 3 and 4 59

Figure 24. Visualization of the Markov Chain (left) and Stationary Distribution (right) of Table 10 66

Figure 25. Visualization of Environments, Configurations, and Actions over Multiple Time Periods. 72

Figure 26. Evolution of configuration path [3 3 2 1]....... 76

Figure 27. Trade between total demanded services performed and total cost for all open-loop paths. 78

Figure 28. Depiction of state space, actions, rewards, and transition probabilities for the MDP example. . 85

Figure 29. Mapping of Flexibility Framework Components into a Markov Decision Process. 87

Figure 30. “Spindle” of $N \times K$ Total States. 89

Figure 31. Trade between total demanded satellites available and total cost for MDP policy solutions. 92

Figure 32. Evolution of states and objectives for Pareto-optimal policy #14 (defined in Table 16). 93

Figure 33. Initial configurations as a function of expected Pareto-optimal path or policy total cost. 97

Figure 34. Number of Available Transitions for Pareto-Optimal Initial Configurations....... 98

Figure 35. Multivariate plot of multi-objective genetic algorithm policy results....... 101

Figure 36. Multivariate plot of genetic algorithm results with constraints imposed. 103

Figure 37. Φ of Initial Configurations of Optimal Policies for varying Environment Entropy Rates. 106

Figure 38. Trade between satellites available and cost for solutions subject to deterministic demand. 107

Figure 39. Trade between satellites available and cost for solutions subject to uniform random demand 109

Figure 40. Possible system configurations. 116

Figure 41. Planning periods and decision points for the distributed-payload satellite example....... 117

Figure 42. Switching cost vs. initial cost from Config. 1. 122

Figure 43. Available transitions for three example 30-month budgets....... 123

Figure 44. Available configuration transitions as a function of the available 30-month budget. 124

Figure 45. Visualization of the demand environment Markov chain described by Table 26....... 128

Figure 46. Evolution of path [4 9 15 15 1], representative of incremental capability buildup. 133

Figure 47. Trade between total demanded services performed and total cost for all open-loop paths. 135

Figure 48. Evolution of configuration path [1 11 11 11 1], a Pareto-optimal path. 136

Figure 49. “Spindle” of Total States. 138

Figure 50. Trade between total demanded services performed and total cost for MDP policy solutions. . 140

Figure 51. Evolution of states and objectives for Pareto-optimal policy #3 (defined in Table 28). 140

Figure 52. Evolution of states and objectives for an anticipatory policy (defined in Table 29)....... 143

Figure 53. Initial configurations for Pareto-optimal paths and policies vs. expected total cost. 147

Figure 54. Multivariate plot of multi-objective genetic algorithm policy results. 150
Figure 55. Multivariate plot of genetic algorithm policy results with constraints imposed. 151
Figure 56. Architectural Projection for the Human Space Exploration Application 164
Figure 57. Notional Sequence of Configurations (Architecture Triplets). 167
Figure 58. Transition Cost Model for Human Space Exploration Configurations. 168
Figure 59. Illustrations of Disallowed Transition Rules. 171
Figure 60. Cost Transition Matrix Visualization (left) and Distribution of Costs (right). 172
Figure 61. Available configuration transitions for a $500 million 2-year budget. 175
Figure 62. Available configuration transitions for a $1 billion 2-year budget. 176
Figure 63. Available configuration transitions for a $2 billion 2-year budget. 176
Figure 64. Available configuration transitions for a $5 billion 2-year budget. 176
Figure 65. Available configuration transitions for a $12.9 billion 2-year budget. 177
Figure 66. Switching cost vs. initial cost from Architecture 1. 179
Figure 67. Available architecture transitions for $0.5, 1, 5, 10, 15, and 20 billion budgets. 181
Figure 68. Available architecture transitions vs. available time-independent budget. 183
Figure 69. Number of available transitions vs. available budget over two years. 185
Figure 70. Distributions of available transitions for $12.9B per period, by operations architecture. 187
Figure 71. Distributions of available transitions for $25B per period, by operations architecture.. 188
Figure 72. Example of a "Flexibility Reversal" in the NASA Application. 189
Figure 73. Visualization of Markov chain median for condition that demand is fulfilled, Δt = 8 years. 194
Figure 74. Visualization of Markov median for the condition that demand is *not* fulfilled, Δt = 8 years. 195
Figure 75. Visualization of Decisions and Environment Evolution over Time for NASA Application. 203
Figure 76. Flowchart describing the total-state-to-total-state probability used in NASA application. 208
Figure 77. Evolution of states and objectives for the anticipatory reference policy. 213
Figure 78. Trade between missions to demanded destinations and cost for MDP policy solutions. 215
Figure 79. Evolution of states and objectives for policy W0.125-N1. 216
Figure 80. Evolution of states and objectives for policy W0.050-N2. 220
Figure 81. Evolution of states and objectives for policy W0.600-N∞. 222

Figure 82. Trade between missions and cost subject to a $12.9 billion budget constraint........................ 224

Figure 83. Evolution of states and objectives for max-performance policy given cost constraint. 225

Figure 84. Initial architectures of configurations implied by Pareto-optimal policies vs. total cost.......... 228

Figure 85. Initial architectures of Pareto-optimal policies given $12.9 billion per-period cost constraint. 230

Figure 86. Trade between missions and cost for constrained solutions subject to alternate demand. 232

Figure 87. Initial architectures of cost-constrained Pareto-optimal policies subject to alternate demand. 233

Figure 88. Multivariate plot of multi-objective genetic algorithm policy results.................................... 237

Figure 89. Multivariate plot of genetic algorithm results with 90th percentile cost constraint imposed. ... 238

Figure 90. Example of an objective space with a concave frontier of nondominated points. 260

Figure 91. Nondominated solutions identified via aggregate function of Eq. (A1) for varying n............ 260

Figure 92. Contours of the aggregate objective J as n increases. Γ_1 and Γ_2 are weighted equally. 260

Figure 93. Flowchart for this appendix's heuristic multi-objective optimization algorithm..................... 266

Figure 94. Performance trades and the full frontier for the example application. 269

Figure 95. Graphical representation of the targets listed in Table 40 and Table 41. 270

Figure 96. Comparison of Pareto frontiers generated using heuristic and full factorial methods. 271

Figure 97. Normalized and discretized development and first-period production cost distribution.......... 280

Figure 98. Summary of Development and First-Period Production Cost Estimation Logic. 283

Figure 99. Histogram of absolute value of discrepancies from Table 46. 288

Figure 100. Histogram of absolute value of discrepancies from Table 47. 290

Figure 101. Geographic view of invitee NASA center affiliations. 293

Figure 102. Screen Shot of Main Survey Website (http://www.flexibility.gatech.edu). 297

Figure 103. Screen Shot of Survey Submission Website. 297

Figure 104. Consent Form (Worksheet #1) from Survey Excel File.. 298

Figure 105. Figures of Merit Section (Worksheet #2) from Survey Excel File. 299

Figure 106. Demand Evolution Section (Worksheet #3) from Survey Excel File. 300

Figure 107. Final Submission Instructions (Worksheet #4) from Survey Excel File 301

Figure 108. Dialog Box for Part I of Worksheet #3. 301

Figure 109. Dialog Box for Part II of Worksheet #3.. 301

Figure 110. Distributions of Participant Years of Experience. 303

Figure 111. Multivariate Plot Illustrating Years of Experience Correlations. 305

Figure 112. Aggregate Figure of Merit Rating Results. 308

Figure 113. Summary of Figure of Merit Response Interquartile Ranges and Medians.......................... 309

Figure 114. Summary of transition matrix and time responses, given demand is fulfilled. 311

Figure 115. Summary of transition matrix and time responses, given demand is not fulfilled. 312

Figure 116. Aggregated Interquartile Range Statistics from Figure 114 and Figure 115. 313

Figure 117. Visualization of Markov chain mean for condition that demand is fulfilled, Δt = 8 yrs. 320

Figure 118. Visualization of Markov chain mean for condition that demand is *not* fulfilled, Δt = 8 yrs... 321

Figure 119. Stationary distribution of Markov chain, for condition that demand is *always* fulfilled. 323

Figure 120. Stationary distribution of Markov chain, for condition that demand is *never* fulfilled. 323

Figure 121. Visualization of non-dominated expert model median given that demand is fulfilled........... 325

Figure 122. Visualization of non-dominated expert model median given that demand is *not* fulfilled..... 326

Figure 123. Stationary distribution of non-dominated expert median given demand is *always* fulfilled... 327

Figure 124. Stationary distribution of non-dominated expert median given demand is *never* fulfilled..... 327

LIST OF SYMBOLS

A_i	i^{th} architecture
a	decision-maker action
$a_{DDT\&E}$	DDT&E cost regression multiplier coefficient
a_{TFU}	TFU cost regression multiplier coefficient
b	budget level
$b(\tau)$	time-period-referenced schedule of budget levels
$b_{DDT\&E}$	DDT&E cost regression power coefficient
b_{TFU}	TFU cost regression power coefficient
C	total cost transition matrix
C_{dev}	development cost transition matrix
C_{devpct}	development cost savings matrix
C_{ops}	operations component of recurring cost transition matrix
C_{prod}	production component of recurring cost transition matrix
C_{rec}	recurring cost transition matrix
c_{ij}	element (i,j) of the total cost transition matrix
$c_{dev,ij}$	element (i,j) of the development cost transition matrix
$c_{devpct,ij}$	element (i,j) of the development cost savings matrix
$c_{ops,ij}$	element (i,j) of the operations cost transition matrix
$c_{prod,ij}$	element (i,j) of the production cost transition matrix
$c_{DDT\&E}$	component DDT&E cost estimate
c_{TFU}	component TFU cost estimate
$c_{ops,annual}$	annual architecture operations cost estimate
D	column vector of shutdown costs for system subcomponents
$E(\bullet)$	expected value operator
$E_{\Delta h,i}$	transition elasticity

F	vector-valued function with random-variable argument and deterministic value output
g	Wright model learning percent
H'	entropy rate
h	forward difference budget increment
J_τ	aggregate reward-to-go at the beginning of period τ
K	number of states in demand environment state space
LC	learning percent for TFU cost model
M	number of design variables required to define each system or configuration
N	number of candidate systems or configurations under consideration
N_{total}	total number of states that exist in the configuration state space
n	aggregate objective function power
P	Markov chain probability transition matrix
P_{ij}	element (i,j) of Markov chain probability transition matrix
$P(\bullet)$	probability operator
$P*$	intermediate matrix in continuous- to discrete-time Markov chain conversion
P_{CTMC}	probability transition matrix for continuous-time Markov chain
Q	column vector of per-unit production costs for system subcomponents
q_k	k^{th} element of vector Q
q_{TFU}	quantity of theoretical first production units
R	column vector of development costs for system subcomponents
S_i	i^{th} system of the set of candidate configurations
$s_{i,k}$	number of the k^{th} subcomponent in system S_i
T	number of time periods in the time horizon of interest
t	time
t_0	initial time
U	performance matrix
u_{ij}	element (i,j) of performance matrix
v	reference transition rate

v_{CTMC}	transition rate vector for continuous-time Markov chain
w_i	weight placed on the i^{th} objective
$w_{vehicle}$	system weight at sea level on Earth
x_k	k^{th} design variable
Y	demand environment state random variable
y	demand environment state deterministic instantiation
$\alpha^{\tau+}$	policy of actions adopted for all future states
β	discounting factor
Γ_i	normalized cumulative performance in terms of the i^{th} objective
γ_i	normalized per-period performance in terms of the i^{th} objective
Δt	time step of interest
η	aggregate per-period expected reward
Λ	set of possible actions
Ξ	set of possible total states
ξ	total state
π_j	long-run probability of existing in demand environment j
τ	discrete time index
$\Phi_i(b)$	number of transitions available from configuration i for budget b
Φ'	total number of transitions accessible in a T-period tree
Ω	number of per-period objectives

LIST OF ABBREVIATIONS

AAP	Apollo Applications Program
AC	Average Cost
ARC	Ames Research Center
ATLO	Assembly, Test, and Launch Operations
CCLM	Commercial Cargo Logistics Module
CCLV	Commercial Cargo Launch Vehicles
CITI	Collaborative Institutional Training Initiative
CLV	Crew Launch Vehicle
CNC	Computer Numerical Controlled
CSM	Command and Service Module
DARPA	Defense Advanced Research Projects Agency
DDT&E	Design, Development, Test, and Evaluation
DSEA	Difficulty Scale for Evolvability Analysis
EELV	Evolved Expendable Launch Vehicle
ESAS	Exploration Systems Architecture Study
EVA	Extravehicular Activity
F6	Future Fast, Flexible, Fractionated, Free-Flying Spacecraft united by Information eXchange
FMS	Flexible Manufacturing System
FY	Fiscal Year
GEO	Geosynchronous Earth Orbit
GRC	Glenn Research Center
GSE	Ground Support Equipment
GSFC	Goddard Space Flight Center
GT-FAST	Georgia Tech F6 Architecture Synthesis Tool
HAT	Human Spaceflight Architecture Team

HEFT	Human Exploration Framework Team
HLV	Heavy Lift Launch Vehicle
HQ	Headquarters
IACO	Integration, Assembly, and Checkout
IRB	Institutional Review Board
ISRU	In Situ Resource Utilization
JPL	Jet Propulsion Laboratory
JSC	Johnson Space Center
LaRC	Langley Research Center
LEO	Low-Earth Orbit
LOOS	Launch and Orbital Operations Support
MADM	Multi-Attribute Decision-Making
MC	Marginal Cost
MDP	Markov Decision Process
MF	Manufacturing Flexibility
MOCM	Mission Operations Cost Model
MRO	Mars Reconnaissance Orbiter
MPCV	Multi-Purpose Crew Vehicle
MSFC	Marshall Space Flight Center
NASA	National Aeronautics and Space Administration
NEO	Near-Earth Object
PL	Payload
PM	Program Management
SE&I	Systems Engineering and Integration
SMAD	*Space Mission Analysis and Design*
STO	System Test Operations
SVLCM	Spacecraft/Vehicle Level Cost Model
TDN	Time-Expanded Decision Network

TFU	Theoretical First Unit
TRL	Technology Readiness Level
U.S.	United States
VWFO	Value-Weighted Filtered Outdegree

SUMMARY

The past decades have seen the state of the art in aerospace system design progress from a scope of simple optimization to one including robustness, with the objective of permitting a single system to perform well even in off-nominal future environments. Integrating flexibility, or the capability to easily modify a system after it has been fielded in response to changing environments, into system design represents a further step forward. One challenge in accomplishing this rests in that the decision-maker must consider not only the present system design decision, but also sequential future design and operation decisions. Despite extensive interest in the topic, the state of the art in designing flexibility into aerospace systems, and particularly space systems, tends to be limited to analyses that are qualitative, deterministic, single-objective, and/or limited to consider a single future time period.

To address these gaps, this thesis develops a stochastic, multi-objective, and multi-period framework for integrating flexibility into space system design decisions. Central to the framework are five steps. First, system configuration options are identified and costs of switching from one configuration to another are compiled into a cost transition matrix. Second, probabilities that demand on the system will transition from one mission to another are compiled into a mission demand Markov chain. Third, one performance matrix for each design objective is populated to describe how well the identified system configurations perform in each of the identified mission demand environments. The fourth step employs multi-period decision analysis techniques, including Markov decision processes from the field of operations research, to find efficient paths and policies a decision-maker may follow. The final step examines the implications of these paths and policies for the primary goal of informing initial system selection.

Overall, this thesis unifies state-centric concepts of flexibility from economics and engineering literature with sequential decision-making techniques from operations research. The end objective of this thesis' framework and its supporting tools is to enable selection of the next-generation space systems today, tailored to decision-maker budget and performance preferences, that will be best able to adapt and perform in a future of changing environments and requirements. Following extensive theoretical development, the framework and its steps are applied to space system planning problems of (1) DARPA-motivated multiple- or distributed-payload satellite selection and (2) NASA human space exploration architecture selection.

The point we wish to make is that in modern life, in economic, industrial, scientific and even political spheres, we are continually surrounded by multi-stage decision processes. Some of these we treat on the basis of experience, some we resolve by rule-of-thumb, and some are too complex for anything but an educated guess and a prayer.

Richard E. Bellman, Ph.D., 1957
Preface to *Dynamic Programming*

CHAPTER 1

INTRODUCTION: FLEXIBILITY IN SPACE SYSTEMS

In the late 1960s, as it was preparing to mount the first manned lunar mission, the nascent National Aeronautics and Space Administration (NASA) encountered its first major budget cuts. Between 1965 and 1970, the agency's annual budget was cut by over $1.5 billion, or 29% [1]. In a political environment focused on other national priorities and on lowering federal spending, support for future space exploration enterprises fell short of NASA's post-Apollo ambitions [2]. On January 4, 1970, it was announced that Apollo 20 would be cancelled [3]-[4], and on September 2, two more Apollo lunar missions were cancelled [5]. With development of a space shuttle years away from approval, NASA's Apollo Applications Program (AAP), formed in 1965 to develop alternative mission options using Apollo architectural components, was the only measure available to mitigate a disastrous gap in human spaceflight and loss of the 400,000-person Apollo workforce [6]-[8].

Thus was born the Skylab space station program. Skylab itself (see Figure 1), launched atop a Saturn V rocket in May 1973, was a modified S-IVB stage originally intended to fly as the upper stage of a Saturn IB launch vehicle [8]. Three-man crews were transported to the space station using an adapted Apollo command and service module (CSM) that incorporated 23 major modifications [7] and that launched on the Apollo-heritage Saturn IB. Plans even existed – and in one instance saved a mission from a premature end – to modify the CSM further for a rescue capability to accommodate five crew [6]-[7]. Over the course of eight months, nine crew launched to the station and accumulated a total of 513 crew-days in space [7], five times more than all previous U.S. spaceflight and providing a wealth of data on long-duration effects of spaceflight.

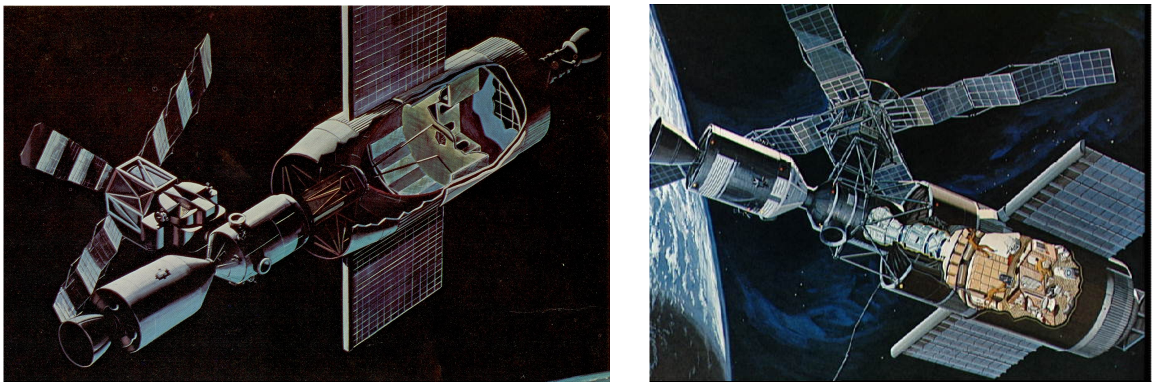

Figure 1. An early concept for Skylab which included a Lunar Module converted for use as the solar observatory (left) [9], and the Skylab configuration as launched (right) [10].

Importantly, Skylab and the Apollo Applications Program demonstrated the capability of the Apollo lunar architecture to be easily modified after it had been fielded in response to a changing environment and changing requirements. This property is what the present thesis will refer to as *flexibility*. However, Apollo's flexibility was largely accidental: The components of the Apollo architecture had been selected to meet the goal of landing a man on the Moon by 1970, with little emphasis on other objectives like flexibility. Of the 98 pages discussing possible Apollo architectures in the authoritative 1962 architecture decision document [11], only three pages are devoted to implications on growth potential – and these three pages indicate the chosen lunar orbit rendezvous architecture offered the *least* growth potential compared to others under consideration.

Over the four decades since Apollo, the world's civil and military space programs have given increasing emphasis to flexibility when designing new systems, but flexibility tends to remain an intangible and abstract concept to engineers. Many of the techniques used to evaluate this elusive property are **qualitative**, **subjective**, **deterministic**, **single-objective**, and/or **limited to consider a single time step in the future**. The question remains: How can engineers and decision-makers systematically, quantitatively, and objectively consider flexibility in the design of a new space system? How can the space system design community reduce future occurrences where flexibility is desirable but is unattainable, or where it exists but is accidental?

These questions drive the present thesis, which begins by investigating the state of the art of considering flexibility in space and other engineering system design decisions in order to substantiate the characterization of this art as described in the preceding paragraph. In the process of this investigation, it is found that, although disparate, some threads of thought on flexibility are common in the literature. These threads, which entail the construction of state spaces to describe a system's flexibility of movement in a two-period setting, form a foundation for the framework that this thesis presents.

The new framework that this thesis introduces consists of five practical steps intended for implementation by engineering systems analysts, the first three of which focus on defining and characterizing a set of state spaces representing system options and environment demands. The fourth step employs multi-period decision analysis techniques, including Markov decision processes from the field of operations research, to find efficient paths and policies a decision-maker may follow. The final step examines the implications of these paths and policies for initial system selection. The end product is a quantitative, stochastic, multi-objective, and multi-period framework for integrating flexibility into engineering system design decisions. Moreover, this thesis illustrates that not only is a state-centric notion of flexibility prevalent in the literature compatible with a comprehensive decision support framework, but that it is naturally adapted for use with Markov decision process solution techniques from the operations research community.

After theoretical development using a simple satellite system example, applications are illustrated using a fractionation-motivated multi-payload defense satellite example, as well as the substantially more complex example of NASA human space exploration architecture selection. These latter two applications in particular substantiate the relevance of this framework in informing decisions for problems of interest to the space industry today. Before proceeding further, however, it is necessary to establish a preliminary background on several related topics.

1.1. Uncertainty in Modern Space Missions

Flexibility in space systems design is relevant largely because of the uncertainties involved when planning space missions, whether the goal of the mission is as limited as low-resolution Earth observation or as ambitious as human exploration of Mars. These uncertainties can be divided into categories of downside risk and upside potential:

The foremost uncertainties in most system engineers' minds are generally those involving downside risk, or the possibilities of off-nominal situations causing undesirable consequences. These include risks of launch failure, component failure or degradation, physical or directed energy attack, funding cuts, cost growth, or decrease in satellite service demand.

Another important but less commonly considered form of uncertainty is upside potential, or the possibility for an off-nominal situation to present opportunities upon which a program can capitalize with desirable consequences. Examples include increases in satellite service demand, new initiatives and increases in funding, and unforeseen scientific opportunities.

While by definition these risks and opportunities are not predictable with certainty, some may be more probable than others, and some may entail greater consequences. Decisions made during design have the potential to mitigate or exacerbate such consequences when or if these events occur, and ideally a decision-maker will make the proper choices during design to allow the system to adequately respond to requirement or environment changes later. In the words of former and current Defense Advanced Research Projects Agency (DARPA) program managers Owen Brown and Paul Eremenko, "an uncertain future does not mean that we throw up our hands, and simply wait to react to future shocks. It does mean that we must explore a variety of potential futures, and create strategies and policies, as well as technical and architectural solutions that provide hedges for a variety of circumstances that could occur." [12]

1.2. Flexibility Defined

The Merriam-Webster Dictionary definition of flexibility is the “ready capability to adapt to new, different, or changing requirements.” [13] This thesis adopts a similar definition, namely that flexibility is **the capability to easily modify a system after it has been fielded in response to a changing environment or changing requirements** (cf. [14]). Central to this notion of flexibility are the conditions that (1) a system’s environment or requirements may change in the future and (2) the system can, to some degree, be modified to accommodate such change. This definition also includes the notion of ease of modification, which means that the effort required to effect a change (whether measured in dollars, manpower, or other resource-representative metrics) is also relevant to discussions of flexibility. These important properties of the flexibility definition will become more clearly defined throughout Chapters 2 and 3. Also, note that while some techniques developed in this thesis may apply to incorporating flexibility into system development phases (i.e., permitting system modifications prior to system fielding), this thesis focuses on modifications that are to be available *after* the system is fielded. Flexibility during the development process is another important area of work, described by Refs. [14] and [15] and often linked to the desire to preserve design freedom and delay cost commitment between design and manufacturing [16]-[19]. While Chapter 6 shows that approaches developed herein are extensible to the consideration of flexibility in the development process, the principal intent of these techniques is toward considering modifications to present or future fielded systems.

1.2.1. Flexibility in the Context of Optimization and Robustness

The definition of flexibility above may be enhanced with a graphical comparison to the more established engineering concepts of optimization and robustness. Figure 2 illustrates a helpful way of visualizing these concepts, with each concept shown in terms of a notional performance metric plotted against an environment (or requirement) variable:

Traditional optimization involves the minimization or maximization (as in Figure 2) of the performance metric assuming a system subject to the nominal operating environment. Off-nominal environments are not considered, and it is possible for performance to degrade significantly in these environments.

In contrast, a robust system is designed such that, when the system is exposed to an off-nominal environment, performance remains close to the nominal level. Robust design techniques, popularized by Taguchi in the 1980s [20]-[24], has been well explored over the past few decades (for helpful surveys on this topic, see Refs. [24] and [25]). By definition, however, a robust system cannot have a nominal performance better than the optimized case; and generally such a system will have a lower nominal performance.

A flexible system is distinguished by the fact that modifications can allow the system to effectively change its performance curve at the operator's discretion. If the system is in a particular configuration at time t_1 and the environment changes, the operator can choose to make a modification to the system (at some cost in resources) and achieve a new performance characteristic at time t_2. This dynamic behavior introduces a host of challenges in modeling and decision-making.

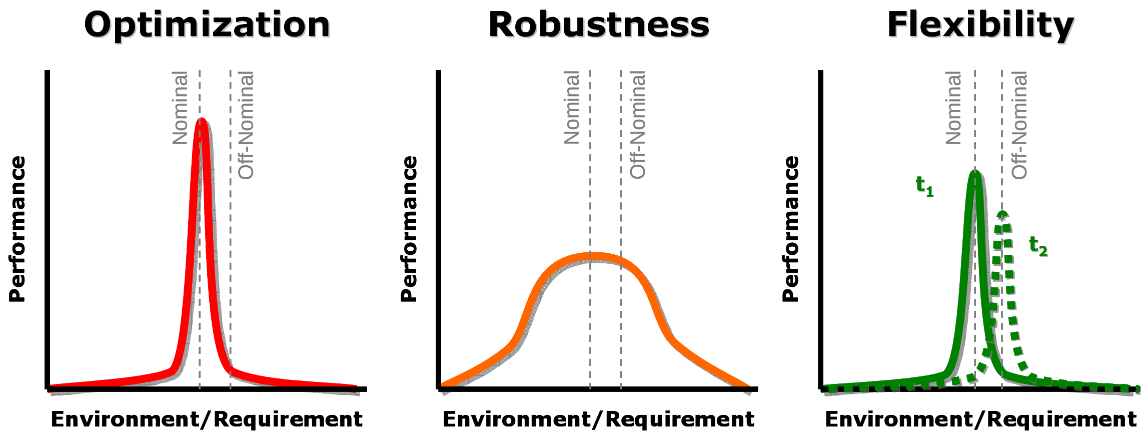

Figure 2. Optimization, Robustness, and Flexibility Notionally Compared.

1.2.2. Observed and Observable Flexibility

As defined here, system flexibility is essentially unobservable until required to manifest itself in response to requirement or environment changes. As a result, flexibility tends to remain an intangible and abstract concept to engineers. To help address this limitation, specific examples of observed flexibility in the history of space exploration and human spaceflight have been previously documented [26]-[28]. Analysis of these examples has highlighted the classification of space system flexibility into the two categories of intra- and inter-mission flexibility [28]. In cases of intra-mission flexibility, a one-of-a-kind system is fielded and then modified over time to adapt to a changing environment or requirements (examples include the Hubble Space Telescope, International Space Station, and the *Mir* space station). In cases of inter-mission flexibility, multiple vehicles are fielded in series and adapted from one mission to another during the course of a program (examples include the Space Shuttle, Apollo, and Venera programs). For both categories, decisions made during design affect the system's ability to adapt to new mission environments and requirements. Examples in the following pages illustrate how this thesis' framework can be applied to both intra-mission flexibility (see Chapter 5) and inter-mission flexibility (see Chapter 6).

1.3. Recent Examples from Industry and Government

Interest in codifying, quantifying, and integrating flexibility in space system design has grown in recent years. Highlighted here are three examples from recent DARPA and NASA programs, representing what may reasonably be considered state of the art (or state of the practice) in incorporating flexibility into space system design.

1.3.1. Exploration Systems Architecture Study (NASA)

In May 2005, NASA Administrator Michael Griffin commissioned the Exploration Systems Architecture Study (ESAS) [29] to recommend an architecture to support sustained human and robotic lunar exploration. In its trade studies, ESAS used five categories of figures of merit, one of which was Extensibility/Flexibility. Within this category were considerations of lunar mission flexibility, Mars mission flexibility, extensibility to other exploration destinations, commercial extensibility, and national security extensibility. ESAS characterized these flexibility considerations in terms of qualitative high (green), medium (yellow), low (red) ratings based on expert judgement. One example of these qualitative ratings for an evolved expendable launch vehicle (EELV) derived crew launch vehicle (CLV) is shown in Figure 3.

The ESAS methodology largely reflects of the state of the practice in designing for space system flexibility today. The approach has positive qualities in that it considers flexibility during conceptual design process, and it does so with the recognition that flexibility must be traded against other objectives such as cost. As a result, this approach is amenable to application of common multi-attribute decision-making techniques. However, this approach has clear disadvantages in its subjectivity and, more importantly, its use of a Likert-like qualitative scale with no physical units. This inhibits the analysis' repeatability and allows the method's results to be readily disputed. More fundamentally, the method treats flexibility as a scalar metric of the same class as cost or performance; it might reasonably be argued that the decision-maker does not actually care about flexibility itself (in whatever units one chooses for it), but rather cares about the *effects* that designed-in flexibility may have on future cost or performance. These shortcomings will be addressed by the framework proposed by this thesis.

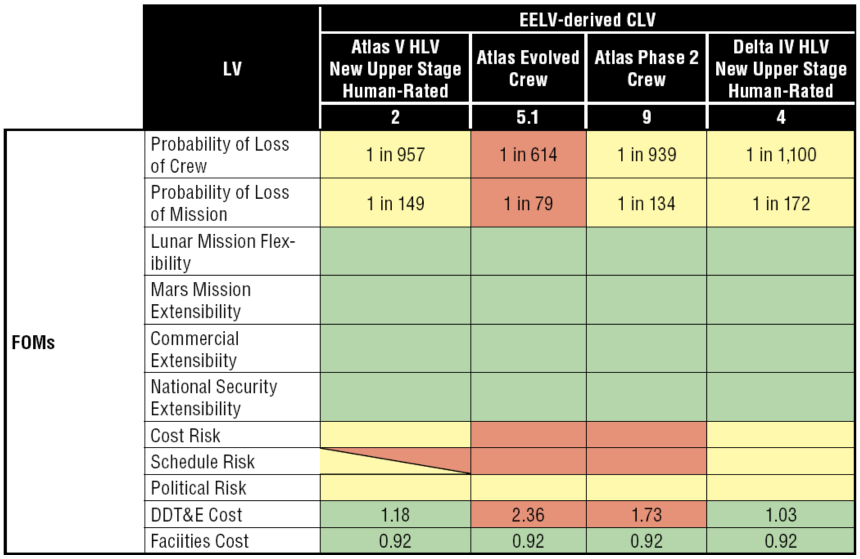

	LV	EELV-derived CLV			
		Atlas V HLV New Upper Stage Human-Rated	Atlas Evolved Crew	Atlas Phase 2 Crew	Delta IV HLV New Upper Stage Human-Rated
		2	5.1	9	4
FOMs	Probability of Loss of Crew	1 in 957	1 in 614	1 in 939	1 in 1,100
	Probability of Loss of Mission	1 in 149	1 in 79	1 in 134	1 in 172
	Lunar Mission Flexibility	🟩	🟩	🟩	🟩
	Mars Mission Extensibility	🟩	🟩	🟩	🟩
	Commercial Extensibility	🟩	🟩	🟩	🟩
	National Security Extensibility	🟩	🟩	🟩	🟩
	Cost Risk	🟨	🟥	🟥	🟨
	Schedule Risk	🟨🟥	🟥	🟥	🟨
	Political Risk	🟨	🟨	🟨	🟨
	DDT&E Cost	1.18	2.36	1.73	1.03
	Faciities Cost	0.92	0.92	0.92	0.92

Figure 3. Sample summary of figure of merit ratings for concepts in the ESAS report. *Note the qualitative red/yellow/green ratings for flexibility.* **[29]**

1.3.2. System F6 (DARPA)

In July 2007, DARPA issued a Broad Agency Announcement for the development of System F6, a flight demonstration of a satellite architecture in which the functionality of a traditional monolithic satellite is fulfilled with a fractionated cluster of free-flying, wirelessly interconnected modules. A purpose of this program was to demonstrate the potential benefits of a system with a built-in capability to respond to mid-mission requirement changes. Four industry teams participated in Phase 1 of the F6 project, and an emphasized component of the project was the development of value-centric design methodologies to account for the full range of benefits (beyond cost savings) available through the fractionated spacecraft approach. All four teams developed discrete event simulations to track cost, revenue, and performance metrics throughout simulated spacecraft lifecycles [30]-[33]. Some teams tracked net present value of the satellite investment (in cases where monetary revenue was an appropriate measure of satellite performance), while others combined performance benefits into an aggregate utility [33]. In order to simulate system operator behavior, the methods tended to assume ad-hoc decision policies (e.g., rules regarding when to replace or upgrade satellites) while exploring the space of possible satellite designs. The efforts of the F6 industry teams represent a considerable step forward in the space industry's ability to quantitatively consider benefits of intangible properties like flexibility.

1.3.3. Flexible-Path Human Space Exploration (NASA)

In 2009, the White House Office of Science and Technology Policy called for the formation of the 10-member Review of U.S. Human Spaceflight Plans Committee (better known as the Augustine Committee) to independently assess the current status and future direction of NASA's human spaceflight program. The committee's final report was released in October 2009 [34]. One of the report's major findings was that "no [human spaceflight] plan compatible with the FY 2010 budget profile permits human exploration to continue in any meaningful way" and that "it is possible to conduct a viable exploration program with a budget rising to about $3 billion annually in real purchasing power above the FY 2010 budget profile." [34]

One of the viable exploration programs the committee proposed was an innovative "flexible path" for human space exploration involving the development of systems to enable mission options for a variety of inner solar system destinations. Highlighted in the committee's report is an example of how missions to the lunar vicinity, Earth-Moon and Sun-Earth Lagrange points, near-Earth objects, Mars vicinity, and the moons of Mars could be accomplished in consecutive years using similar architectural components (e.g., see Figure 4). Figure 5 shows the committee's mapping of possible paths from one destination to another. Although only the green path in Figure 5 was costed and evaluated during the committee's study, a variety of other paths exist. Furthermore, changing political and economic conditions may make demand for particular paths higher than others at different periods in the future. These observations are incorporated later in this thesis when human space exploration architecture selection is used as a demonstration of the proposed decision support framework. In particular, the aim of this example application is to recommend what human spaceflight architecture should be developed initially in order to allow low cost and high return in an environment of uncertain and changing mission demand.

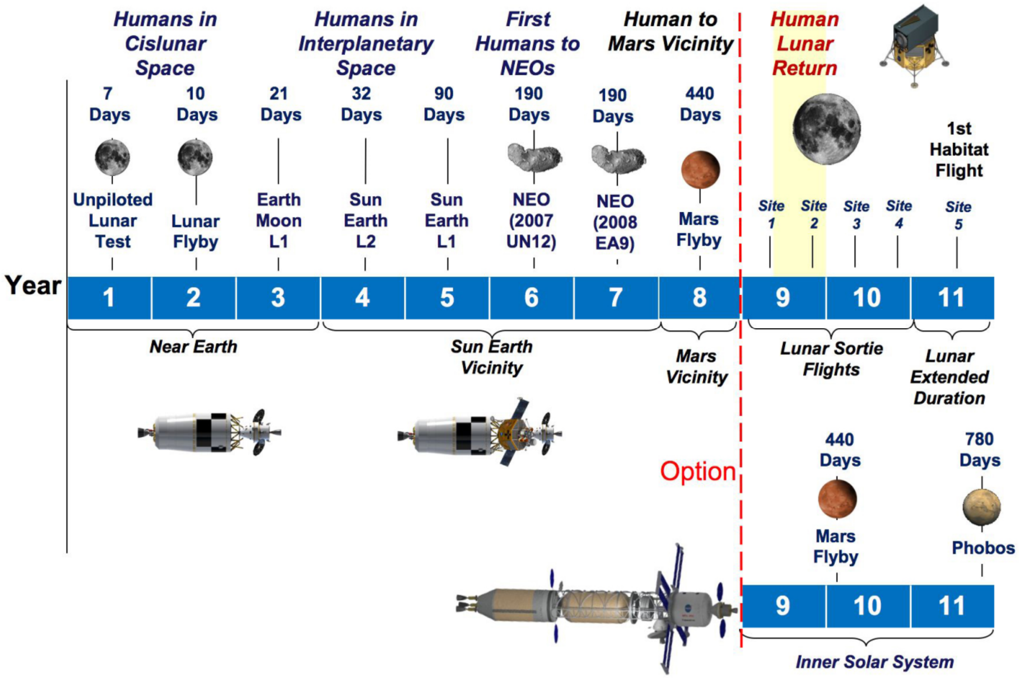

Figure 4. Possible "Flexible Path" Mission Sequence. [34]

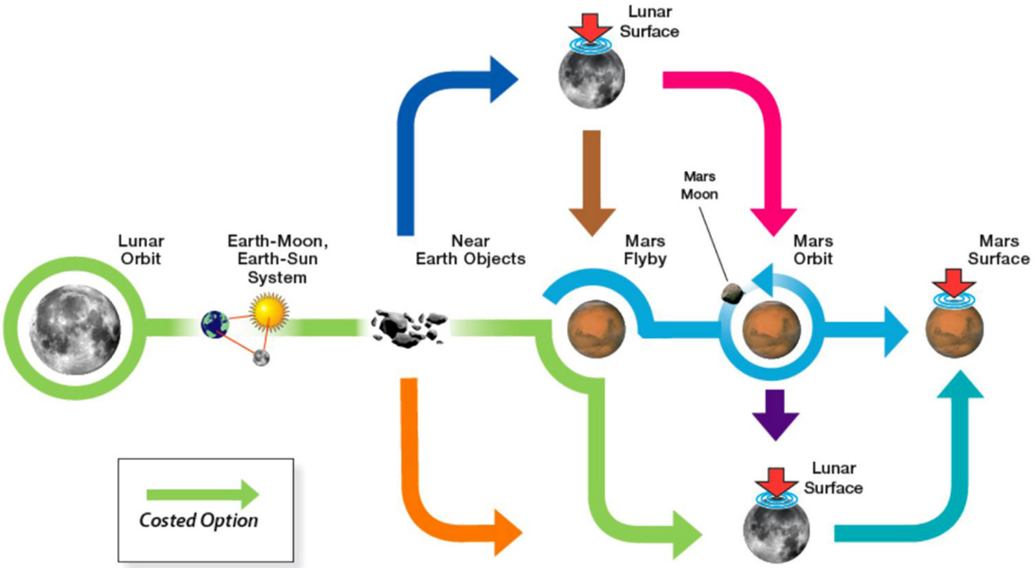

Figure 5. Possible Flexible Path Destination Sequences. [34]

1.4. Flexibility in Space System Design Decisions

In short, flexibility is a property that is sought after by many space system decision-makers but also one that is intangible and difficult to define opertionally. Conceptually, flexibility is the capability to easily modify a system after it has been fielded in response to a changing environment or changing requirements. In the space industry, where environment and requirement changes are prevalent and typically cannot be predicted with certainty, flexibility has been increasingly recognized as important to success. This recognition has been exemplified recently by the fact that DARPA and NASA have proposed flexible spacecraft and flexible paths, respectively, as future program directions with substantial budgetary and resource implications.

However, certain aspects of flexibility, such as its distinction from robustness in the requirement to consider a system's ability to be modified over multiple time periods, present challenges to analysis and decision-making. These challenges add to an already demanding task for space system analysts and decision-makers, involving the enumeration and modeling of many engineering options, understanding the technical, programmatic, and political implications of these options, and making system decisions that strike the proper balance among multiple objectives of differing priorities. To assist in informing the substantially more complex decision facing the decision-maker considering flexibility, flexibility-related challenges are confronted and addressed comprehensively in this thesis, with the objective of enabling selection of the next-generation space systems today that will be best able to adapt and perform in a future of changing environments and requirements.

CHAPTER 2

REVIEW OF LITERATURE ON FLEXIBILITY

Before describing this thesis' proposed framework, it is helpful to review the history of thought on flexibility. This chapter is divided into five sections: The first reviews early notions of flexibility in the economics literature, leading to a brief discussion of decision tree analysis. The second section reviews a common representation of flexibility in terms of next-period decisions within a state space. The third section reviews relevant literature on flexible manufacturing systems, and the fourth reviews recent efforts to consider flexibility in aerospace engineering academia. The latter two sections in particular contain examples that reflect the limited current practice of treating flexibility as a separate scalar metric in a larger decision-making process. The final section identifies this and other gaps in the present literature and state of the practice, establishing the motivation for the framework described in Chapter 3.

2.1. Early Economic Notions of Flexibility

Some of the earliest discussions on flexibility in a decision-making context originate in the economics literature. As early as 1921, economist Frank Knight[*] observed that, compared to agricultural production, which requires commitment at the beginning of each growing season, the supply of manufactured goods "is more flexible over short periods of time" since these goods can be stored and the decision about whether to bring them to the market can be delayed. [35] Sixteen years later, Hart recognized that the postponement of decisions until additional information becomes

---

[*] Among his accomplishments, Knight is known for his distinction between risk and uncertainty. Knight characterized risk as a situation with an uncertain result but certain probability density or mass functions, whereas Knightian uncertainty involves situations with both uncertain results and uncertain distributions.

available is a normal occurrence and preserves flexibility in a business plan. [36] However, he also recognized that this flexibility generally comes at a cost:

> … an entrepreneur who was obliged to make all his decisions as to volume of operations in the present would be unable to use fuller information as it came in, and would have to act on what was available. But the normal case is that the business man expects to be in receipt of additional information bearing on markets at most future dates long before he will have been forced to make all the decisions affecting output … The entrepreneur's fundamental means of meeting uncertainty is the postponement of decisions till more information comes in – that is to say, the preservation of *flexibility* in his business plan. But flexibility involves costs … ordinarily a given production-schedule can be produced at lower cost if the entrepreneur has adapted his input to it well in advance than if plans are improvised. [36]

In 1939, Stigler[*] developed economic thought on flexibility somewhat further. He too recognized that "flexibility will not be a 'free good'" [37] but also illustrated how, in terms of marginal cost and average cost plots (see Figure 6), a flexible plant might have a smaller variability in average and marginal costs as a function of output compared to an inflexible plant. Figure 6, from Stigler's 1939 article, illustrates how at a nominal output *F*, a flexible plant (represented by the dashed line) would incur a higher average cost to produce each item than would an inflexible plant; however, at an off-nominal output *A*, the flexible plant would have lower average costs.[†]

---

[*] In 1982, Stigler would earn the Prize in Economic Sciences in Memory of Alfred Nobel.

[†] In fact, in a competitive market, the inflexible plant would need to close since output *A* falls on a decreasing part of the marginal cost curve, meaning price per item (equivalent to marginal cost) would be less than average variable costs. [37]

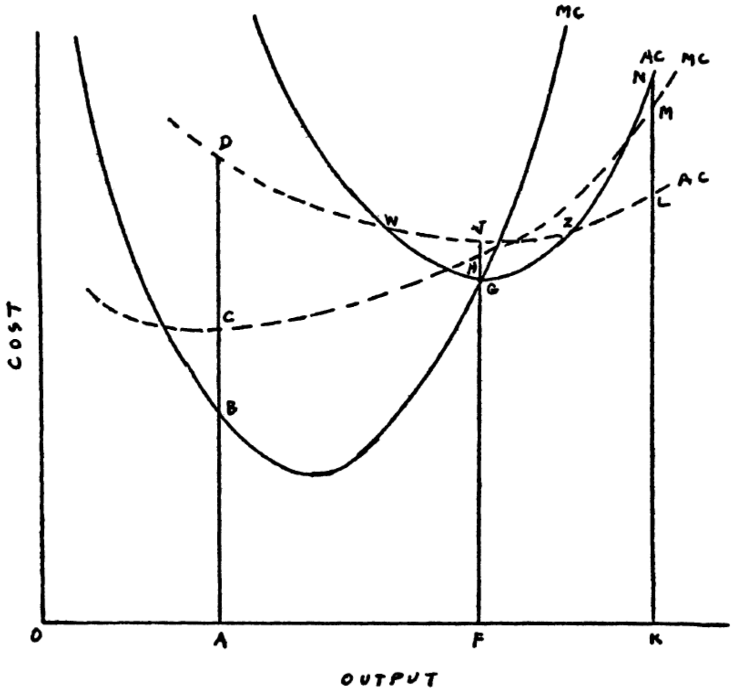

Figure 6. Marginal Cost (MC) and Average Cost (AC) Curves for an Inflexible Plant (solid line) and Flexible Plant (dashed line). [37]

Twenty-five years later, in 1964, Koopmans reiterated the relevance of flexibility by observing that "almost all choices occurring in real life are sequential, 'piece-meal,' choices between alternative ways of narrowing down the presently existing opportunity rather than 'once-and-for-all' choices between specific programs visualized in full detail." [38] Koopmans introduced the notion of "partitioning of opportunities" which, as shown in Figure 7, modeled the narrowing of opportunities with time as a tree of opportunity nodes spaced at discrete times in the future. Koopmans' partitioning of opportunities resembles decision tree analysis, introduced in the late 1950s and 1960s within the broader field of decision analysis. [39]-[45] Decision tree analysis has been used substantially in management, economics, and engineering contexts (for examples, see [44]-[48]), typically for the cases in which a user's objective is minimization or maximization of the expected value of a single profit, cost, or utility metric.[*] A common drawback is that the analysis (and even simply populating the tree's probability inputs) can quickly become unwieldy as the number of options and time periods grow into a "decision bush" rather than a "decision tree" [47]-[48]. Also, typically the focus of decision tree analysis is on valuating existing options rather than recommending which options should be embedded into the system initially. [49] Nevertheless, recognition that the options provided by flexibility can be visualized in a rapidly-expanding tree structure provides a useful model for discussion and thought. It also hints that dynamic programming techniques, which are well-suited to optimizing paths within networks of nodes, may be particularly useful in analysis of flexibility.[†] This idea is incorporated into the approach proposed in Chapter 4; however, first it is important to introduce a second important concept from the economics literature.

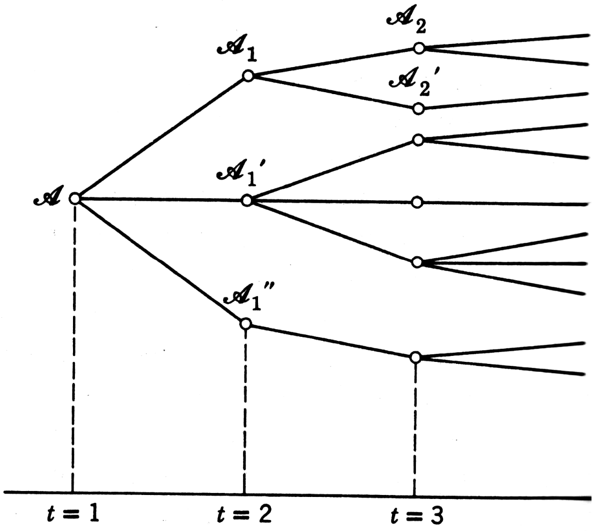

Figure 7. Visualization of Koopmans' Partitioning of Opportunities. [38]

---

[*] An irony of this approach is that, in assuming a single expected-value objective, traditional decision-tree analysis leaves little or no prerogative for the decision-maker to trade different objectives or risks against each other.

[†] For example, one application of stochastic dynamic programming in the later economics literature involves consideration of an individual's labor supply flexibility in order to maximize total discounted lifetime expected utility. [50]

2.2. The Two-Period State-Centric Notion of Flexibility

A second and largely separate body of literature in economics and industrial engineering considers flexibility within a framework of period-to-period transitions between options in a state-space. Epitomizing this view is a paper written in 1984 by Jones and Ostroy [51] which suggested, "Flexibility is a property of initial positions. It refers to the cost, or possibility, of moving to various second period positions." Jones and Ostroy also suggested, "One position is more flexible than another if it leaves available a larger set of future positions at any given level of cost." This was mathematically formalized with Eqs. (1) and (2). Eq. (1) defines $G(a,s,\alpha)$ as the set of next-period positions b attainable from position a at a cost c that does not exceed some value α, in the context of some state s of the operating environment. Eq. (2) formalizes that position a is more flexible than a' (denoted by $a >_F a'$) if the set of positions attainable from a always contains the set attainable from a', excluding the zero-cost option to stay in a'.

$$G(a,s,\alpha) \equiv \{b : c(a,b,s) \leq \alpha\} \tag{1}$$

$$\begin{aligned} &a >_F a' \quad \text{when} \\ &G(a,s,\alpha) \supset G(a',s,\alpha) \setminus g(a') \end{aligned} \tag{2}$$

Thus, an important recognition in Jones and Ostroy's work is that the relative flexibility of two positions is budget-dependent (or resource-dependent). For an infinite budget, two positions would be equally flexible because each can reach the same set of [all possible] future positions. At lower budgets, this may not be true.

However, Eq. (2) has a limitation: It defines relative flexibility only for the case where the set of second-period positions from a' is fully contained within the set of second-period positions from a. As illustrated in Figure 8, no conclusion can be drawn if one of the sets is not fully contained within the other. This is appropriate in principle, as the positions available from *a'* that are not available from *a* may be very important (e.g., may perform particularly well in meeting a particular new requirement or environment), and it illustrates the need to consider more than cost when making decisions regarding flexibility.

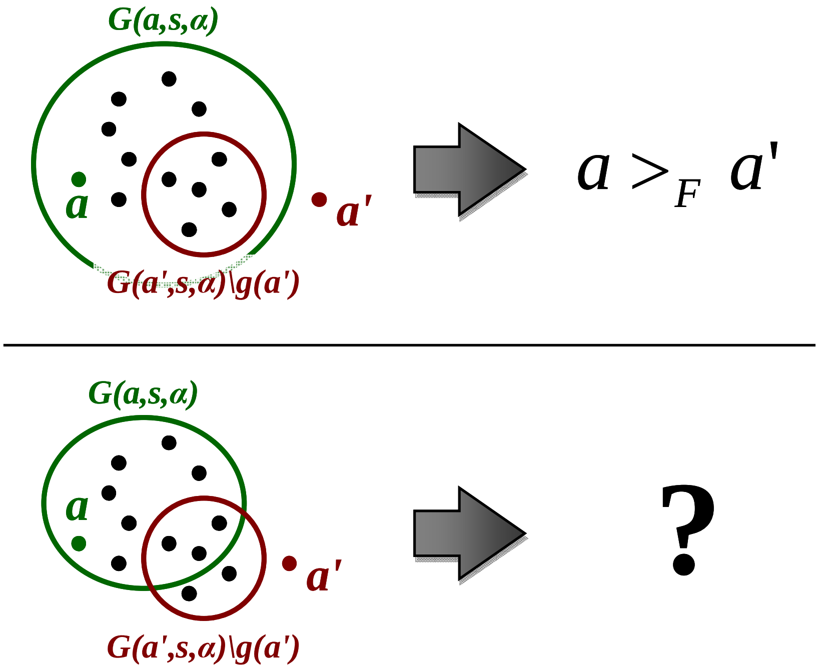

Figure 8. Graphical interpretation of Jones and Ostroy's interpretation [51] of relative flexibility of positions.

Other works which have discussed similar state-centric frameworks include Christian and Olds [52]-[53], Gupta and Rosenhead [54], Baykasoğlu [55], Silver and de Weck [56]-[57], and Mandelbaum and Buzacott [58]. Saleh's visualization of Mandelbaum's and Buzacott's basic concept [49] illustrates the interesting difference from Christian and Olds' and that of Jones and Ostroy in that the system's allowable states are not necessarily the same between periods (i.e., that the definition of a system's state space may change with time, which can complicate analysis). In general, these frameworks and others of this class are helpful because they provide a visualization of the concept of flexibility itself (as opposed to the value of flexibility), which is intuitively related to the number of options that exist for a system as time progresses. However, unlike decision tree analysis (see Section 2.1), these frameworks become difficult to visualize and apply for decisions consisting of more than two periods. In essence, the framework detailed later in this thesis combines the intuitive concept of flexibility provided by these state-centric frameworks with extended variants of the multi-period analysis available through decision trees detailed earlier.

2.3. Flexible Manufacturing Systems

A large body of literature exists within the manufacturing community on the selection and operation of flexible manufacturing systems. A flexible manufacturing system, or FMS, can be defined as a computer-controlled production system capable of processing a variety of part types. [59]-[60] These systems generally consist of computer numerical controlled (CNC) machines, loading and unloading stations, transportation systems for parts and tools, and computerized planning and control systems (e.g., see Figure 9). [59],[61] Key concerns in this field revolve around (1) how to select the appropriate pieces of equipment and layout for an FMS and (2) how to optimally operate an already existing FMS, in both cases to allow the system to optimally (e.g., quickly and inexpensively) respond to changing production requirements.

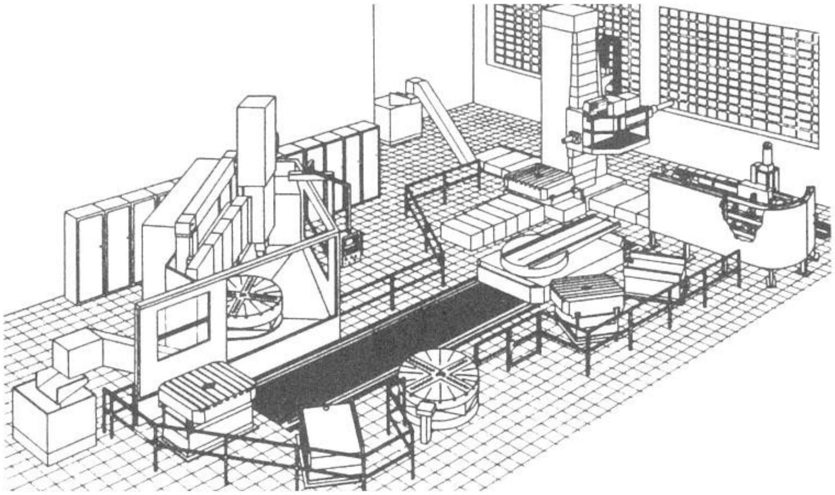

Figure 9. Example setup of a small flexible manufacturing system. [59]

Unfortunately, the literature in this field tends to be highly specialized to the modeling of machining systems, and analysis approaches suggest modeling and optimization techniques specific to different types of equipment rather than one overarching methodology. As a result, FMS measures of flexibility are numerous (including machine flexibility, operation flexibility, routing flexibility, process flexibility, product flexibility, volume flexibility, and more [59],[61]-[63]; one survey identifies 28 different types [64]). In addition, these flexibility metrics are often measured either on a Likert-like (e.g., 1-5) scale (e.g., see [64]) or on a scale whose physical meaning is difficult to interpret [55],[65]. This combination of disparate, qualitative metrics can complicate decision-making, an issue which has been recognized in the past. For example, Gupta and Goyal [63] note, "a single all encompassing measure of MF [manufacturing flexibility] seems to be an evasive issue and such a measure is yet to be developed," and Mohamed [62] and Cox [66] note, "the concept of flexibility is new, with no acceptable measurement, and consequently is treated on an abstract basis rather than a concrete basis." [62]

One interesting detail raised by the work of Tempelmeier [61] and Tetzlaff [59] is that dynamic programming techniques may be used to find the lowest-cost route within a network in which nodes are time periods and arcs are paths of fixed FMS configurations (see Figure 10). However, an important limiting assumption behind this approach (and throughout the thesis of Tetzlaff [59]) is that the required production rate in each period is known in advance. This highlights that the definition of flexibility posed in Section 1.2 encompasses situations in which future changes in environments and requirements are precisely known in advance (i.e., the deterministic limit of the more general case where environments and requirements are not known in advance). Thus, Tempelmeier and Tetzlaff illustrate the curious concept of flexibility with respect to deterministic requirement changes.

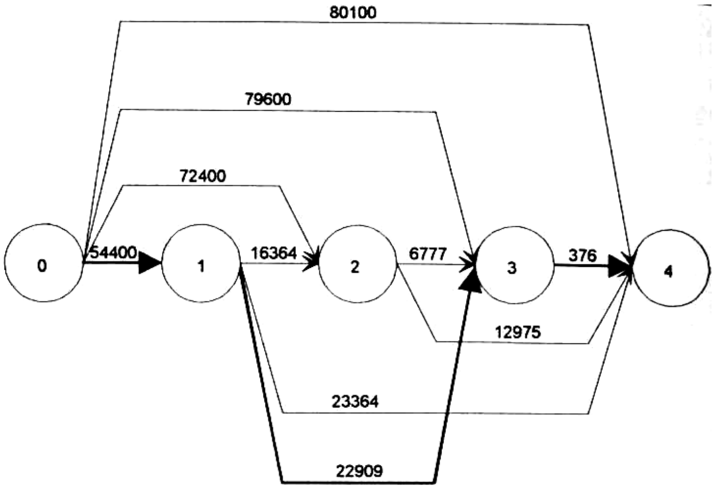

Figure 10. The deterministic flexibility example of Tempelmeier and Tetzlaff [61]. *The optimal path through the network is shown by the dark path 0 → 1 → 3 → 4.*

2.4. Examples from Aerospace Engineering Academia

While the aerospace industry's consideration of flexibility in design of new space vehicles is largely reflected by the examples provided in Section 1.3, the aerospace engineering academic community has recently proposed a variety of additional techniques.

Ross, Viscito, and Rhodes [67]-[68] propose the analysis of flexibility in terms of epochs and eras, where an epoch is a time period of "fixed context and fixed value expectations" [68] and an era is a time-ordered sequence of epochs (i.e., one possible timeline of expectations). Once an era is defined, Ross and Viscito [67] propose quantifying flexibility with a metric called value-weighted filtered outdegree (VWFO) defined in Eq. (3). In this equation, u_j^{k+1} indicates the utility of system design option j in epoch $k+1$ (i.e., the next epoch), and $Arc_{i,j}^{k}$ is a binary 0 or 1 depending on whether the transition is possible for a given budget (or "filter"). As a result, systems with many high-utility next-epoch (next-period) options and few low-utility next-epoch options receive high VWFO scores. However, this metric has limitations. First, the use of the signum function in the summation of Eq. (3) permits a system design with many high-utility options and many low-utility options to have a VWFO indistinguishable from one with only medium-utility options. Second, VWFO is computed from epoch to epoch, making it difficult to assess for an entire era. Finally, the metric convolves the notion of flexibility with the value (or utility) of that flexibility, preventing the two from being distinguished. However, the metric contributes a clear example employing a two-period state-centric concept of flexibility similar to that of Jones and Ostroy (see Section 2.2), including use of a budget constraint.

$$VWFO_i^k = \frac{1}{N-1}\sum_{j=1}^{N-1}\left[\operatorname{sgn}\left(u_j^{k+1} - u_i^{k+1}\right) * Arc_{i,j}^k\right] \tag{3}$$

More recently, in 2010 Olthoff, Cunio, Hoffman, and Cohanim [69] proposed a seven-step procedure for applying flexibility, in an attempt to develop a practical example of designing flexibility into a small guidance, navigation, and control testbed. The group identified two strategies for flexibility, namely modularity and maximum overhead capacity (i.e., system margin). Limitations exist in that, at present, the group does not appear to distinguish flexibility from robustness, and development of the testbed appears to have occurred in parallel with development of the flexibility decision procedure. As a result, the example application used the decision procedure in retrospect to justify decisions already made, and the procedure to date lacks detail on the tools needed to fully inform decision-making.

Substantially greater depth on the nature and complexity of the flexibility problem was covered in theses by Saleh [48] and later Mark [70] and Nilchiani [71]. In 2002, Saleh [48] extensively motivated the need for flexibility in space systems and examined its definition, in particular contrasting it against the more static property of robustness. Specific examples were provided to illustrate the need for flexibility in modern space systems, including instances of historical requirements change, market demand change, and obsolescence. Saleh applied techniques from decision tree and real options analysis to demonstrate the existence of net-present-value-optimal design lifetimes for revenue-generating satellites and used these techniques further to quantify the value of satellite servicing.

In 2005, Mark [70] further explored designing flexibility into systems for the application of an unmanned aerial vehicle, proposing to consider flexibility in the context of platforms and frames (where the platform is the set of common elements between modified designs, and the frame is the set of changed elements). Mark proposed to define flexibility as "the ratio of performance enhancement (output) to the cost and time required to realize such an enhancement (inputs)". [70] Later in 2005, Nilchiani [71] proposed a 12-step process for assessing the value of flexibility in a space system, which included using decision trees as well as creating a "flexibility tradespace" for visualizing alternatives' cost-revenue (and/or cost-benefit) trades one period into the future. Nilchiani also addressed how the proposed methodology could be integrated into a multi-attribute trade-space exploration in a merged methodology named FlexiMATE. [71]

In 2009, Lim [72]-[73] also proposed a general approach to design evolution, focusing on aircraft and using example applications of evolving the F/A-18 Hornet fighter as well as a simpler cantilever beam design. Lim adopted the framework of stochastic programming with recourse in order to optimize the initial design of a system while probabilistically considering events that could unfold one period in the future. Lim suggested a combination of deterministic scenario-based optimization, stochastic programming, and interactive decision support tools to design evolvable systems using a 9-step process named EvoLVE.

The work of Christian and Olds [52]-[53] is another recent example of aerospace literature considering flexibility. In their work, Christian and Olds describe flexibility in terms of a system's ability to move between different end states in a lawful state space (similar to the two-period state-centric framework described in Section 2.2). An example application evaluates two competing human exploration architectures in terms of their ability to easily achieve extended lunar missions, if future requirements dictate such a need. Three state variables are defined to describe the performance requirements of the extended lunar mission,[*] and a Difficulty Scale for Evolvability Analysis (DSEA) is formulated to permit expert judgement to rate the difficulty (on a 1-3-9-27-81 power scale) of evolving each architecture to meet various second-period performance states. The authors observed that "a single metric cannot capture the sensitivity of an architecture's capability to evolve" since that capability depends on the final evolved state that is desired. For example, architecture *A* may be able to easily adapt to requirement *x* but not requirement *y*, while architecture *B* may be able to easily adapt to requirement *y* but not requirement *x*. In such a scenario, it cannot be said that either *A* or *B* is more flexible (or evolvable, or adaptable) unless the future requirement is known a priori.

In 2006, Silver and de Weck [56]-[57] proposed an analysis of evolvability based on expansion of a network of system operating and switching costs through several time periods. A set of particular deterministic exogenous demand scenarios was assumed, and an optimizer was used to find the least-cost path through the network for each scenario. Silver and de Weck refer to the method as a time-expanded decision network (TDN) and apply it to selection of an example NASA heavy-lift launch vehicle. One notable limitation to the method is its single-objective and deterministic solution approach: Since the exact present and future demands of each scenario are known in advance to the decision-maker (or optimizer), paths through the time domain are able to fully specify any optimal solution. No explicit consideration is given to the possibility that a decision-

---

[*] Contrary to Jones and Ostroy, whose state-space "positions" appear to refer to future options, the state space of Christian and Olds is defined by the *performance* of those future options. For reasons that should become apparent in Section 4, this thesis primarily supports the view of Jones and Ostroy. However, as will be shown, incorporation of future requirements into the state space is required in order to apply the Markov decision process approach.

maker will make choices in part to hedge against uncertain future events. In this sense, the approach is similar to the deterministic flexibility considered by Tempelmeier and Tetzlaff (see Section 2.3).

A final note should be made on recent work from Daniels, Tracey, Irvine, Schram, and Paté-Cornell [74], presented in March 2011 and developed independently of the present thesis. Motivated by recent DARPA efforts toward developing value-centric frameworks to address the business case for fractionated spacecraft, the authors propose heuristic and dynamic-programming-optimized decision rules for the operation of future fractionated spacecraft. Using the example of a fractionated 3-module weather satellite, the work simulated the state of the satellite and support systems (e.g., which modules were functional, whether spares existed on the ground) and optimized the procurement or launch of new modules in order to achieve the highest expected net present value under a set of assumptions to assign a dollar value to incoming weather data. While the authors' goals differ substantially from those of the present thesis (e.g., they do not seek to operationally define or measure flexibility, nor are they interested in informing initial system design decisions), their use of Markov decision processes from the operations research community is common.

This set of literature from aerospace academia is listed in order in Table 1 and summarized in terms of several important characteristics that have arisen in the preceding discussion. Each element of the table indicates whether each work either implemented/provided (⦿), recognized (△), or did not address (no mark) each of the six characteristics represented by the columns:

Beginning with the first column, it is noted that many of the aerospace works surveyed here have arisen as a result of efforts to provide further definition to the concept flexibility, and some of these works have explored this topic in great depth. As Table 1 indicates, six of these nine works provide definitions for flexibility (or, in one case, the equivalent term evolvability). Fewer of these works recommend metrics for flexibility, although in some cases the metrics of others are recognized in literature reviews. Also in terms of this second category, it is worth noting that neither of the two works that provide objective metrics for flexibility clearly distinguish flexibility from its value.

The third column indicates whether a work considered trades among multiple distinct objectives. In general, this topic tends to be covered unsystematically or not at all among present works on flexibility, perhaps in part because much of the aerospace flexibility literature has focused on application to systems with priced services (cf. Daniels, Tracey, Irvine, Schram, and Paté-Cornell [74], Nilchiani [71], and Saleh [48]). In cases where multiple decision-maker objectives are considered, few, if any, mentions are made of efforts to seek Pareto-optimal trades among these objectives to ensure that the decision-maker is making an objectively good decision.

The fourth column indicates whether a work utilizes or considers stochastic models, and the fifth indicates whether a work considers decisions at multiple future periods. With the exception of the very recent work of Daniels, Tracey, Irvine, Schram, and Paté-Cornell [74], note that implementation of these two characteristics is mutually exclusive. This major limitation reflects the fact that posing a stochastic single-future-decision problem and a deterministic multiple-future-decision problem are each relatively simpler than posing a stochastic multiple-future-decision problem. However, overcoming the complexity of solving this more realistic problem brings with it corresponding benefits.

The final column indicates whether a work implements or proposes a framework by which an engineer or decision-maker is intended to make an initial system decision. In many cases, such a framework is the intent of the work, but in some cases (e.g., Daniels, Tracey, Irvine, Schram, and Paté-Cornell [74], Saleh [48], and Ross, Viscito, and Rhodes [67]-[68]) it is not.

The final row in Table 1 indicates that it is the intent of the present thesis to contribute toward each of the key characteristics that have been here identified. More importantly, however, it is the intent of this present work to do so in an integrated fashion. In considering all aspects of this problem simultaneously, it is intended that this thesis will contribute not only improvements in previous works' considerations of individual aspects of the flexibility topic, but that it will contribute a more coherent understanding of flexibility as a whole.

Table 1. Summary of Flexibility-Related Literature in Aerospace Academia.

⊙ (green) = Implemented or Provided
△ (orange) = Recognized
No Mark (red) = Not Addressed

Authors	Year	Characteristic Definition of Flexibility (or Evolvability)	Objective Metric(s) for Flexibility (or Evolvability)	Trades among Multiple Distinct Objectives	Stochastic	Considers Decisions at Multiple Future Periods	Framework for Initial System Design or Selection
Ross, Viscito, and Rhodes	2009	⊙	⊙	△		△	
Olthoff, Cunio, Hoffman, and Cohanim	2010	⊙		△			△
Saleh	2002	⊙	△		⊙	△	
Mark	2005	⊙	⊙	△		△	⊙
Nilchiani	2005	⊙	△	△	⊙	△	⊙
Lim	2009			⊙	⊙		⊙
Christian and Olds	2004	⊙		⊙			△
Silver and de Weck	2006				△	⊙	⊙
Daniels, Tracey, Irvine, Schram, and Paté-Cornell	2011			△	⊙	⊙	
Present Thesis	2011	⊙	⊙	⊙	⊙	⊙	⊙

2.5. Gaps in the Flexibility Literature

In summary, this chapter has surveyed a broad set of literature spanning economics, industrial and systems engineering, and aerospace engineering. In combination with the state of the practice in industry described in Chapter 1, certain gaps are evident in current thinking on flexibility and current implementation of methods to consider this property in system design:

- In much of the literature (esp. cf. Sections 2.3 and 1.3.1), there appears a tendency for engineers to consider flexibility as a system-dependent scalar quantity. This concept has driven the invention of numerous scalar measures for flexibility that are often subjective and expressed on a scale with no units or clear physical interpretation. Further, when or if these measures are used in trade studies, they imply that flexibility is a property of the system separate from all others (such as cost and performance measures). ***However, the decision-maker likely has little interest in flexibility for the sake of flexibility: He or she cares about flexibility primarily because of cost and performance benefits it may enable in the future.***

- ***Few existing methods for considering flexibility look at decisions more than one period in the future.*** While considering one future period is an important first step, it is only one period less myopic than the traditional single-period horizon. If a system or program is to be operated for many decades (as is often the case in the aerospace industry), the prudent decision-maker cares not only to consider options for the first time that requirements or environments change, but also for many subsequent changes.

- ***Furthermore, of methods that do consider implications of flexibility more than one period into the future, few utilize stochastic models.*** Some methods assume a deterministic schedule of future requirements, while others select a handful of deterministic scenarios upon which to evaluate the system of interest. However, the probability of any one scenario occurring may be nearly (or, if continuous random variables are involved, exactly) zero. Without an understanding of the underlying probabilities of transition between demand or requirement environments, it may be problematic to assume a handful of scenarios can properly represent the entire space of possible futures.

- While some existing methods (such as decision trees) permit valuation of the avenues of flexibility provided by a system, they typically operate by assuming a single expected-value objective function. ***In reality, engineering design involves trades among multiple cost and performance metrics as well as measures of dispersion for these parameters when subject to a stochastically changing environment.***

- Finally, the flexibility literature contains little discussion about the policies that flexible system operators should use to decide whether to exercise the options provided by flexibility. Some appear to assume that the appropriate policy is to always modify the system to precisely meet the anticipated demand or requirement. However, this is a very special case, and it may be in the program's best interests not to meet this demand if it is likely to be transient,* or to over-perform if doing so is likely to boost performance in a later period of high demand. ***The policy by which the system will be operated is an important part of system design, especially for a flexible system.*** It would be imprudent to design a flexible system and "throw it over the fence"

---

* For example, in 1983 Bernanke illustrated a class of stochastic problem in which optimal investment must involve at least one period of *no* investment at all. [75]

to the operators with no guidance on how or when to exercise the options that were so carefully embedded.

In summary, **today there exists no quantitative, stochastic, multi-objective, and multi-period framework for integrating flexibility into space system design decisions**. It is such a framework that this thesis proposes. It is fully recognized that in order to be practicable, this framework must (1) originate from an intuitive and easily communicable operational understanding of flexibility, (2) provide enough structure and tools to guide analysis but not so much as to lead to "process tunnel vision" for the engineer in the field, (3) require a reasonable number of inputs, and (4) provide for clear interpretation of results. To accomplish this, the framework draws from literature and tools from operations research, engineering, and economics in order to operationally define flexibility and transform its consideration into a tractable problem of stochastic optimal control.

CHAPTER 3

OBJECTIVES AND CONTRIBUTIONS

The contributions of this thesis largely address the gaps in the present literature identified in Section 2.5. These gaps suggest that at least four components are critical for a decision framework that integrates flexibility into space system design decision-making: First, a **stochastic** model for the evolution of system demand over **multiple future time periods** must be developed; such a model must describe what a system may be expected to accomplish (or what a decision-maker may be rewarded for performing) in the future. Second, a set of **candidate system designs** or configurations must be developed that is valid for multiple time periods in the future; this describes the future options available to the decision-maker and is suggested by the two-period state-centric notion of flexibility in the literature. **Quantitative performance measures** are required to evaluate how well the configuration that is fielded at a given time fulfills the demand or mission requested of it; in some scenarios, multiple performance measures may be required to capture trades among **multiple objectives**. Finally, since decisions regarding which system(s) to develop and field next must be made at multiple future time periods, a process must exist for providing **sequential decision support** in an easily interpretable manner. Since the framework developed in this proposal is intended to be used by decision-makers facing an immediate system selection problem, of particular interest is to aid in **informing initial system selection**. These components are illustrated graphically in Figure 11.

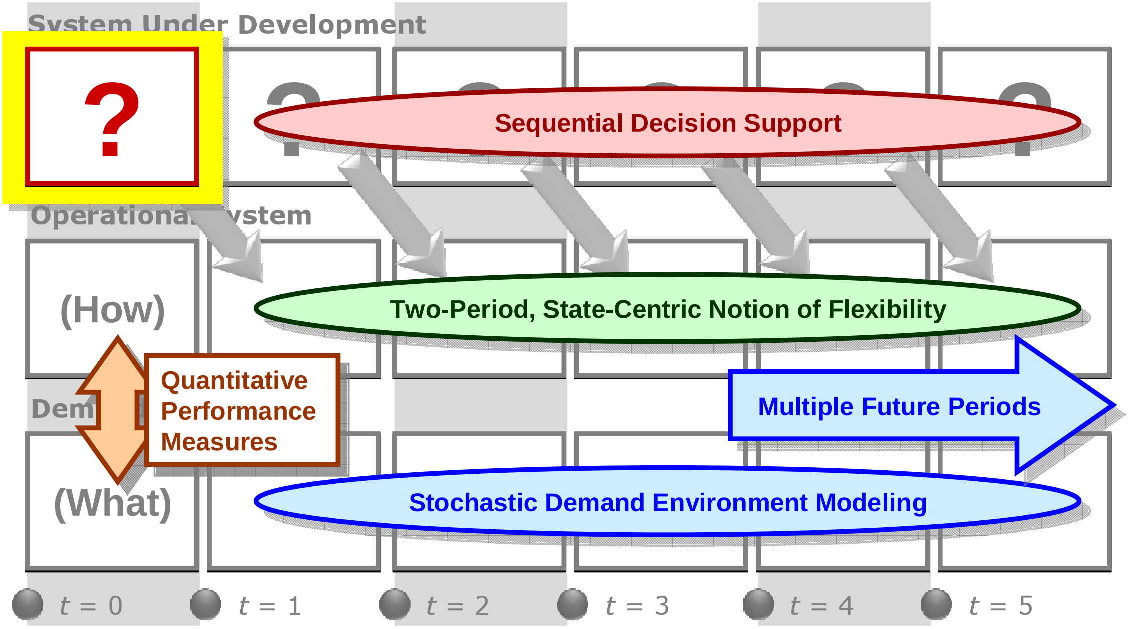
System Under Development
?
Sequential Decision Support
Operational System
(How)
Two-Period, State-Centric Notion of Flexibility
Dem
Quantitative Performance Measures
Multiple Future Periods
(What)
Stochastic Demand Environment Modeling
t = 0
t = 1
t = 2
t = 3
t = 4
t = 5

Figure 11. Critical components for decision frameworks addressing gaps in present flexibility literature.

Guided by the present gaps in the literature and the critical components above, this thesis develops a particular set of steps that engineers and decision-makers in the future can follow not only to better understand modes and implications of flexibility for their particular engineering systems, but also to identify best possible initial system or architecture designs. Considering flexibility in a way that addresses these gaps in current methods will enable the selection of systems today, tailored to the decision-maker's budget and preferences, that will be best able to perform when subject to a future of changing environments and requirements. To accomplish this, core objectives and contributions of this thesis include:

- Formulation of the two-period state-centric notion of flexibility (see Section 2.2) as a formal configuration-state-based concept for space system analysis and design.

- Formulation of a state-centric stochastic multi-period model capable of describing evolution of the demand environment in which an engineering system operates.

- Incorporation of system modification policy into initial system selection by using the above formulation to pose integration of flexibility in design as a solvable sequential decision-making problem.

- Implementation and demonstration of the utility of solving for the multi-objective (Pareto-) optimal sequential decisions enabled by flexibility, including:

    - Optimal "open loop" sequential system configuration paths, in which future system configurations are changed according to a preset schedule.

    - Optimal "closed loop" system configuration policies, in which future system configurations are chosen based on a combination of the current configuration and current demand environment. Enabling tools from the operations research community are the formulation and probabilistic dynamic programming solution techniques for Markov decision processes. In addition, Appendix A contributes a new heuristic technique for identifying concave portions of Pareto frontiers in dynamic programming problems.

- Systematic use of multi-objective (Pareto-) optimal configuration paths and policies to recommend initial system configuration decisions.

- Application and illustration using the examples of (1) communications and reconnaissance satellite system selection, (2) multiple- or distributed-payload satellite selection, and (3) NASA human space exploration architecture selection. The latter two examples in particular use this thesis' framework to provide practical insights regarding current problems of interest to the space industry. Also, Appendices B and C include contributions of a human space exploration transition cost model and Markov chain expert judgement elicitation implementation that permit the NASA example to be executed.

The remainder of this thesis is organized as follows: Chapter 4 introduces the five-step framework central to the thesis and extensively establishes its theoretical basis. Chapter 4 also includes a demonstration of the framework for a simple example in which a small government must decide upon whether to develop and field 0, 1, or 2 communications or reconnaissance satellites at various future time periods. Chapter 5 applies the newly developed framework to a more current fractionation-related question of whether to distribute payloads among multiple free-flying modules for an Earth-orbiting satellite. Chapter 6 applies the framework to a current NASA question of what human spaceflight architecture decisions will result in maximum long-run return for minimum long-run cost. The latter example introduces several modeling complexities to the framework, demonstrating significant extensibility beyond the simple examples of Chapters 4 and 5. Chapter 7 concludes with a summary and suggestions for future work.

CHAPTER 4

THEORETICAL BASIS FOR A MARKOVIAN STATE-SPACE FLEXIBILITY FRAMEWORK

Based on the gaps in the current literature observed in Chapter 2, Chapter 3 established that the three rungs of Figure 11, plus the quantitative performance measures linking the bottom two rungs of the figure, are critical for a decision framework that integrates flexibility into space system design decision-making. To accommodate these requirements, this chapter presents a framework consisting of five basic steps, outlined in Figure 12. First, system configuration options are identified and costs of switching from one configuration to another are compiled into a cost transition matrix. Second, probabilities that demand on the system will transition from one mission to another are compiled into a mission demand Markov chain. Third, one performance matrix for each design objective is populated to describe how well the identified system configurations perform in each of the identified mission demand environments. Fourth, possible future sequences of system configurations are simulated and sequences that are Pareto-optimal in terms of the decision-maker's objectives are identified. In a complementary approach, the system decision problem is formulated as a multi-objective variant of a Markov decision process, and Pareto-optimal decision policies are identified. Finally, the paths and policies from the latter step are synthesized into a set of data to inform initial system selection.

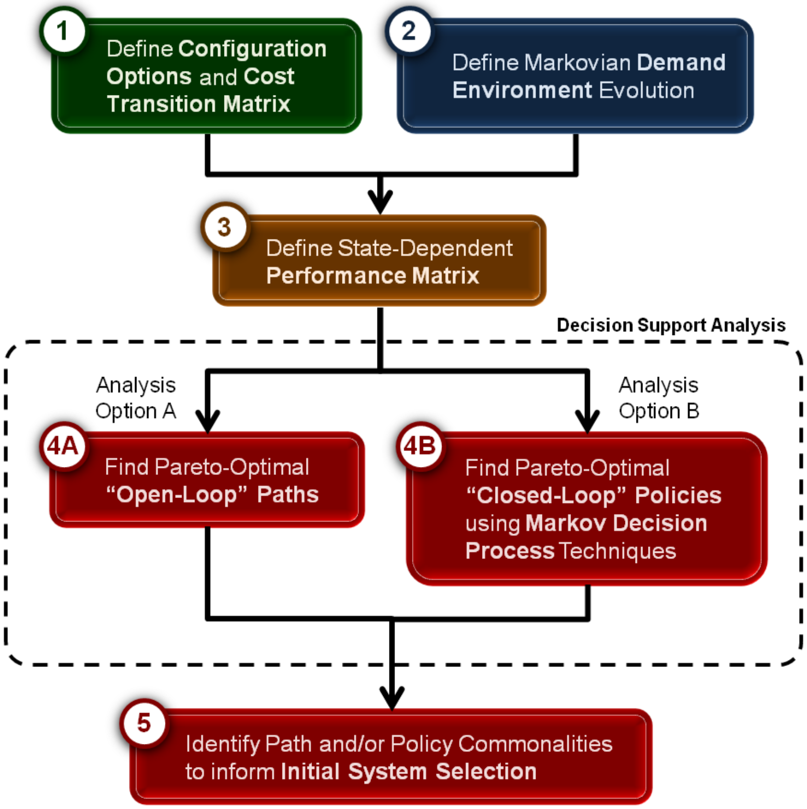
1
Define Configuration Options and Cost Transition Matrix
2
Define Markovian Demand Environment Evolution
3
Define State-Dependent Performance Matrix
Decision Support Analysis
Analysis Option A
Analysis Option B
4A
Find Pareto-Optimal "Open-Loop" Paths
4B
Find Pareto-Optimal "Closed-Loop" Policies using Markov Decision Process Techniques
5
Identify Path and/or Policy Commonalities to inform Initial System Selection

Figure 12. Five major steps of this thesis' framework.

4.1. Step 1: Define Configuration Options and the Cost Transition Matrix

As noted in Section 2.2, in 1984 economists Jones and Ostroy [51] suggested, "Flexibility is a property of initial positions. It refers to the cost, or possibility, of moving to various second period positions." Thus, Step 1 of this proposed framework begins by defining: What are the possible "positions" of an engineering system?

4.1.1. Defining the Configuration Space

This framework proposes that the "positions" of an engineering system are its possible configurations, or its possible design options. This choice for the position definition has the reasonable implication that given enough resources, the engineer or decision-maker can choose to field any particular system configuration (or be at any particular "position") in the future.

What precisely defines such a set of configurations is application-specific but may be guided by the fact that systems, by definition, consist of combinations of lower-level components or characteristics. In the example of an airplane, each system configuration might be defined by a combination of characteristics such as wing sweep angle and aspect ratio, engine type, and fuselage diameter. In the case of a satellite constellation, each configuration might be defined by characteristics like number of satellites, number of orbit planes, angular spacing between satellites in a plane, and the inclinations and right ascensions of the ascending nodes of the orbit planes. In other words, each system S_i in a set of systems $\{S_i, i = 1, 2, \ldots, N\}$ may be defined by a set of design variables $\{x_k, k = 1, 2, \ldots, M\}$, where N is the number of candidate systems under consideration and M is the number of design variables required to uniquely define each system. Written concisely, S_i is defined by the ordered M-tuple $(x_1, x_2, \ldots, x_M)$.

Thus, the available configurations for an engineering system may be considered to comprise a configuration state space that can be visualized as a set of design points in a hyperspace in which each dimension represents a particular design variable or design characteristic. A simple two-dimensional example is illustrated in Figure 13. Here, the configuration state-space consists of five systems of interest defined by particular combinations of values of the design variables x_1 and x_2. An important concept conveyed by this visualization is that different discrete systems may not be equally distinct from each other. For example, in Figure 13 it is clear that S_1, S_2, and S_3 are physically quite alike in the sense that their defining design variables have similar values; in contrast, S_4 and S_5 lie in different areas of the configuration state space and are physically different. Thus, within the configuration state space, distance (euclidean or otherwise) is an indicator of the physical similarity of two systems.

Unfortunately, in many cases, the design variables may not have cardinal or even ordinal properties. For example, one design variable for a satellite might be the type of battery used for energy storage (e.g., nickel-cadmium, nickel-hydrogen, or lithium-ion batteries). In such cases, engineering judgement may still suggest that some of the options are more alike than others and at least a qualitative notion of distance may still exist.[*]

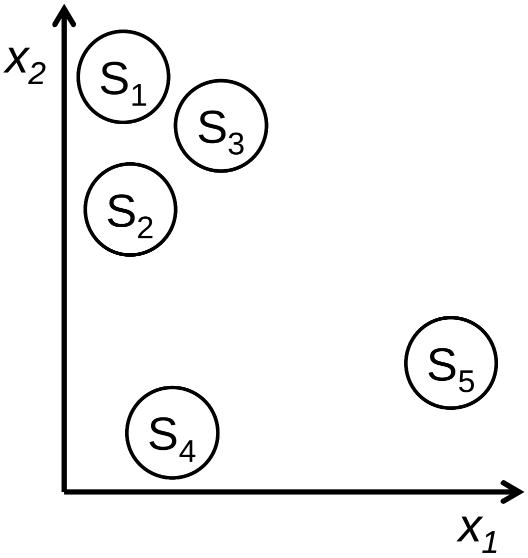

Figure 13. Example configuration state space in which five systems are defined by two design variables.

From where do the discrete systems of the configuration state space originate? In some applications, an engineer may be faced with a problem in which many candidate configurations have already been defined. In others, a systematic process may be required to identify these configurations. This is a common early step in multi-attribute decision-making (MADM) methodologies. Since the combinatorial space of alternative system configurations can be quite large, previous works have proposed the use of morphological matrices as a brainstorming tool (and occasionally as a tool to enumerate the entire combinatorial space) [17],[76]-[79].

An example morphological matrix is shown in Table 2. Each row denotes a particular design variable x_k for the system, and possible discrete values for each variable

---

[*] In many cases, this qualitative notion of distance might be rigorously quantified by defining each of the nominal (non-ordinal and non-cardinal) options in terms of their own internal design variables.

are listed as options. A single configuration S_i is defined once one value is selected from each row. Thus, if values for each design variable may be selected independently, the total number of states N_{total} in the configuration state space is the product of the cardinality of each design variable's set of possible discrete values. This relationship is expressed via Eq. (4).

This total number of states can be quite large, depending on the number of design variables considered and the number of values each may take. One way to restrict the architectures considered to a representative but manageable set is to use the morphological matrix to assist in brainstorming themed configuration options [77]. In this case, the number of configurations N considered in the analysis will be less than full-factorial N_{total}.

Table 2. Example Form of a Morphological Matrix.

Design Variable	Discrete Values						Number of Discrete Values
x_1	$x_{1,1}$	$x_{1,2}$	$x_{1,3}$	$x_{1,4}$			$\lvert x_1 \rvert = 4$
x_2	$x_{2,1}$	$x_{2,2}$	$x_{2,3}$				$\lvert x_2 \rvert = 3$
x_3	$x_{3,1}$	$x_{3,2}$	$x_{3,3}$	$x_{3,4}$	$x_{3,5}$	$x_{3,6}$	$\lvert x_3 \rvert = 6$
.							.
.							.
.							.
x_M	.	.	.	.	.	.	$\lvert x_M \rvert$

$$N_{total} = \prod_{k=1}^{M} \lvert x_k \rvert \tag{4}$$

4.1.2. Defining the Cost Transition Matrix

Recalling that flexibility "refers to the cost, or possibility, of moving to various second period positions" [51], to proceed it is necessary to incorporate cost information in addition to information on the composition of each system configuration. For engineering systems, this cost typically consists of two temporally distinct components: recurring and nonrecurring costs.

4.1.2.1. The Development Cost Transition Matrix

The cost information most central to the concept of flexibility falls in the category of nonrecurring costs. These costs, which typically account for the one-time costs required to develop a new engineering system, are related to the resources required to develop the system given existing components and technologies. In other words, these are the transition costs incurred due to a switch from one configuration to another. As a result, these switching costs are naturally defined in a pairwise manner. This thesis proposes the definition of a matrix C_{dev} where the elements $c_{dev,ij}$ are the costs incurred to develop configuration j for the next time period given that the configuration in the current period is i. Table 3 illustrates the format of such a matrix.

Table 3. Example Format for a Cost Transition Matrix *C*.

		To Configuration						
		S_1	S_2	S_3	·	·	·	S_N
From Configuration	S_1	$c_{1,1}$	$c_{1,2}$	$c_{1,3}$	·	·	·	$c_{1,N}$
	S_2	$c_{2,1}$	$c_{2,2}$	$c_{2,3}$	·	·	·	$c_{2,N}$
	S_3	$c_{3,1}$	$c_{3,2}$	$c_{3,3}$	·	·	·	$c_{3,N}$
	·	·	·	·	·			·
	·	·	·	·		·		·
	·	·	·	·			·	·
	S_N	$c_{N,1}$	$c_{N,2}$	$c_{N,3}$	·	·	·	$c_{N,N}$

Costs in this matrix may be calculated element-by-element using available parametric models or other cost estimation techniques. However, in some cases, especially when large numbers N of possible configurations are under consideration and N^2 elements must be populated, this technique may become time-prohibitive. In such cases, simplifying approximations for these transition costs may be warranted. For example, consider a configuration state space in which each system S_i is defined by a set of binary design variables. That is, $S_i = [x_1\ x_2\ \ldots\ x_M]^T$, where $x_k \in \{0,1\}\ \forall\ k$. Let each design variable x_k denote whether or not ($x_k = 1$ or $x_k = 0$, respectively) independent subcomponent k of the system has been developed and exists for system S_i. Let the cost of developing each of the M independent subcomponents be defined by the M elements of a column vector R. In this case, Eq. (5) provides a simple formula for computing costs for each element $c_{dev,ij}$ of the matrix. Note that in Eq. (5), the open circle ($\circ$) denotes the Hadamard entrywise product operator and the dot ($\cdot$) indicates the dot product operator; in short, this equation simply adds the costs of developing each of the previously undeveloped components.

$$c_{dev,ij} = \vec{R} \cdot \left(\max\left(\vec{S}_j - \vec{S}_i, \vec{0}\right)\right) = \vec{R} \cdot \left(\vec{S}_j - \vec{S}_i \circ \vec{S}_j\right) \quad (5)$$

Note that Eq. (5) inherently assumes that the retirement of subcomponents in the transition from S_i to S_j has a negligible cost. If this assumption is unrealistic and subcomponent retirement (or shutdown) costs can be defined by the M elements of a column vector D, the formula may be modified as in Eq. (6).

$$c_{dev,ij} = \vec{R} \cdot \left(\vec{S}_j - \vec{S}_i \circ \vec{S}_j\right) + \vec{D} \cdot \left(\vec{S}_i - \vec{S}_i \circ \vec{S}_j\right) \quad (6)$$

In the case of Eq. (6), for example, the number of data elements that must be provided by a cost estimation analyst has been changed from N^2 to $2M$. As Figure 14 helps to illustrate, typically this change serves to substantially reduce the pieces of data that a cost analyst must provide: For instance, in the case where a system is defined by $M = 5$ subcomponents, Eq. (6) is more efficient as long as more than $N = 3$ configurations are under consideration. If shutdown costs are not relevant, Eq. (5) is more efficient as long as more than $N = 2$ configurations are under consideration. Generally, N will be substantially larger than these break-even values if a significant trade-space exploration is to be conducted. In fact, if the *full* space of possible configurations is to be explored for this binary example, then by definition $N = 2^M$.

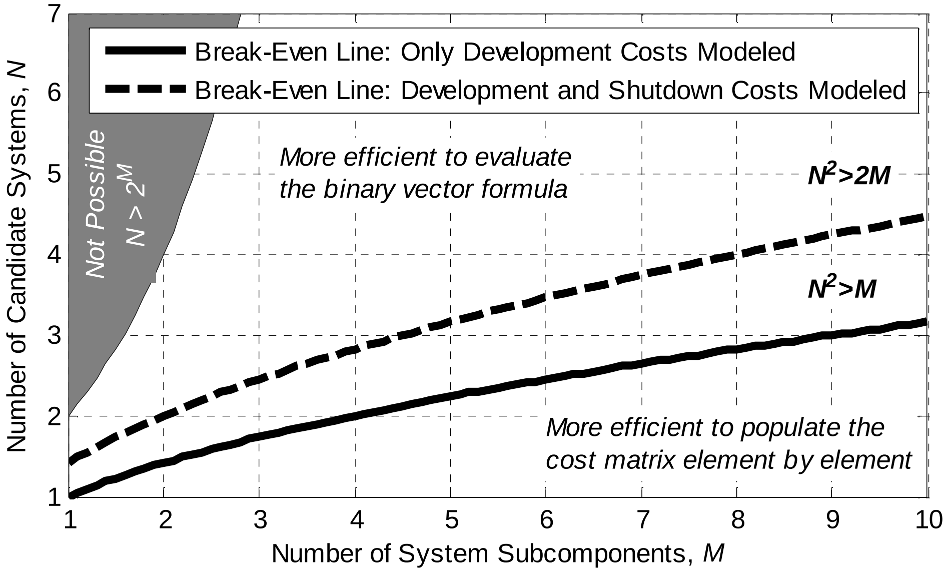

Figure 14. The *M-N* space, indicating for the binary subcomponent scenario regions in which it is more efficient to use Eqs. (5) or (6) rather than populate the cost matrix element by element.

It should be emphasized that this binary subcomponent example is intended only to illustrate one straightforward method for populating the development cost transition matrix from more basic pieces of information. Extensions to this basic form are clearly possible (for example, if the subcomponents are not independent and having developed one for the current time period offsets costs of developing another for the next time period). For the remainder of this thesis, no assumptions are made regarding how the development cost matrix is populated.

4.1.2.2. The Recurring Cost Transition Matrix

A second component to transition costs is the recurring cost, which typically accounts for the production and operation of a fully developed system. These costs too can be represented in an $N \times N$ matrix, and can be decomposed into two lower-level costs of production and operation.

If we extend the binary subcomponent model and define a configuration by a set of nonnegative integer-valued design variables, we have $S_i = [x_1\ x_2\ \ldots\ x_M]^T$, where x_k

$\in \{0, 1, 2, \dots\}\ \forall k$. Such a representation would be useful in describing, for example, an engineering system consisting of multiples of subcomponents. In this case, the design variables of S_i directly describe the number of subcomponents that must be produced in order to produce the system by the next time period. If subcomponents in existence during the current period cannot be effectively reused into the next period, then production costs become a function only of the configuration decision for the next period, or the column of the cost matrix.[*] In this case, the production component of the recurring cost transition matrix can be represented by the $N \times N$ matrix C_{prod}, described by Eqs. (7)-(8). Note that Q is a column vector of per-unit production costs, with dimensions $M \times 1$ and individual elements q_k. Equation (7) applies linear algebra and assumes no learning effects during production, while Eq. (8) demonstrates how learning effects can be incorporated using the Wright learning curve model with learning percent g [80]. Note that $s_{j,k}$ denotes individual element k of column vector S_j.

$$C_{prod} = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix}_{N\times 1} \vec{Q}^T \left[\vec{S}_1\ \vec{S}_2 \dots \vec{S}_N\right] \tag{7}$$

$$c_{prod,ij} = \sum_{k=1}^{M} q_k s_{j,k}^{g+1} \tag{8}$$

If it is assumed that the configuration under development in the current period is to become the operational system in the next period, then operations costs can be accounted for through an $N \times N$ matrix C_{ops} in which each column is identical. Unfortunately, in many scenarios it is unrealistic to simply add operations costs for each subcomponent of the configuration; for example, operations costs may be nonlinear functions of the total system investment cost. In this case, per-period operations costs

---

[*] If this is not the case, an extension to Eqs. (7)-(8) to include dependence on the rows is possible.

$c_{ops,i}$ for each configuration i must be estimated with application-specific tools and converted to C_{ops} via Eq. (9).

$$C_{ops} = \left[c_{ops,i}\right]_{N\times 1}\left[1\cdots 1\right]_{1\times N} \qquad (9)$$

The recurring cost matrix is simply the sum of the production and operations cost matrices, as in Eq. (10).

$$C_{rec} = C_{prod} + C_{ops} \qquad (10)$$

4.1.2.3. The Total Cost Transition Matrix

With nonrecurring (development) and recurring (production and operation) costs now defined in matrix format, the two can be added (as in Eq. (11)) to form the total cost transition matrix *C*. This matrix accounts for all costs incurred over the subsequent time period as the result of the decision to transition from developing system configuration S_i to developing system configuration S_j.

$$C = C_{dev} + C_{rec} \qquad (11)$$

4.1.3. Analyzing the Cost Transition Matrices

The data represented by the cost transition matrices can be analyzed, visualized, and related to flexibility in several useful ways. To illustrate, this section will assume a simple, notional scenario in which a government of a small country is contemplating options for government satellite systems to develop within the next eight years. Two different types of satellites are under consideration: communications satellites and reconnaissance satellites. Producing up to two of each satellite is considered feasible. In this case, a "configuration" will be defined by the number of communications satellites ($x_1 \in \{0,1,2\}$) and number of reconnaissance satellites ($x_2 \in \{0,1,2\}$) to be developed; that is, $S_i = [x_1\ x_2]^T$ with $i \in \{1,2,3,4,5,6,7,8,9\}$ as noted in Table 4.

Suppose the cost of development is \$50 million for a communications satellite and \$300 million for a reconnaissance satellite, and that each of these costs is independent of whether the other satellite has been developed. Shutdown costs will be neglected. In this case, Eq. (5) can be applied directly, with $R = [50\ 300]^T$ and with the resulting C_{dev} matrix shown in Table 5. Similarly, assume that production (including launch) costs are \$100 million for the communications satellite and \$200 million for the reconnaissance satellite, and that learning effects are negligible such that $Q = [100\ 200]^T$ and Eq. (7) can be used directly to calculate the production costs in Table 4. Finally, assume that operations costs are a nonlinear function of the total development cost of the system, such that per-period operations costs are as given in Table 4. The resulting matrices C_{rec} and C are shown in Table 6 and Table 7, respectively.

Table 4. Configuration Definitions for Satellite Example.

Config. ID No. (i)	Number of Communications Satellites (x_1)	Number of Reconnaissance Satellites (x_2)	Production Costs ($c_{prod,xi}$), \$M	Operations Costs ($c_{ops,ix}$), \$M / period
1	0	0	0	0
2	0	1	200	310
3	1	0	100	133
4	1	1	300	372
5	0	2	400	392
6	2	0	200	191
7	1	2	500	449
8	2	1	400	412
9	2	2	600	486

Table 5. C_{dev} for Satellite Example. Costs are in millions of dollars.

		To Configuration								
		1	2	3	4	5	6	7	8	9
From Configuration	1	0	300	50	350	300	50	350	350	350
	2	0	0	50	50	0	50	50	50	50
	3	0	300	0	300	300	0	300	300	300
	4	0	0	0	0	0	0	0	0	0
	5	0	0	50	50	0	50	50	50	50
	6	0	300	0	300	300	0	300	300	300
	7	0	0	0	0	0	0	0	0	0
	8	0	0	0	0	0	0	0	0	0
	9	0	0	0	0	0	0	0	0	0

Table 6. C_{rec} for Satellite Example. Costs are in millions of dollars.

		To Configuration								
		1	2	3	4	5	6	7	8	9
From Configuration	1	0	200	100	300	400	200	500	400	600
	2	310	510	410	610	710	510	810	710	910
	3	133	333	233	433	533	333	633	533	733
	4	372	572	472	672	772	572	872	772	972
	5	392	592	492	692	792	592	892	792	992
	6	191	391	291	491	591	391	691	591	791
	7	449	649	549	749	849	649	949	849	1049
	8	412	612	512	712	812	612	912	812	1012
	9	486	686	586	786	886	686	986	886	1086

Table 7. C for Satellite Example. Costs are in millions of dollars.

		To Configuration								
		1	2	3	4	5	6	7	8	9
From Configuration	1	0	500	150	650	700	250	850	750	950
	2	310	510	460	660	710	560	860	760	960
	3	133	633	233	733	833	333	933	833	1033
	4	372	572	472	672	772	572	872	772	972
	5	392	592	542	742	792	642	942	842	1042
	6	191	691	291	791	891	391	991	891	1091
	7	449	649	549	749	849	649	949	849	1049
	8	412	612	512	712	812	612	912	812	1012
	9	486	686	586	786	886	686	986	886	1086

4.1.3.1. Development Cost Transition Matrix

A helpful visualization of the switching or development cost data in Table 5 is provided in Figure 15. In this figure, each vertical line indicates the range of switching costs from a given configuration, defined by the rows of Table 5. Solid dots indicate minimum and maximum values, open circles indicate mean values, and triangles indicate median values. Each vertical line is located horizontally at the cost needed to develop the configuration from scratch (in this case, Config. 1). For example, if no system currently exists and a decision-maker chooses to develop Config. 2 (involving only the reconnaissance satellite), a cost of \$300 million is incurred (on the *x*-axis), and the cost to switch configurations in the future varies from \$0 to \$50 million, depending on which future configuration is chosen. In contrast, if the decision-maker instead chooses to develop Config. 3 (involving only the communications satellite), a cost of $50 million is initially incurred, and the cost to switch configurations in the future varies from $0 to $300 million. Thus, Figure 15 empirically confirms the intuitive trend that future switching costs can often be reduced by earlier investments. More abstractly, this confirms the early observations of Hart [36] and Stigler [37] that flexibility (the ability to easily modify a system, of which switching cost is an inverse indicator) comes at a cost.

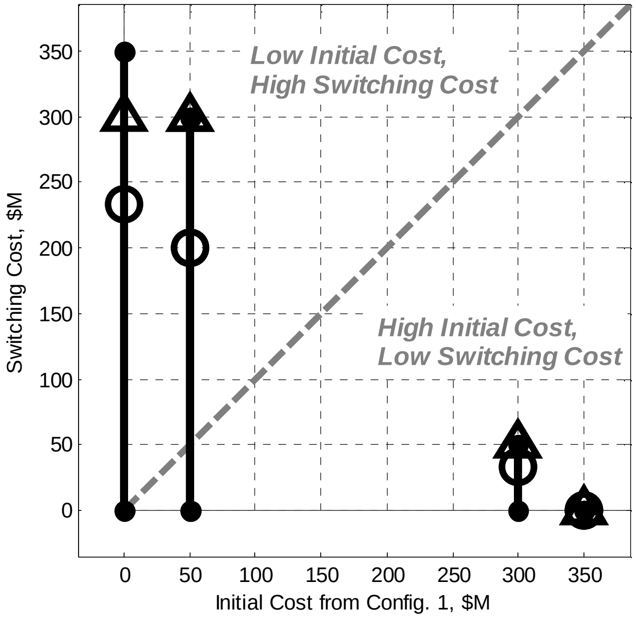

Figure 15. Switching Cost vs. Initial Cost from Config. 1 (S_1) for the satellite example. *Each vertical line indicates the range of switching costs from a given configuration; some configurations overlap. Solid dots indicate minimum and maximum values, open circles indicate mean values, and triangles indicate median values.*

A similarly interesting set of data that can be obtained from the development cost transition matrix is shown in Table 8. This table shows the ratio of $c_{dev,ij}$ to $c_{dev,1j}$, expressed as a percentage. In other words, recalling that configuration $i = 1$ refers in this example to the “do nothing” configuration, this is the cost savings that results from starting with Configuration i to reach Configuration j rather than starting with nothing. For example, Table 8 indicates that starting with Config. 3 (the one-communications-satellite configuration) makes development of Config. 4 (the communications-plus-reconnaissance-satellite configuration) 14% less expensive than if Config. 4 were developed from scratch. Values of 100% in Table 8 indicate that no additional development is required to reach Configuration *j* from Configuration *i*; this occurs frequently toward the bottom of the example matrix because, as they were numbered in Table 4, more capable and demanding configurations were generally listed later. Also by definition, values of 100% occur along the diagonal, where no additional development is required to remain in the same configuration.

Table 8. Development Cost Savings Matrix. *Elements* $c_{devpct,ij}$ *indicate the percent of development costs saved in reaching Configuration j by starting from Configuration i rather than nothing (Config. 1).*

		To Configuration								
		1	2	3	4	5	6	7	8	9
From Configuration	1	▨	▨	▨	▨	▨	▨	▨	▨	▨
	2	▨	100%	0%	86%	100%	0%	86%	86%	86%
	3	▨	0%	100%	14%	0%	100%	14%	14%	14%
	4	▨	100%	100%	100%	100%	100%	100%	100%	100%
	5	▨	100%	0%	86%	100%	0%	86%	86%	86%
	6	▨	0%	100%	14%	0%	100%	14%	14%	14%
	7	▨	100%	100%	100%	100%	100%	100%	100%	100%
	8	▨	100%	100%	100%	100%	100%	100%	100%	100%
	9	▨	100%	100%	100%	100%	100%	100%	100%	100%

Before continuing, it is worth making one final note about this development cost matrix. Recall that earlier it was mentioned the configuration space can be roughly conceptualized as a multidimensional map, with similar system configurations grouped together on the map and unlike configurations distant from each other. While it is tempting to consider the possibility that the switching cost matrix of Table 5 might form the basis for drawing such a map, this is not possible. Note first that Table 5 is not symmetric: Movement between two configurations might be expensive in one direction and inexpensive (or zero) in the other direction. This alone precludes the use of switching cost as a true distance measure. Furthermore, although the triangle inequality is indeed fulfilled in the Table 5 example, it is not necessarily satisfied for all reasonable development cost matrices. For example, starting at a particular Configuration A, although improbable, it is not impossible for the development of an intermediate Configuration B to dramatically reduce costs of arriving at Configuration C, such that $c_{dev,AC} > c_{dev,AB} + c_{dev,BC}$. This further precludes the use of switching cost as a distance measure. Because of these observations, this thesis retains the definition of configuration space distances in the sense described in Section 4.1.1 and Figure 13; in particular, Figure 16 shows a visualization of this satellite example's configuration state space.

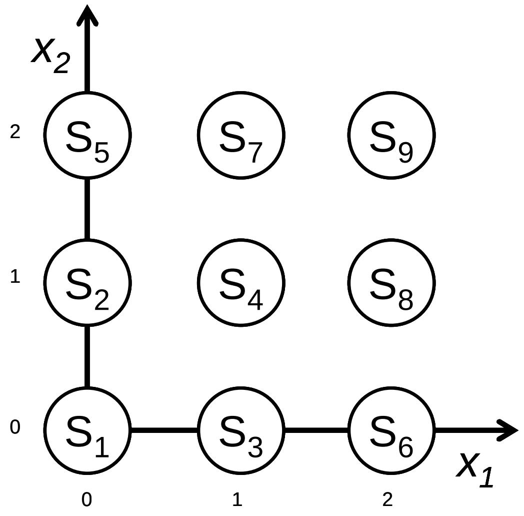

Figure 16. Configuration state space for the satellite example.

4.1.3.2. Total Cost Transition Matrix

Recall that the total cost transition matrix C accounts for all costs incurred over a subsequent time period as the result of the decision to transition from developing system S_i to developing system S_j. This matrix will be particularly important in steps 4-5 of this framework, and a helpful synthesis of the information contained in this matrix with the two-period state-centric notion of flexibility is shown in Figure 17. Here, each node represents one of the configurations considered in the design space, the color of which is indicative of the cost to develop and produce it from the "nothing" configuration (Config. 1). The nodes are arranged in a configuration similar to that in Figure 16, but with slight geometric modifications to avoid confusion when ascertaining which arrows connect which nodes. Above each of the four groups of nodes is a budget, and for every element of the total cost transition matrix less than or equal to the given budget, a directed link is drawn. In cases where the total cost on the diagonal of the matrix is less than or equal to the budget, a dark circle is drawn around the appropriate node. For example, the top left portion of Figure 17 shows that, if the currently-fielded architecture is Config. 3, a $400 million budget for a given eight-year period would allow the decision-maker to transition to Configs. 1 or 6, or to remain in Config. 3.

A natural observation from Figure 17 is that, as budget is increased, more links become available. That is, as the decision-maker has more resources available, more options exist. The total number of links in the graphs of Figure 17 increases from 12 at the $400 million budget to 22 at the $550 million budget, 40 at the $700 million budget, and 60 at the $850 million budget. Eventually, at a large enough budget, all 81 links would appear. Linking this to the two-period state-centric concept of flexibility, a clear indicator of the flexibility of a given configuration i is the number of links or transitions available to it for a given budget b (the number of "outs" available, which will be denoted $\Phi_i(b)$). This indicator is plotted in Figure 18. Here, the starting configuration (node) is shown on the x-axis, and the number of available transitions from that node to other nodes is shown on the y-axis. Note that these available transitions do not increase linearly with budget; for example, adding $150 million of budget to $400 million results in no increases to the transitions available from Configs. 3 and 6, while adding the same amount to $700 million increases the transitions available to each by 75%.

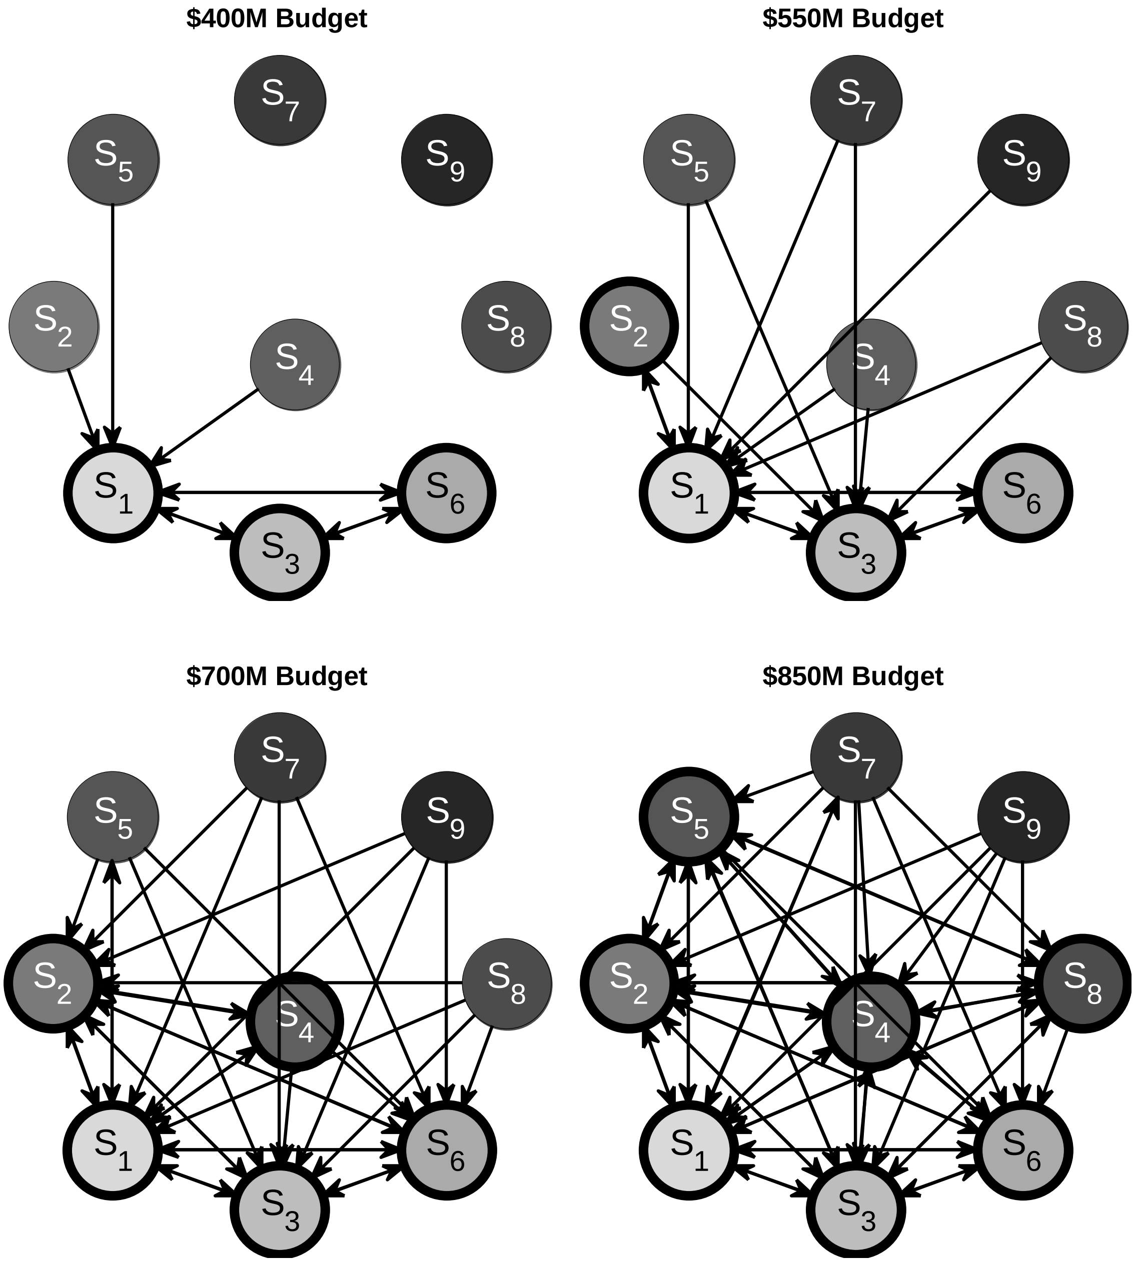

$400M Budget

$550M Budget

$700M Budget

$850M Budget

Figure 17. Available configuration transitions for $400, 550, 700, and 850 million budgets.

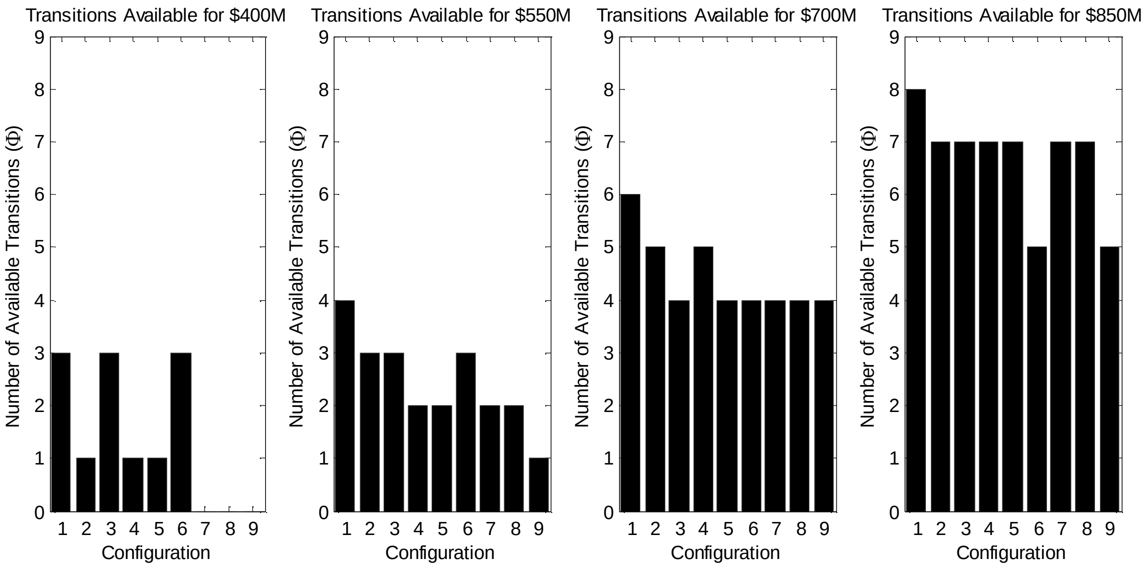

Figure 18. Number of available transitions (Φ) for $400, 550, 700, and 850 million budgets.

If available budget is considered on a continuum instead of four discrete intervals, the data in Figure 19 result. This figure shows the number of available transitions as a function of available budget, where data for each configuration is represented by a single line. For example, the figure shows that for a per-period budget of $200 million, Config. 1 (the "nothing" configuration) has Φ = 2 transitions available, Configs. 3 and 6 each have Φ = 1 available transition, and all other configurations have no available transitions (i.e., the available budget is insufficient even to support operation of the current configuration into the next period). It also shows that by a budget of $1.1 billion, any configuration can be reached from any other configuration since all configurations have 9 available transitions.

An interesting characteristic visible in Figure 19 is that Configs. 1, 3, and 6 tend to have significantly more transitions available than the other configurations for per-period budgets below $500 million. These configurations have in common the fact that they have no reconnaissance satellites to incur large operations costs; as a consequence, while all other configurations must spend between $300 and $500 million simply to operate, Configs. 1, 3, and 6 are able to use this budget to effect transitions to other configurations.

To develop this observation more fully, Figure 20 shows a subset of the configurations visible in Figure 19, in particular Config. 3 (one communications satellite only) and Config. 4 (one communications satellite and one reconnaissance satellite). As expected from Figure 19, Config. 4 has fewer transitions available than Config. 3 at low budgets because of its operations cost requirements. However, at a per-period budget of about $570 million, a reversal occurs. Above this budget, Config. 4 always has at least as many transitions available as Config. 3. At high budgets, the greater developed capability of Config. 4 (i.e., the existing reconnaissance capability) translates into lower development transition costs. Thus, this graph serves to illustrate that *flexibility is not solely a function of the engineering configuration a decision-maker selects, but also a function of the resources that are available to change that configuration*. In this case, the low-development-cost (no-reconnaissance-satellite) configurations are equally or more flexible than the high-development-cost configurations when resources are scarce because the low-development-cost configurations incur lower fixed operations costs. However, as financial resources become more abundant, the configurations that include reconnaissance satellites permit more flexibility because they already have a reconnaissance capability which the low-development-cost (no-reconnaissance-satellite) must spend resources to develop.

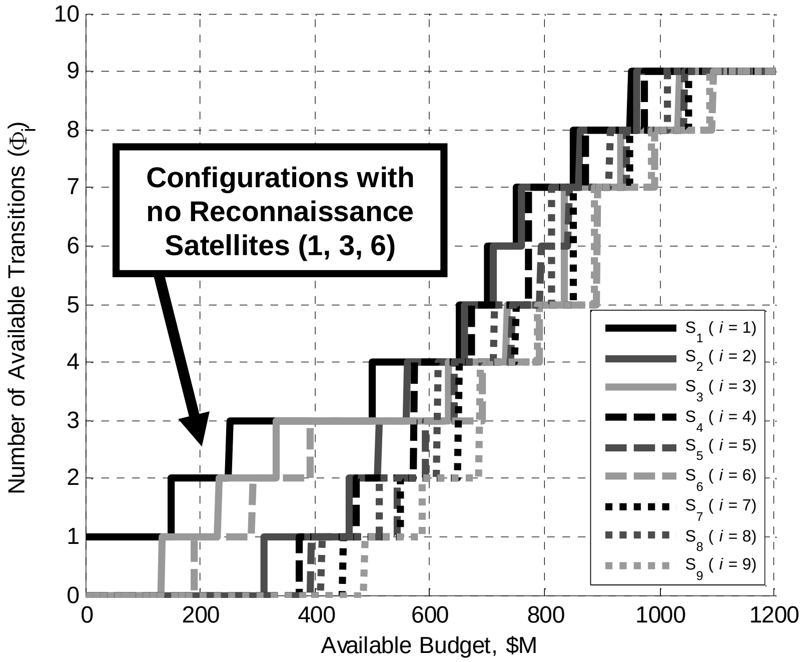

Figure 19. Available configuration transitions as a function of available per-period budget.

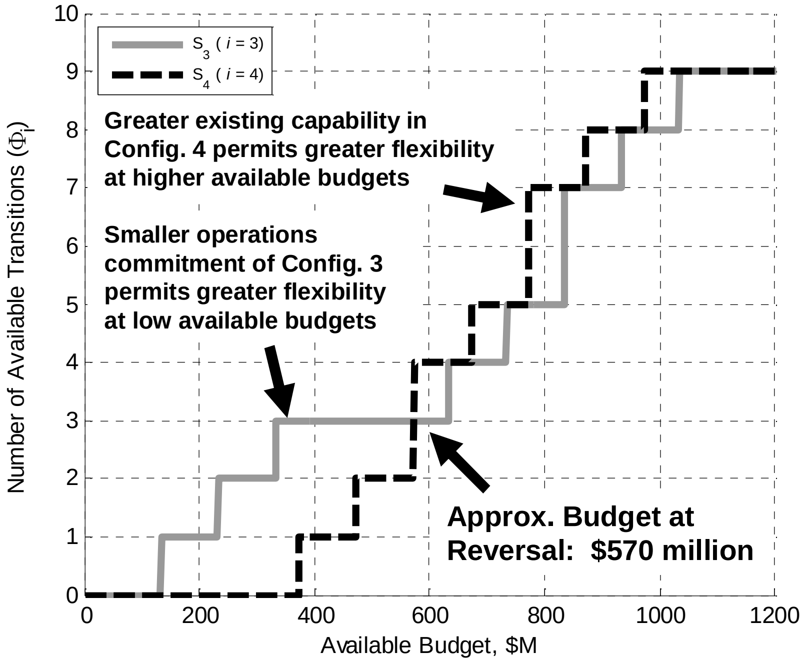

Figure 20. Available configuration transitions for Configs. 3 and 4 as a function of available per-period budget, illustrating a "flexibility reversal".

4.1.3.2.1. Sensitivity to Budget

One analysis that Figure 19 and Figure 20 enable is examination of the sensitivity of a configuration's number of available transitions Φ (roughly speaking, the sensitivity of a configuration's flexibility) to changes in the allocated per-period budget. Such an analysis is particularly useful to entities interested in selecting an appropriate budget level for a multi-period program or project.

Figure 21, which is derived from Figure 20, tracks the number of transitions gained from a \$250 million per-period budget increase at each of the budget levels in Figure 20. Formally, this is the forward difference $\Delta_h\Phi_i(b)$ given in Eq. (12) (cf. [81]-[85]), with h = \$250 million. This difference is used in lieu of the derivative (that is, the limit of $\Delta_h\Phi_i(b)/h$ as $h \to 0$) because the derivatives of the functions in Figure 20 take values only of zero or infinity and are not insightful to examine.

$$\Delta_h\Phi_i(b) = \Phi_i(b+h) - \Phi_i(b) \tag{12}$$

Note that this forward derivative, plotted on the y-axis of Figure 21, illustrates a distinct difference between Config. 3 (one communications satellite only) and Config. 4 (one communications satellite and one reconnaissance satellite): While $\Delta_h\Phi_4$ exhibits an overall unimodal behavior, having from \$525-570 million a gain of 5 transitions per \$250 million budget added, $\Delta_h\Phi_3$ exhibits a more bimodal behavior. Config. 3 has a high forward difference for low budgets that then disappears to $\Delta_h\Phi = 0$ before rising again at higher budget levels. In effect, the $\Delta_h\Phi = 0$ valley illustrates that there exists a capability gap for Config. 3: A certain threshold of resources must be invested in order to permit any options beyond a communications-satellite-only capability. For a decision-maker considering a \$250 million per-period budget increase above an existing \$350 million budget in a situation where Config. 3 already exists, Figure 21 clearly indicates that such a budget increase would provide no additional options. On the other hand, for a decision-maker considering a \$250 million per-period budget increase above an existing \$625 million budget, Figure 21 indicates that such an increase may be justified, as it adds four more options.

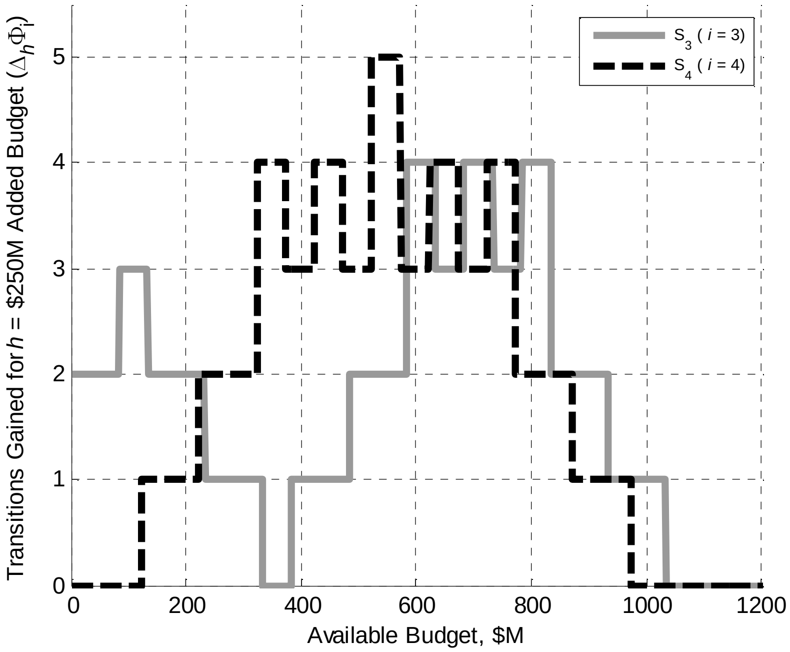

Figure 21. Number of transitions gained when budget on the *x*-axis is raised by $250 million.

Another way to represent the data in Figure 21 is via an elasticity metric. This metric, defined in Eq. (13) for the case of the discrete step size h, indicates the percentage change in the number of transitions Φ that can be achieved by adding budget h divided by the percentage change in the current budget that h represents. Note that this metric is undefined when $\Phi = 0$. If this metric is plotted for Configs. 3 and 4, Figure 22 results. This indicates, for example, that at its peak elasticity, adding $250 million to a $470 million budget for Config. 4 results in the number of transitions increasing relatively 7.5 times more than the budget. Thus, this is a region where flexibility can be very significantly impacted by budget increases. Note that, as in Figure 21, the elasticity curves also fall to zero at the highest budgets since, at these budgets, any configuration state can be reached from any other configuration state.

$$E_{\Delta_h,i}(b)=\frac{\Delta_h\Phi_i(b)/\Phi_i(b)}{h/b} \qquad (13)$$

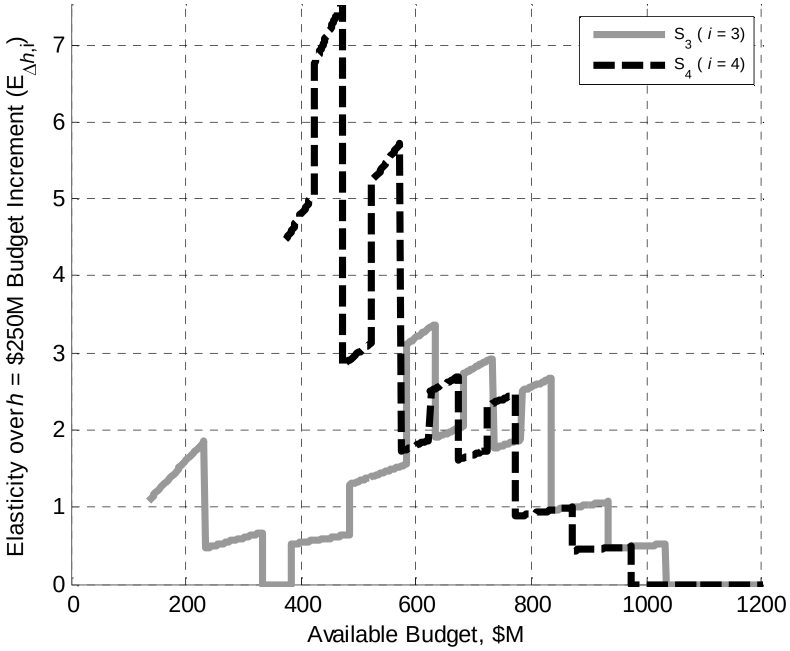

Figure 22. Transition Elasticity for a forward difference of $250 million.

4.1.3.2.2. Transitions for Uncertain Costs

A final extension of the cost transition concepts presented here is provided to illustrate how probabilistic analysis can assist in understanding the robustness of the deterministic transition results illustrated thus far. Suppose, for example, that the $50 and $300 million development cost and $100 and $200 million production cost estimates assumed for *R* and *Q* in Section 4.1.3 are associated with significant degrees of uncertainty. Suppose that each of these four parameters can take values from 25% below to 50% above their baseline values and can be modeled by independent triangularly-distributed random variables with modes equal to the baseline values above. As a result, the available transitions indicated by Figure 19 and Figure 20 are no longer properly described by single deterministic lines, but rather by bands of uncertainty.

These bands are shown in the upper plot of Figure 23 for Configs. 3 and 4. This plot has a format identical to Figure 20, except that bands of uncertainty surround a median near the baseline deterministic result. The bands in Figure 23 are drawn to encompass the 5th to 95th percentile results as obtained from 1,000 Monte Carlo simulations of the triangular input uncertainties. Note that, even in the presence of these substantial uncertainties, the error bands for Config. 3 below a budget of about $550 million do not overlap with those of Config. 4. This is reflected as well in the lower plot of Figure 23, which shows that between a budget of $170 million and $500 million there is near certainty that Config. 3 will have more transitions available than Config. 4. As the available budget is increased beyond $500 million, there is a sharp decline in this probability, until by $720 million the reverse occurs. In this region, there is near certainty that Config. 3 will have fewer transitions available than Config. 4. Thus, in addition to assisting the decision-maker in visualizing the uncertainty in the transition numbers from the deterministic analysis, the results of this probabilistic cost analysis can reveal the existence of "flexibility reversal" scenarios even in the presence of cost estimate uncertainties.

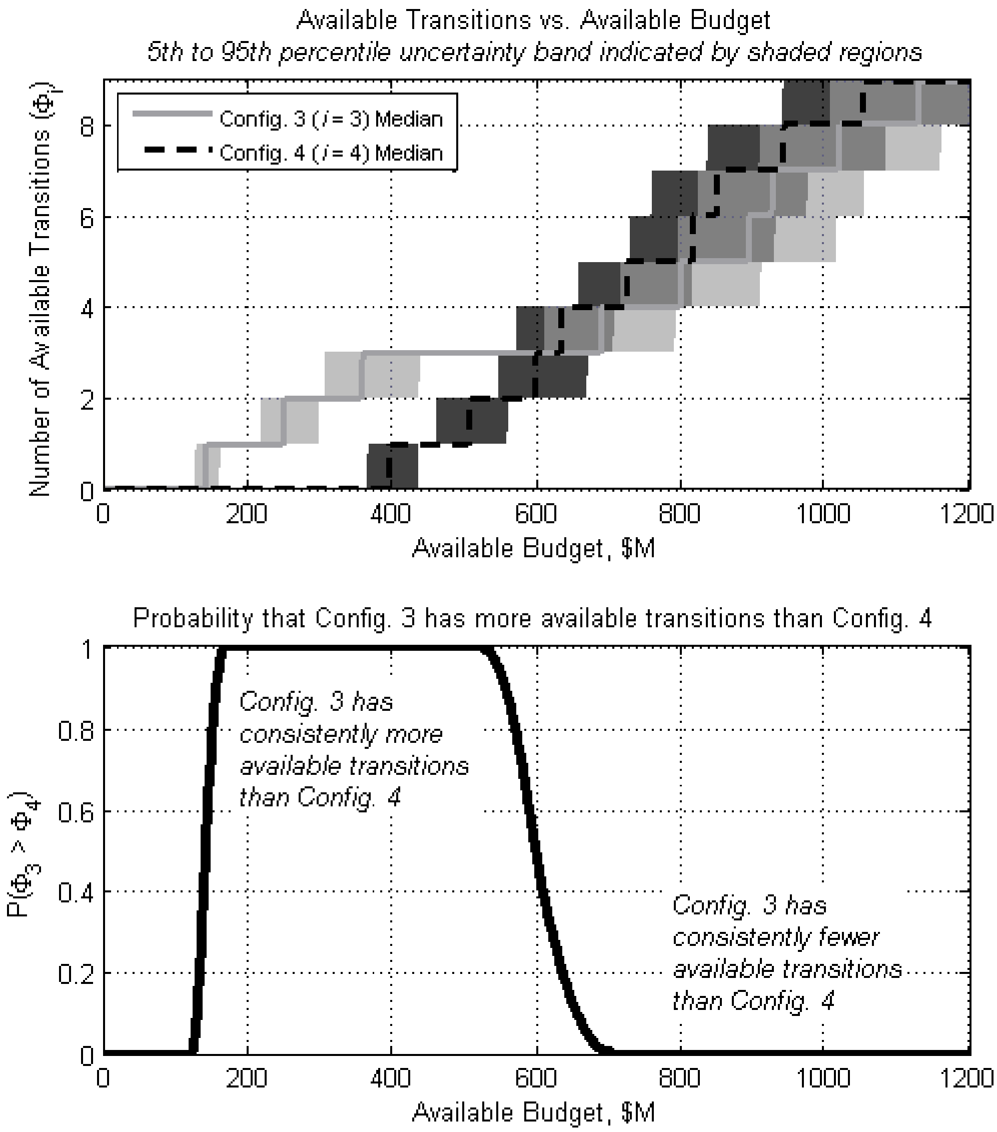

Figure 23. Probabilistic comparison of the transitions available to Configs. 3 and 4.

4.1.4. Limitations of Cost-Only Considerations

In summary, this step of the framework has shown that the two-period state-centric notion of flexibility from previous literature can be adapted to apply to configuration changes for engineering systems or architectures. For an engineering system in which multiple configuration options exist over time, a cost transition matrix can be formed and used to visualize the options that exist for changing the system as a function of available budget. If a single, relatively constant per-period budget is likely to exist for the foreseeable future, that budget can be selected and a diagram such as one of the graphs in Figure 17 can be useful in tracing possible configuration pathways. If the available budget is likely to be subject to change or partially under the control of the decision-maker, the available transitions can be plotted as a function of budget to determine if additional budget would make a substantial difference in the available options. Analysis of these graphs and associated data illustrate how budget itself can drive whether one configuration is more flexible than another, and sensitivity and uncertainty analyses can both be conducted to yield additional useful insight.

At the conclusion of Step 1, it is reasonable to ask: From this information, what conclusions can be drawn about the best initial system configuration to select? Unfortunately, none. To do so requires overcoming two limitations of considering only configurations and cost over a two-period time interval. First, the time horizon of the analysis must be expanded to more than two periods to avoid potentially myopic decision-making.[*] Second, the benefits of being in a given configuration at a given time must be quantified. A limitation of using number of available transitions as an indicator for flexibility is that it contains no information about the value of each configuration in each future time period. As a result, it is possible to manipulate this metric to make certain configurations appear relatively more or less desirable by either (1) including in the state space a large number of physically similar configurations or (2) including in the state space a large number of configurations that are unlikely to have any value in the future. These limitations are resolved in the following two steps of the framework.

---

[*] If this were the only limitation, it might be overcome for example by defining a time horizon of *n* periods and expanding the configuration space into a tree, tracking in total how many configurations are accessible over the *n*-period horizon given an initial configuration.

4.2. Step 2: Define Markovian Demand Environment Evolution Probabilities

Discussed in Step 1 was how the two-period state-centric flexibility framework of previous literature can be adapted to apply to configuration changes for engineering systems or architectures. That step focused on the options available to and under the control of the engineer or decision-maker. However, as recognized in Step 1, information about the utility or value of being in each configuration at a given time period is needed in order to make meaningful conclusions about the suitability of each configuration decision. In order to do this, it is necessary to have information about the environment (in particular, the demand environment) in which system is operating, which is generally *out* of the control of the engineer or decision-maker. Rarely is deterministic prediction of this environment possible, and so this information must generally be in the form of a stochastic model.

Mathematically, this is equivalent to the statement in Eq. (14), i.e., that a performance or utility u (preferably a metric with physical meaning, but not precluded from being a normalized aggregate metric) that is gained in time increment t is a function not only of the configuration S that is operational at that time, but also of the demand environment y that materialized at that time. This environment evolves according to some stochastic process $\{Y(t)\}$. The fact that $u(t)$ is not a function of $S(t)$ alone concisely explains why conclusions about value or utility cannot be drawn from Step 1 alone.

$$u(t) = f\left(S(t), y(t)\right) \tag{14}$$

Thus, the two questions that arise in Step 2 are: What are the environments that $\{Y(t)\}$ can describe, and what kind of stochastic model should be used to describe $\{Y(t)\}$?

4.2.1. Definition of the Demand Environment

As with the definition of configurations in Step 1, what precisely defines a demand environment is application-specific. Ideally, the environment definition would completely describe the current state of the world (or universe). However, since such an extensive definition of the state of the environment would be far from tractable, the analyst may be guided by two practical considerations: First, what major external factors or combination of factors tend to describe demand for the system being considered? For example, in defense applications this might involve the terrain of the theater of operations or the type of enemy combatant, while in disaster relief applications this might involve the type and frequency of various natural and manmade disasters. Second, of these factors or combinations of factors, are some likely to distinguish the performance of some configurations over others, or do they affect all configurations equally? In general, factors that would have little effect on the configuration decision can be neglected.

This framework assumes that the set of possible demand environments $\{Y_i, i = 1, 2, \ldots, K\}$ to which the system of interest may be subject is finite and discrete or can be reasonably approximated as finite and discrete, and that this environment evolves stochastically with time. In the case of the example satellite application, suppose that the demands upon the satellite system are primarily driven by (1) the existence of armed conflicts and (2) the degree to which existing commercial capacities reduce the need for a government satellite capability. In this case, the demand environment might be described by six states, summarized in Table 9. This table characterizes the set of demand environments $\{Y_i, i = 1, 2, \ldots, K\}$, where $K = 6$, by the two drivers above, and translates these qualitative descriptions into a reasonable quantitative implication in terms of the number of government communications and reconnaissance satellites needed. In general, within the table, hostile conflict environments with no available commercial capacity produce the greatest demand for government satellites, while environments in which full capacity is provided by commercial entities produce the least demand.*

Table 9. Demand Environment Definitions for Satellite Example.

Env. ID No. (*i*)	Conflict Environment	Available Commercial Capacity	Implied No. of Government Communications Satellites Needed	Implied No. of Government Reconnaissance Satellites Needed
1	Hostile	None	2	2
2	Hostile	Some	1	1
3	Hostile	Full	0	0
4	Quiescent	None	1	1
5	Quiescent	Some	0	1
6	Quiescent	Full	0	0

4.2.2. The Markovian Stochastic Model

To continue, we address the second question of this step: What sort of stochastic model should be used to describe $\{Y(t)\}$? No doubt the simplest stochastic model for $\{Y(t)\}$ would be a time-ordered set of independent random variables; however, the implication of such a model is that the past has no influence on the future, and it is questionable whether such an assumption is reasonable in most practical situations faced in the space industry. A more general stochastic model for $\{Y(t)\}$ is a Markov chain. Formally, a Markov chain is a time-ordered† set of random variables $\{Y(t)\}$ for which the probability that $Y(t)$ takes some value a depends only on the value of $Y(t\text{-}\Delta t)$, i.e., Y in the previous time period. The possible values for Y must be finite or countable. In a Markovian stochastic process, the past influences the future only through the present

---

* Note that some of these environments, such as environments 3 and 6, have identical satellite requirements. It would be equally valid to define the demand environment states in terms of these requirements, if the transition probabilities were more easily estimable between these environment states.

† Strictly speaking, the variables need not be time-ordered as long as they are ordered by some other monotonically increasing parameter. For the present application, this parameter will be time.

state; if it is necessary to build additional memory into the process, it is possible to do so by expanding the chain's state space (i.e., the definition of the possible values of Y). The conditional probabilities $\mathrm{P}[Y(t) = a \mid Y(t-\Delta t) = b]$ with which values of Y at time $t-\Delta t$ evolve to other values of Y at time t are organized in a probability transition matrix, which for Markov chains is typically assumed constant with time. For further familiarization with Markov chains and its traditional applications, the reader is referred to Refs. [86] and [87].

If sufficient historical data exists, a Markov chain's probability transition matrix can be populated by statistically mining the historical data for the appropriate conditional probabilities. However, if this data does not exist or would take too much in time or resources to obtain, a positive quality to the use of a Markov chain is that the probability transition matrix can be populated via expert judgement without excessive complication.[*]

Suppose that, by use of historical data or expert elicitation, a Markov chain probability transition matrix P for the satellite application is populated as in Table 10. In this matrix, each element P_{ij} indicates the probability that demand will transition from environment i to environment j over one time increment (in this case, eight years, corresponding to the time step assumed in Step 1).

Table 10. Sample Markov Chain Transition Matrix for the Satellite Example Application.

		To Demand Environment, Y_j					
		- 1 - Hostile, None	- 2 - Hostile, Some	- 3 - Hostile, Full	- 4 - Quiescent, None	- 5 - Quiescent, Some	- 6 - Quiescent, Full
From Demand Environment, Y_i	- 1 - Hostile, None	0.10	0.15	0.05	0.40	0.20	0.10
	- 2 - Hostile, Some	0.20	0.10	0.10	0.10	0.30	0.20
	- 3 - Hostile, Full	0.10	0.20	0.05	0.05	0.50	0.10
	- 4 - Quiescent, None	0.20	0.10	0.05	0.30	0.30	0.05
	- 5 - Quiescent, Some	0.10	0.20	0.10	0.15	0.30	0.15
	- 6 - Quiescent, Full	0.05	0.10	0.10	0.20	0.30	0.25

---

[*] One application of such a process is available in Ref. [88].

The Markov chain of Table 10 can be visualized as a set of demand environment states as in Figure 24. In this figure, high-probability transitions are represented as thick dark links and low-probability transitions are represented as thin light links.* From each state, a green link identifies the most likely (highest-probability) transition(s). For example, from this diagram it can be seen that environments 4 and 5 act much like a sink: The most likely transitions from each environment all lead to one of these two states.

It is also useful, for reference, to observe two properties of the demand environment Markov chain. First, most practical demand environment Markov chains will involve probability transition matrices with strictly positive elements; rarely will the probability of transition from one environment to another be *exactly* zero. These will thus be single-class chains that are irreducible, positive recurrent, and aperiodic (therefore ergodic). As a consequence, a unique long-run probability π_j of being in demand environment j will exist [86] and can be found via Eq. (15). While behavior of the demand environment an infinitely long time into the future is not often of principal interest to the decision-maker, it can provide the analyst helpful intuition regarding the direction toward which the demand will eventually tend as a consequence of the assumed matrix P. In the case of the Markov chain in Table 10, the right half of Figure 24 displays the stationary probabilities of existing in each demand environment. Note that in the long term, the quiescent conflict environment with some available commercial capacity is the most likely (30.3%), while the hostile conflict environment with full available commercial capacity is nearly four times less likely (7.9%). Also, it is insightful to note that the Markov chain of Table 10 implies that the conflict environment in the long term is more often than not (64.7% of the time) quiescent.

---

* In the field of combustion, a similar type of visualization, called a reaction pathway diagram, is used to convey information about the relative importance of elementary reactions in more complex reaction mechanisms (e.g., see Ref. [89]).

$$\lim_{n\to\infty} P_{ij}^{n} = \pi_j = \sum_{i=1}^{K} \pi_i P_{ij} \qquad (15)$$

Second, it can be insightful to calculate the entropy rate H' of the demand environment Markov chain. Entropy rate is a quantity with origins in the field of information theory that serves as an indicator of the degree of uniform randomness introduced at each time step (or other index) of a stochastic process. In the case of a Markov chain with a stationary distribution defined by π_j, $\forall j \in \{1,2,\ldots, K\}$, the entropy rate is calculated as in Eq. (16) [90] and is reported in bits. Note that the maximum entropy rate of a K-state Markov chain occurs when its transition matrix is completely uniform, i.e., $P_{ij} \equiv 1/K$. In such a case, $H' = \log_2 K$, which serves as a helpful upper bound for understanding the randomness indicated by the entropy rate. In the case of the Markov chain of Table 10, the entropy rate is a relatively high 2.36 bits (of a possible $\log_2 6 = 2.58$ bits) per eight-year time period. As a result, this particular demand environment model can be characterized by significant uncertainty over its eight-year time step.

$$H'(\{Y(t)\}) = -\sum_{i=1}^{K} \pi_i \sum_{j=1}^{K} P_{ij} \log_2 P_{ij} \qquad (16)$$

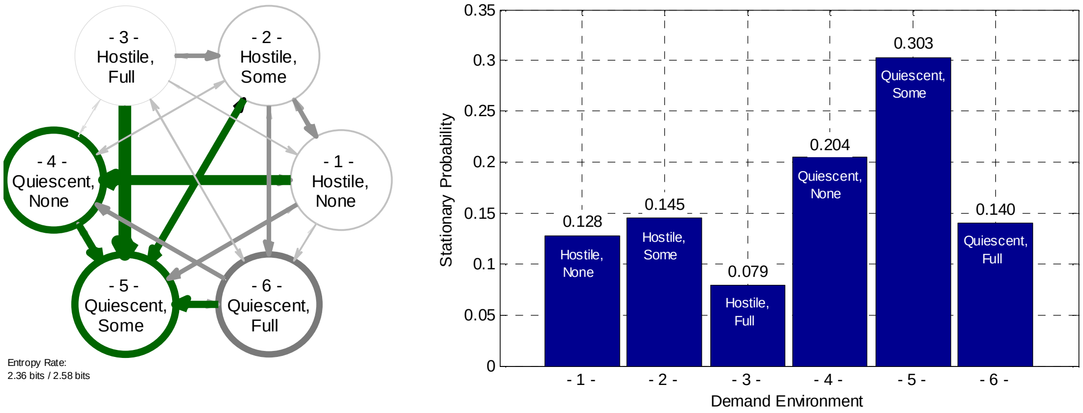

Figure 24. Visualization of the Markov Chain (left) and Stationary Distribution (right) of Table 10.

4.3. Step 3: Define State-Dependent Performance Matrix

Step 2 presaged the fact that at least one utility or performance metric u is necessary to make decisions that are properly informed by both performance and cost considerations. The functional dependence of the performance metric as provided in Eq. (14) (copied below for convenience) assumes an explicit dependence on the operational configuration S and demand environment y.[*] As a result, a natural representation of the function $u = f(S, y)$ is in the form of an $N \times K$ matrix (recalling that N is the cardinality of the set of possible configurations and K is the cardinality of the set of possible demand environments).

$$u(t) = f\big(S(t), y(t)\big) \qquad (14)$$

Such a matrix, denoted U, is shown in Table 11 for the satellite illustration. In this case, the chosen performance metric is the number of demanded satellites that are available (and utilized). For example, if the demand in one time period is associated with a hostile conflict environment and no available commercial satellite capacity (Environment 1, in column 1) and the operational configuration during that time period has two communications and two reconnaissance satellites available (Config. 9, in row 9), then the decision-maker accumulates the successful utilization of all four available satellites. As a consequence of the specification of this matrix, the decision-maker will be incentivized to place satellites into orbit that will meet likely demands.

---

[*] Note that time is not explicitly captured in this dependence, i.e., $u(t) \neq f(S(t),y(t),t)$. In many cases, this lack of explicit time dependence in u has few or no practical modeling limitations. However, if such a dependence is indeed important, it can be incorporated by integrating time into the definition of the demand environment and/or configuration state.

Table 11. *U* for Satellite Example. Metric indicates the number of demanded satellites available.

	Demand Environment					
Configuration	- 1 - Hostile, None	- 2 - Hostile, Some	- 3 - Hostile, Full	- 4 - Quiescent, None	- 5 - Quiescent, Some	- 6 - Quiescent, Full
1	0	0	0	0	0	0
2	1	1	0	1	1	0
3	1	1	0	1	0	0
4	2	2	0	2	1	0
5	2	1	0	1	1	0
6	2	1	0	1	0	0
7	3	2	0	2	1	0
8	3	2	0	2	1	0
9	4	2	0	2	1	0

Before concluding the discussion of Step 3, it is worth making two final notes: First, the example metric of Table 11 is just one of many that might be considered for this example. For example, a decision-maker may also be interest in a cumulative binary metric that indicates a 1 or 0 in each time period depending on whether satellite needs were fully met; over the long term, such a metric would indicate the percentage of time that the system fully meets the demands placed on it. Another two examples would be metrics that are specific to communications or reconnaissance satellites (e.g., (1) number of demanded communications satellites available or (2) number or demanded reconnaissance satellites available). Any such metric can easily be accounted for via a matrix such as in Table 11. Second, although the example application shown here employs only one performance metric, the theoretical development of Steps 4 and 5 should make it evident that incorporation of multiple performance metrics and matrices U_1, U_2, …, U_m can be fully accommodated and integrated within this thesis' framework.

4.4. Step 4: Decision Support Analysis

At this point, enough information has been specified in Steps 1-3 to simulate a system as it changes in response to a decision-maker's actions over time. The space of possible configurations and the costs of moving between them over each time step have been defined in Step 1; the space of possible demand environments and the probabilities of moving between them over each time step have been defined in Step 2; and the performance accumulated when a given configuration is subjected to a given demand environment has been specified in Step 3.

However, it has not yet been specified which actions the decision-maker will (or should) take during such a simulation. As a result, Step 4 has the dual purposes of (1) defining this simulation and (2) solving for the decision-maker actions that will result in the "best" possible outcome.

The definition of the simulation, given a T-period time horizon, Δt time step, and a configuration $S(t_0)$ and demand environment $y(t_0)$ at initial time t_0, is provided by the following two dynamics equations (or "equations of motion") in Eq. (17). The first indicates that the configuration selected for development by the decision-maker's action a at time t becomes the operational configuration S at time $t+\Delta t$. The second indicates that the demand environment that materializes at time $t+\Delta t$ is distributed as indicated by the row of the Markovian probability transition matrix P corresponding to the demand environment $y(t)$ in the current time period.

$$\begin{aligned} S(t+\Delta t) &= a(t) \\ Y(t+\Delta t) &\sim P_{y(t),1..K} \end{aligned} \tag{17}$$

The goal of the more challenging decision support task is mathematically expressed in Eq. (18). This task is multiobjective in nature: The decision-maker will typically wish to minimize cost and maximize one or more performance metrics[*] over the T-period horizon of the system. This is represented in Eq. (18) by the maximization of a vector-valued function of a vector sum. The first element of the vector within the summation is the element ($S(\tau)$, $a(\cdot)$) of the total cost matrix C; this expresses the fact that at time index τ, the system configuration is $S(\tau)$ and the decision-maker has incurred some cost via the decision a. The second element is element ($S(\tau)$, $y(\tau)$) of the performance matrix U; this expresses the fact that at time index τ, a performance benefit has accrued as a result of the system being in configuration $S(\tau)$ while the demand environment is in state $y(\tau)$. Additional elements of the vector are allowed, for example, to account for performance measures U_2, U_3, …, U_m as described in Step 3. For notation convenience, in Eq. (18) only the maximum function is used and minimum-preferred cost and performance objectives must be negated. A conversion from the “real-world” time t to the index τ is also provided. The vector-valued function F exists to convert its random-variable argument into a vector of representative deterministic values; in most cases, it will be convenient and computationally necessary to select this function to be the expected-value operator E; however, in principle F can be any function of the long-term cost and performance sums.

$$\begin{gathered}\max_{a(\cdot)\in\{S_i\}} F\left(\sum_{\tau=1}^{T}\begin{bmatrix}-c_{S(\tau),a(\cdot)}\\ u_{S(\tau),y(\tau)}\\ \vdots\end{bmatrix}\right)\\ t=t_0+(\tau-1)\Delta t\end{gathered}\tag{18}$$

In general, it will not be possible to simultaneously minimize cost and maximize performance through any particular set of actions or decisions $a(\cdot)$. Thus, the solution to the problem posed in Eq. (18) is not a single answer for $a(\cdot)$, but a set of decisions that

---

[*] The decision-maker may also have multiple cost metrics, which may be bookkept as “negative” performance measures.

depends on the decision-maker's preferences for one objective over another. This formulation is thus one of a multi-objective optimization problem,* the solution of which comprises a set of non-dominated points in the objective space. These points form a multi-dimensional Pareto frontier, on which the performance of one objective cannot be improved without the sacrifice of another. In mathematical terms, a scalar or vector input x^* to the vector-valued function to be maximized f is said to be Pareto-optimal (non-dominated) if there exists no other input x such that (1) $f_j(x^*) \leq f_j(x)$ for all j and if (2) $f_j(x^*) < f_j(x)$ for at least one j, where f_j is the j^{th} element of the vector-valued f. The literature for single-period multi-objective optimization problems is well-established, and the reader is referred to Refs. [87] and [91] for helpful introductory reference material.

Finally, some precision must be added to specify what is meant by the term $a(\cdot)$. The term a indicates the action or decision made at a given time, as indicated in Eq. (17). As specified in Eq. (18), the values taken by a are drawn from the set of available configurations $\{S_i\}$. Clearly, a is not simply a constant to be solved for and should be a function (to be solved for) of some other variable or variables; otherwise, a single configuration would be fielded for all time, which is in general an unrealistic expectation for decision-maker behavior. Thus, the question addressed by Steps 4A and 4B centers around: Of what variables should the decision-maker's actions be a function? Step 4A takes a view traditionally taken during long-term roadmapping analysis that this variable should be time, akin to open-loop control. Step 4B takes a more complete but computationally more expensive view that this variable should be the total system and environment state, akin to closed-loop control.

---

* Other common names for this problem in the literature include multiple-attribute decision-making, multiple-criteria decision-making, and multiple-objective decision-making.

4.4.1. Find Pareto-Optimal "Open-Loop" Paths

Figure 25 frames the problem posed by Eqs. (17) and (18) graphically and, combined with Eqs. (17) and (18), suggests a method for simulating and solving for a Pareto-optimal path of actions $a(\tau)$. In Figure 25, time progresses in discrete increments of duration Δt along the x-axis. In each period, the bottom two rows indicate the operational configuration and demand environment, which interact to produce per-period performance values $u_{S(\tau),y(\tau)}$. The top row indicates the decision to be made about which configuration to develop in the current period, which directly affects current costs and determines what configuration will be operational in the subsequent time period, thereby affecting subsequent performance and transition costs. This top row, posed as a sequence of T decisions over T timesteps, suggests that a reasonable covariate for the function a is indeed time, i.e., $a(\cdot) = a(\tau)$.

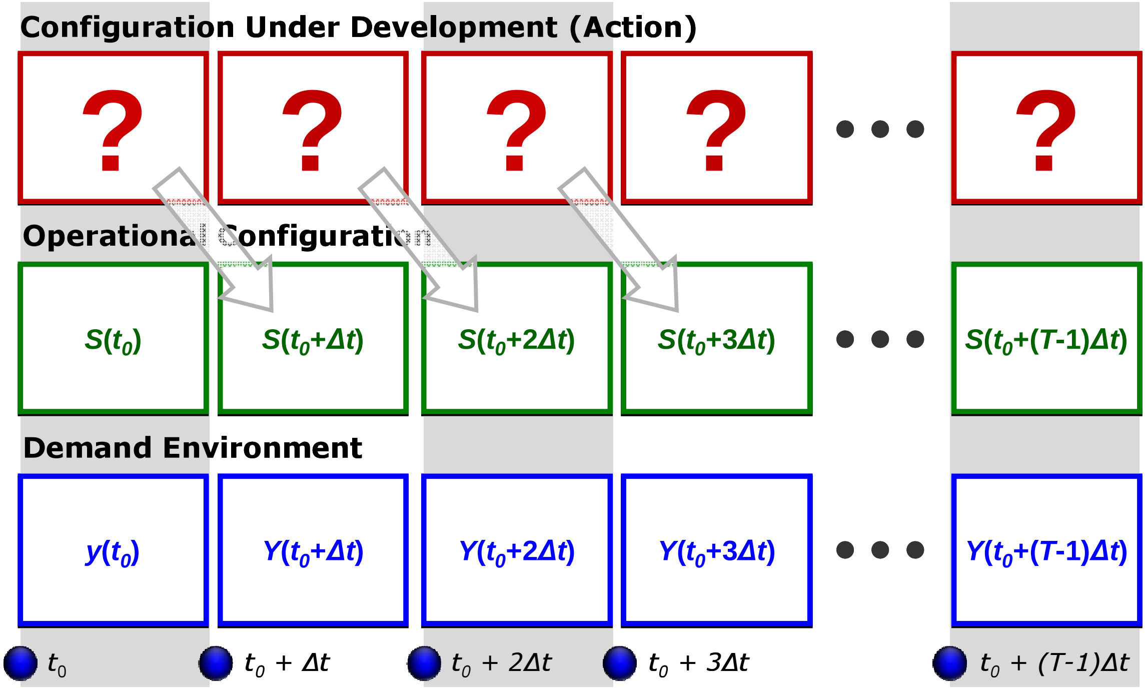

Figure 25. Visualization of Demand Environments, Operational Configurations, and Development Configurations (Actions) over Multiple Time Periods.

Posing the problem in this manner, there exist N^T possible specifications of $a(\tau)$, or possible paths. This is a consequence of the fact that at each of the T periods there exist N configurations that can be selected (recalling that N is the cardinality of $\{S_i\}$). Since the configurations on these paths are specified by the time on the clock at which they are chosen, this type of specification is referred to as an open-loop path.

The number of possible paths N^T may be quite large, depending on the application-specific values of N and T. However, if it is computationally tractable to do so, a Monte Carlo computer simulation may be set up to track the stochastic evolution of cost and performance for each possible path. For each path, a large number of simulations (e.g., several hundred or thousand, depending on the parameters of interest and the confidence desired) is repeated using randomly-generated numbers where required for stochastic propagation of the Markov chain (i.e., according to Eq. (17)) . At each time step in each simulation, the following events and computations occur:

1. Mission demand evolves stochastically according to the Markov chain estimate of Table 10.
2. The operator of the currently operational configuration attempts to use this system to fulfill the new mission demand, earning credit according to the performance matrix.
3. The decision-maker chooses which configuration to develop in the current time period and field in the next time period, incurring a cost according to the cost transition matrix.

For the example satellite application carried through this chapter, $N = 9$ and T will be set to 4 (i.e., an assumed time horizon of 32 years), translating into 6,561 possible paths. The illustrative results that follow assume an initial condition at $t_0 = 0$ in which the operational configuration is Config. 4 (one reconnaissance and one communications satellite) and the demand is characterized by a quiescent conflict environment and full commercial capacity (i.e., Environment 6).

A sample set of Monte Carlo simulation results is shown in Figure 26. This figure shows the result of adopting a configuration path in which a transition is initially made to Config. 3 (the one-communication-satellite configuration) in order to reduce operational costs associated with carrying an unnecessary reconnaissance satellite capability given the relatively high probability that Environment 6 materializes again. Configuration 3 is maintained in the following period, after which the reconnaissance satellite capability is redeveloped and the communications satellite capability is dropped, resulting in Config. 2. Due to the simulation setup, a configuration decision must be made in the final operational time period; since the cost of developing this final configuration will be incurred but no reward will be earned, Config. 1 (the "Nothing" configuration) is selected. As the bottom left portion of Figure 26 shows, this particular path (denoted as [3 3 2 1], by the configuration decisions made at each step) is subject to a stochastically changing demand environment. The size of each yellow dot indicates the likelihood of demand being in a particular state (on the *y*-axis) at a given time (on the *x*-axis); note that all simulations begin in Environment 6 at $t = 0$ years, as specified by the initial condition. The right-hand portion of Figure 26 indicates how per-period cost and performance vary over time. Note that the per-period cost varies between \$233 and \$633 million, and number of demanded satellites available increases from zero to a mean of 0.77 in the final period. The total expected cost for this path over the time horizon is \$1.65 billion, and the total expected number of demanded satellites available is 1.57.

As a theoretical note, it may be observed that the cumulative expected-value results for a given path can be easily computed, without Monte Carlo simulation, as a consequence of the special structure and assumptions of Eqs. (17)-(18). First, note that once a path is chosen, cost in each period (and thus total cost) is fixed, and there is no variability due to future demand environment evolution. Thus, total cost is determined by adding the cost from the total cost transition matrix associated with each pre-specified configuration-to-configuration transition. This is expressed in Eq. (19); note that, within the summation from $\tau = 2$ to $\tau = T$, the row index $S(\tau)$ is substituted with $a(\tau\text{-}1)$ due to the first equation of Eq. (17).

$$E\left(\sum_{\tau=1}^{T} c_{S(\tau),a(\tau)}\right) = c_{S(t_0),a(\tau=1)} + \sum_{\tau=2}^{T} c_{a(\tau-1),a(\tau)} \tag{19}$$

Second, the cumulative performance expectation can be computed analytically as detailed in Eq. (20). In the first step of this short derivation, the expected-value operator is swapped with the period-by-period summation (since, in general, $E(X+Y) = E(X) + E(Y)$). In the second step, the expected performance is expressed as the summation of the environment-conditional performance over all demand environments multiplied by their probabilities of occurrence at time indices τ. Substituting $a(\tau\text{-}1)$ for $S(\tau)$ where appropriate due to the first equation of Eq. (17) yields the final line of Eq. (20). Also note that this line includes a substitution for $P(Y(\tau) = y)$ based on the Chapman-Kolmogorov equations [86] for a Markov chain, in which $P^{(\tau\text{-}1)}_{y(t0),\, y}$ refers to the element in row $y(t_0)$ and column y of the transition matrix P that has been raised to the $(\tau\text{-}1)$ power. Note that this simple Chapman-Kolmogorov substitution is valid because the evolution of the demand environment does not depend upon the configuration path.

$$\begin{aligned} E\left(\sum_{\tau=1}^{T} u_{S(\tau),Y(\tau)}\right) &= \sum_{\tau=1}^{T} E\left(u_{S(\tau),Y(\tau)}\right) \\ &= \sum_{\tau=1}^{T} \sum_{y\in\{Y_i\}} u_{S(\tau),y} \cdot P\left(Y(\tau) = y\right) \\ &= u_{S(t_0),y(t_0)} + \sum_{\tau=2}^{T} \sum_{y\in\{Y_i\}} u_{a(\tau-1),y} \cdot P^{(\tau-1)}_{y(t_0),y} \end{aligned} \tag{20}$$

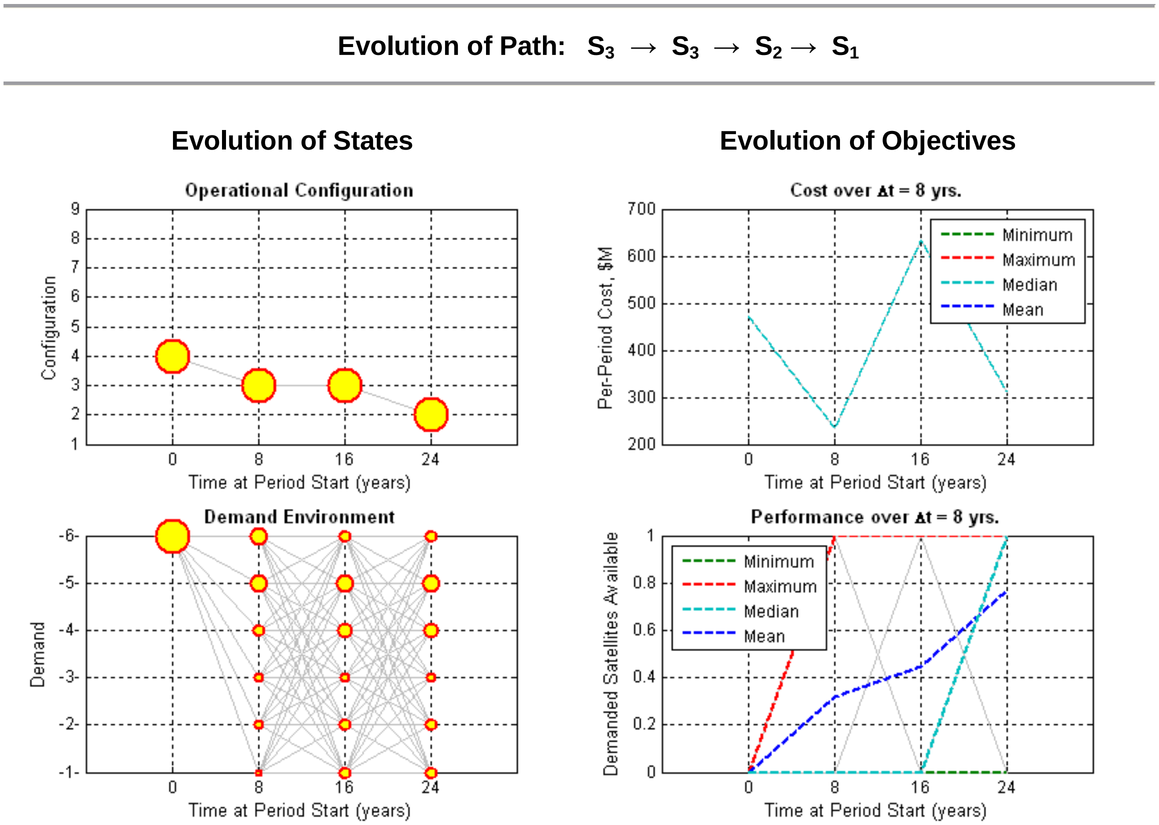

Figure 26. Evolution of configuration path [3 3 2 1]. *In the plots on the left, the size of circles indicates the relative number of Monte Carlo simulation cases that exist in a given configuration or demand environment state (on the y-axes) at a given time (on the x-axes). The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation. Note configuration and cost are deterministic, since a path is specified.*

Obtaining results like those in Figure 26 for each of the 6,561 possible paths in the example satellite application allows the total expected performance to be computed and plotted against total cost for each path as in Figure 27. In this figure, each blue "x" represents the total cost and performance of one path. Notice that, for the population as a whole, there is a general trend that, as more funds are invested, higher performance is expected. However, it is important to recall that the decision-maker has a choice of which path to select. As a result, if he or she cares primarily about total cost and expected total demanded services performed, it would make little sense to select a high-cost, low-performance point toward the lower right of the cluster. Rather, the decision-maker would prefer to choose among the set of nondominated points that comprise the Pareto frontier. This Pareto frontier, shown in red in Figure 27, is composed of the set of possible configuration sequences for which one objective cannot be improved without the sacrifice of another. In this application, the frontier is comprised of just 34 of the 6,561 possible paths and helps to narrow the options considerably.

Listed next to many of the Pareto-optimal points in Figure 27 are associated configuration paths. Note that at the bottom left of the figure is the “do nothing” option in which Config. 1 is fielded for all time periods; this is cost-optimal but also provides the lowest possible performance. At the other extreme is the Pareto-optimal highest-performance option of fielding Config. 9, the two-communications-satellite and two-reconnaissance-satellite option, for all time periods. The Pareto-optimal solutions between these two extremes involve developing Configs. 1, 2, 3, 4, 5, 8, or 9, either immediately or after a delay. Notably absent from the frontier are Configs. 5 and 7, each of which is defined by x_2 = 2 reconnaissance satellites with fewer numbers of communications satellites; the implication of this is that any path that uses these configurations is suboptimal (i.e., is dominated by other paths that can perform at least as well for a cost at least as low).

One additional use of the data in Figure 27 becomes evident when the sample path from Figure 26 is overlaid as the yellow square in Figure 27. Here it can be seen that this path is dominated by solutions on the Pareto frontier. For example, one path, [8 3 3 1], accumulates approximately 26% additional expected performance for a near-identical cost. Another path, [2 3 3 1], accumulates near-identical performance for a 15% lower cost. Thus, the exploration of the possible paths $a(\tau)$ as exemplified by Figure 27 permits candidate paths to be compared and quickly traded against others in terms of relevant figures of merit.

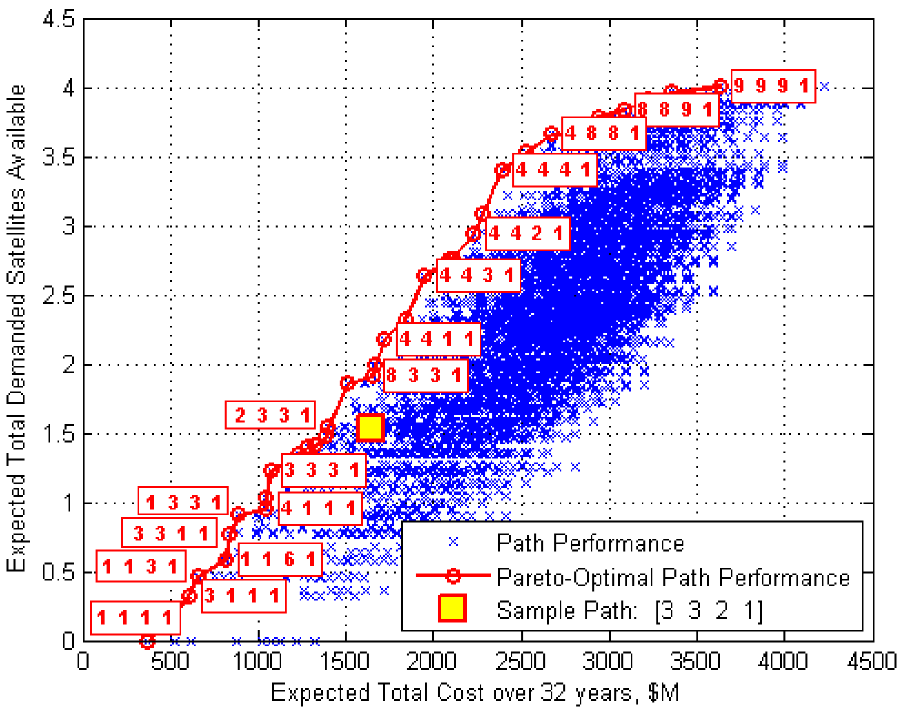

Figure 27. Trade between total demanded services performed and total cost for all open-loop paths. *Selected Pareto-optimal paths are identified by 4-period configuration sequences listed next to red circles.*

In summary, Step 4A of this framework has attempted to pose the multi-period planning problem of Eqs. (17)-(18) in the reasonable manner of asking: What configuration should the decision-maker choose to develop at each time increment? The answer to this question is in general not obvious, particularly since the demand environment evolves stochastically: The decision-maker who wishes to be able to fulfill whatever demand the next period may bring would choose to build the most capable system possible, but this would come at substantial initial expense. The decision-maker who would gamble that tomorrow's demand will be the same as today's would develop few or no new architectural components and in doing so save significant resources; however, this would come with the inability to perform if the next period's demand materializes to require greater capability. Furthermore, whether one period's decision is best (e.g., high-reward or low-cost in the long run) is likely to be dependent on other decisions throughout the system lifetime. In this problem of considering flexibility in design, it is in general necessary to consider all future decisions within a given time horizon in order to judge the appropriateness of any single decision.

This step has illustrated that a straightforward approach to addressing this problem is enumeration of the possible paths $a(\tau)$ over the given time horizon, simulation of these paths, and identification of Pareto-optimal paths in terms of relevant objectives. In the case where the expected values of cumulative cost and performance are of primary interest, this section has shown that analytic computations (see Eqs. (19)-(20)) can be substituted for simulation. The results of this step will be further utilized in the final initial configuration selection step (Step 5).

4.4.2. Find Pareto-Optimal "Closed-Loop" Policies

While straightforward and conceptually similar to an optimization of typical long-term scheduling and roadmapping efforts, the analysis presented in Step 4A has two principal disadvantages. First, for applications with large numbers of configurations N and long time horizons T, it may not be practical to enumerate all N^T possible paths. Second, and conceptually more important, assuming a set path $a(\tau)$ for the entirety of the system's lifetime neglects the ability of the decision-maker to make choices mid-program in response to the evolution of the demand environment.

The latter observation suggests that, for Step 4B, the function a is no longer simply one of time (i.e., $a(\cdot) \neq a(\tau \text{ only})$), but rather also of state, i.e., $a(\cdot) = a(\xi, \tau)$. As this section will show, formulating the action set in this manner permits the state-space framework set forth in Steps 1-3 to be easily integrated and solved within a set of solution techniques for a class of stochastic control processes known as Markov decision processes (MDPs). To begin, however, it is first necessary to define the components and solution procedure for a basic MDP.

4.4.2.1. Markov Decision Processes

To define any MDP, it is necessary to first define (1) a set of states (or state space) Ξ that describes the system of interest for a given time period, (2) a set of decisions or actions Λ available from each state ξ, (3) transition probabilities $p(j|\xi,a)$ given that a particular decision a is made while the system is in state ξ, and (4) expected per-period rewards $\eta(\xi,a)$ associated with actions and/or states [86],[87],[92]. In the case of MDPs on a finite time horizon,* the objective is typically to select the decision policy $a(\xi, \tau)$ that maximizes expected total rewards†. Unfortunately, the number of possible policies $a(\xi, \tau)$ can become much larger than the number of possible paths $a(\tau)$ discussed earlier, and thus enumeration and evaluation of all possible policies is often not practical (to be exemplified later). However, such a problem can frequently be solved by exploiting the computational efficiency of dynamic programming, if the problem exhibits five particular characteristics [87]:

1. The problem can be divided into periods τ with a decision or action a required in each period.
2. Each period τ has a number of system states ξ associated with it. It is desirable for the state to defined such that it contains all the information

---

* There exists some inconsistency in the literature on the definition of MDPs with regard to finite time horizons. Winston [87] adopts the definition that “Infinite horizon probabilistic dynamic programming problems are called Markov decision processes”. Puterman [92], on the other hand, devotes an entire chapter explicitly to finite horizon Markov decision processes. This thesis adopts the latter convention, i.e., that it is the states, decision sets, transition probabilities, and rewards that fundamentally define an MDP and that the time horizon only governs the solution method (e.g., backward induction for the finite horizon problem vs. policy iteration or value iteration for the infinite horizon problem).

† For practical reasons, the decision policies $a(\xi,\tau)$ considered in this thesis are deterministic, i.e., not random variables. More generally, however, MDP formulations exist which can accommodate random policies.

needed to make a decision, since this permits the fourth characteristic to be met.

3. The decision in any period describes how the state in the current period is transformed into the state in the next period.
4. Given the current state, the optimal decision for each of the remaining periods does not depend on previously reached states or previous decisions.[*] In other words, previous decisions and states must not directly influence the optimal path going forward. This is clearly true, for example, when current rewards depend explicitly only on the current decision and/or current system state; and often this can be made true if the system state is properly defined. In many cases such a characteristic is natural. For example, one of the first lessons taught to every economics student is the irrelevance of sunk costs in future planning. As another common example, the shortest route to travel from one city to another has no dependence on how one arrived in the first city.
5. There must exist a recursion that relates the reward earned during periods τ, τ+1, …, T-1 to the reward earned during periods τ+1, τ+2, …, T. In many problems this takes the additive form of $J_\tau = \eta_\tau + J_{\tau+1}$, i.e., that the reward-to-go J at the beginning of period τ is equal to the reward-to-go in the subsequent period plus the reward η earned in period τ itself.

In the case of a multi-period problem exhibiting additive recursion, as more formally shown in Eq. (21), the reward-to-go J from state ξ at period τ is composed of two parts: The first is the reward η earned during time period τ, which is a function of the system's current state ξ and the current action a. The second is the cumulative reward-to-

---

[*] This is a restatement of Bellman's Principle of Optimality: "An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision." [93]

go for all future periods, which is a function of the next-period state $\xi_{\tau+1}$ that is implied by action a (see characteristic 3 above) as well as the policy of actions $\alpha^{\tau+}$ that is adopted for all future states. A discounting factor β can account for effects such as the time value of money or other resources, if an appropriate discount rate is available for use.

$$J_{\tau,\xi}\left(a,\alpha^{\tau+}(\xi,\tau)\right)=\eta_{\tau}(\xi,a)+\beta J_{\tau+1,\xi_{\tau+1}(a)}\left(\alpha^{\tau+}(\xi,\tau)\right) \tag{21}$$

Unfortunately, unless there exist very few states, periods, and possible actions, Bellman's "curse of dimensionality" [93] can make it very difficult to find an optimum specification of the policy α in Eq. (21) via full-factorial analysis or parametric search techniques. However, the form of Eq. (21) permits an optimal policy to be found efficiently via backward induction, a traditional solution procedure for dynamic programming problems. Backward induction begins with the simple problem of optimizing actions in the final period of a multi-period problem, recording the results, and repeating the procedure working backward in time (or other index) until the initial period is reached. For instance, note that in the final period $\tau = T$, by definition no rewards-to-go exist, and the optimal action in this period is specified by the simple problem defined in Eq. (22). In all other periods, the optimization problem can be posed in Eq. (23) as the maximization over all current actions a and policies $\alpha^{\tau+}$ of the expression in Eq. (21). However, if future decisions have no influence on the current reward η, the only role of selecting $\alpha^{\tau+}$ is to maximize the second term in the equation; the maximum possible value of this term at a given period and given state is noted with a hat (^) in Eq. (24). Note that in the next-to-last time period $\tau = T$-1, this second term is already known from the solution to Eq. (22). Similarly, at $\tau = T$-2, the second term is provided by the solution from $\tau = T$-1. This solution process, which continues until the $\tau = 1$ period is reached, is backward induction.

$$\hat{J}_{T,\xi}=\max_{a\in\Lambda(\xi)}\left(\eta_{T}(\xi,a)\right) \tag{22}$$

$$\hat{J}_{\tau,\xi} = \max_{\substack{a \in \Lambda(\xi) \\ \alpha^{\tau+}(\xi,\tau)}} \left(\eta_{\tau}(\xi, a) + \beta J_{\tau+1,\xi_{\tau+1}(a)}\left(\alpha^{\tau+}(\xi,\tau)\right)\right) \tag{23}$$

$$\hat{J}_{\tau,\xi} = \max_{a \in \Lambda(\xi)} \left(\eta_{\tau}(\xi, a) + \beta \hat{J}_{\tau+1,\xi_{\tau+1}(a)} \right) \tag{24}$$

In a case where uncertainties exist in state transitions (i.e., the transformations mentioned in characteristic 3 above), traditional probabilistic dynamic programming operates by considering the expected values (or functions thereof) of reward-to-go treated as a random variable as in Eq. (25). Note that in Eq. (25), the discount factor β is assumed to be deterministic, as is the current-period reward η (although, without loss of generality, η can also be treated as the expected current-period reward).

$$E\left(J_{\tau,\xi}\left(a, \alpha^{\tau+}(\xi,\tau)\right)\right) = \bar{J}_{\tau,\xi}\left(a, \alpha^{\tau+}(\xi,\tau)\right) = \eta_{\tau}(\xi, a) + \beta E\left(J_{\tau+1,\xi_{\tau+1}(a)}\left(\alpha^{\tau+}(\xi,\tau)\right)\right) \tag{25}$$

An analog of the maximization of actions described by Eq. (23) is easily derived as the top line of Eq. (26). However, in the probabilistic problem, note that the next-period state $\xi_{\tau+1(a)}$ is itself a random variable. The expected reward-to-go from period τ +1 can thus be expressed as in the second line of Eq. (26) via the conditioning formula $E(X) = E(E(X|Y)) = \Sigma\, E(X|Y = y)\, P(Y = y)$. The third line mirrors the final step from Eqs. (23) to (24), noting that the optimal policy to adopt for future periods (i.e., the specification of which action to take as a function of future state and time) does not depend on the current action a. The result is again formulation that can easily adopt a backward induction solution procedure, starting at period $\tau = T$.

$$\begin{aligned}
\hat{\bar{J}}_{\tau,\xi} &= \max_{\substack{a\in\Lambda(\xi)\\ \alpha^{\tau+}(\xi,\tau)}} \left(\eta_\tau(\xi,a)+\beta\bar{J}_{\tau+1,\xi_{\tau+1}(a)}\left(\alpha^{\tau+}(\xi,\tau)\right)\right) \\
&= \max_{\substack{a\in\Lambda(\xi)\\ \alpha^{\tau+}(\xi,\tau)}} \left(\eta_\tau(\xi,a)+\beta\sum_{j=1}^{|\Xi|} P(\xi_{\tau+1}=j\,|\,\xi,a)\bar{J}_{\tau+1,j}\left(\alpha^{\tau+}(\xi,\tau)\right)\right) \\
&= \max_{a\in\Lambda(\xi)} \left(\eta_\tau(\xi,a)+\beta\sum_{j=1}^{|\Xi|} P(\xi_{\tau+1}=j\,|\,\xi,a)\hat{\bar{J}}_{\tau+1,j}\right)
\end{aligned} \tag{26}$$

As an example of a Markov decision process on a finite time horizon, consider a system consisting of three states (Ξ = {1, 2, 3}), each of which has two actions available (Λ = {1, 2}), and that rewards and probabilities of transition are as given in Table 12 and Table 13. This system, depicted in Figure 28, simulates a revenue-generating machine with three states (excellent, average, and poor) and a decision-maker with options to operate the machine as usual (action 1) or repair it (action 2) at the beginning of each period. If the machine is repaired from any state, it is instantaneously brought to excellent condition and \$100 of revenue is generated for the period, partially offsetting a \$200 repair cost for a net period reward of -\$100. A four-period time horizon (T = 4) is assumed, and executing the backward induction dynamic programming procedure results in the computations shown in Table 14. Note that the last-period optimization is trivial; since gains from repair are only realized in the long term, it is not optimal to repair the system at the end of the time horizon. As τ is incremented backward, the results of previously computed optimal actions and rewards-to-go are recorded in Table 15. Note that this table indicates that the optimal action (e.g., to repair or not to repair) is a function not only of state but also of time; while it is not optimal to repair a machine in poor condition (ξ = 3) at τ = 3 or τ = 4, it is optimal to do so at τ = 1 and τ = 2. It also indicates, for example, that the expected reward-to-go is a function of the system's initial state; starting in excellent condition (ξ = 1) at τ = 1 entails a \$295.76 cumulative expected reward if the optimal policy is followed, whereas starting in poor condition (ξ = 3) at τ =

1 entails only a \$95.76 cumulative expected reward when the same optimal policy is followed.

Before leaving this example, it is worthwhile to note the computational advantage of dynamic programming. This small problem was solved by hand exactly as reproduced in Table 14, with 12 maximizations and 90 operations (addition or multiplication). In contrast, a full factorial exploration of all possible policies would have required enumerating and evaluating all 2^{12} = 4096 ways of filling out the right half of Table 15.

Table 12. Transition Probabilities for Notional MDP Example.

		Action 1					Action 2		
		To State					To State		
		1	2	3			1	2	3
From State	1	0.7	0.2	0.1	From State	1	0.7	0.2	0.1
	2	0.0	0.5	0.5		2	0.7	0.2	0.1
	3	0.0	0.0	1.0		3	0.7	0.2	0.1

Table 13. Current-Period Reward Function for Notional MDP Example.

		Action	
		1	2
State	1	\$100	-\$100
	2	\$50	-\$100
	3	\$10	-\$100

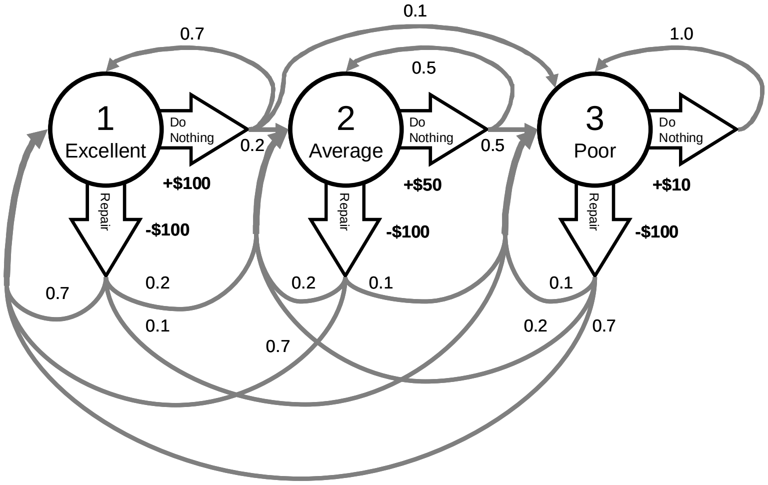

Figure 28. Depiction of the state space, available actions, and action-dependent rewards and transition probabilities for the notional revenue-generating machine MDP example.

Table 14. Backward Induction Calculations for Notional MDP Example.

Period	State	Action	Maximization Argument	Optimal Action?
$\tau = 4$				
	$\xi = 1$			
		$a = 1$	**$100.00**	←
		$a = 2$	**-$100.00**	
	$\xi = 2$			
		$a = 1$	**$50.00**	←
		$a = 2$	**-$100.00**	
	$\xi = 3$			
		$a = 1$	**$10.00**	←
		$a = 2$	**-$100.00**	
$\tau = 3$				
	$\xi = 1$			
		$a = 1$	$100 + 0.7 × $100 + 0.2 × $50 + 0.1 × $10 = **$181.00**	←
		$a = 2$	-$100 + 0.7 × $100 + 0.2 × $50 + 0.1 × $10 = **-$19.00**	
	$\xi = 2$			
		$a = 1$	$50 + 0.5 × $50 + 0.5 × $10 = **$80.00**	←
		$a = 2$	-$100 + 0.7 × $100 + 0.2 × $50 + 0.1 × $10 = **-$19.00**	
	$\xi = 3$			
		$a = 1$	$10 + 1.0 × $10 = **$20.00**	←
		$a = 2$	-$100 + 0.7 × $100 + 0.2 × $50 + 0.1 × $10 = **-$19.00**	
$\tau = 2$				
	$\xi = 1$			
		$a = 1$	$100 + 0.7 × $181 + 0.2 × $80 + 0.1 × $20 = **$244.70**	←
		$a = 2$	-$100 + 0.7 × $181 + 0.2 × $80 + 0.1 × $20 = **$44.70**	
	$\xi = 2$			
		$a = 1$	$50 + 0.5 × $80 + 0.5 × $20 = **$100.00**	←
		$a = 2$	-$100 + 0.7 × $181 + 0.2 × $80 + 0.1 × $20 = **$44.70**	
	$\xi = 3$			
		$a = 1$	$10 + 1.0 × $20 = **$30.00**	
		$a = 2$	-$100 + 0.7 × $181 + 0.2 × $80 + 0.1 × $20 = **$44.70**	←
$\tau = 1$				
	$\xi = 1$			
		$a = 1$	$100 + 0.7 × $244.7 + 0.2 × $100 + 0.1 × $44.7 = **$295.76**	←
		$a = 2$	-$100 + 0.7 × $244.7 + 0.2 × $100 + 0.1 × $44.7 = **$95.76**	
	$\xi = 2$			
		$a = 1$	$50 + 0.5 × $100 + 0.5 × $44.7 = **$122.35**	←
		$a = 2$	-$100 + 0.7 × $244.7 + 0.2 × $100 + 0.1 × $44.7 = **$95.76**	
	$\xi = 3$			
		$a = 1$	$10 + 1.0 × $44.7 = **$54.70**	
		$a = 2$	-$100 + 0.7 × $244.7 + 0.2 × $100 + 0.1 × $44.7 = **$95.76**	←

Table 15. Optimal Actions and Expected Rewards-to-Go for Notional MDP Example.

		Optimal Expected Reward-to-Go, dollars						Optimal Action			
		Time Period, τ						Time Period, τ			
		1	2	3	4			1	2	3	4
State, ξ	1	295.76	244.70	181.00	100.00	State, ξ	1	1	1	1	1
	2	122.35	100.00	80.00	50.00		2	1	1	1	1
	3	95.76	44.70	20.00	10.00		3	2	2	1	1

4.4.2.2. Unification of Flexibility and MDP Frameworks

It may be evident after the above description of a Markov decision process that Steps 1-3 of the flexibility framework proposed in this thesis share many characteristics with an MDP problem. As summarized in Figure 29, Steps 1 and 2 established relevant state spaces, Steps 1 and 3 established rewards (and costs), Step 2 established transition probabilities, and Step 1 established that configuration development options exist for the decision-maker at each point in time. Each of these components – states, rewards, transition probabilities, and possible decisions – is required to define an MDP. However, two slight adjustments must be made to frame the flexibility problem such that MDP solution techniques can be applied directly.

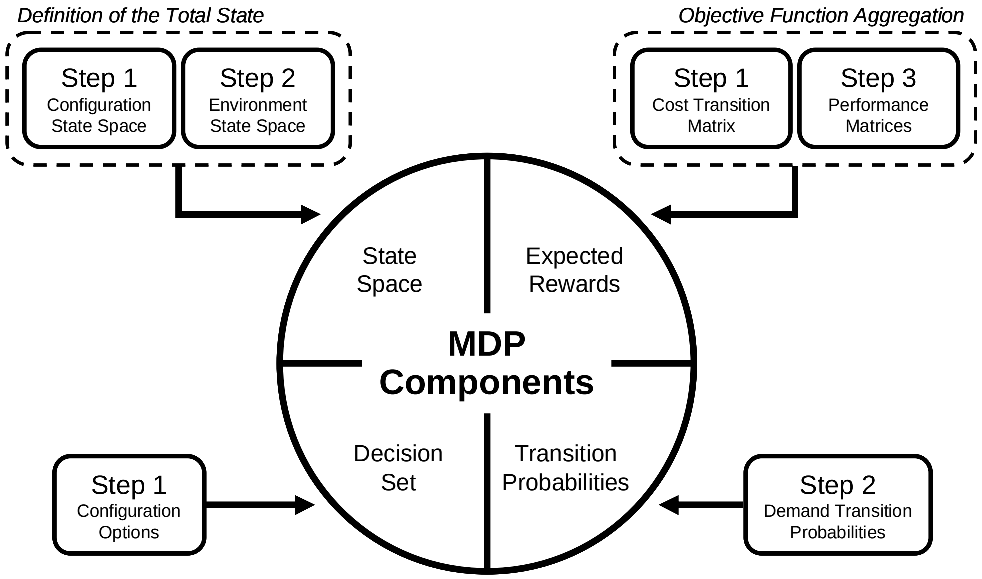

Figure 29. Mapping of Flexibility Framework Components into a Markov Decision Process.

4.4.2.2.1. Definition of the Total State

First, the framework of Steps 1-3 has introduced two *separate* state spaces. Step 1 introduced a configuration state space, and Step 2 introduced a demand environment state space. To utilize an MDP formulation, the problem must be represented in a single state space. It is proposed that a total state be defined as the combination of the configuration and demand states (Total State = {Configuration State, Demand State}, or $\xi = \{S, y\}$). The total state space may be illustrated graphically as in Figure 30 as a three-dimensional "spindle" of total states, in which each vertical layer represents a particular demand environment and each column represents a particular configuration. Thus, it is possible for the fielded system to be in any configuration and operating in any demand environment at any particular point in time. Since configuration is under the control of the decision-maker, he or she can choose to move to any vertical column of the spindle at any point in time (recognizing it takes one time step to make this move).

However, the next-period demand environment is not under the control of the decision-maker. Illustrated in Figure 30 is an instance where Config. 2 is operating in Demand Environment 1. If the decision-maker chooses to develop Config. 7 for the next time period, he or she is assured to move to the column corresponding to Config. 7[*]; however, since the demand environment evolution is stochastic, the layer to which he or she moves is uncertain and depends on the evolution of the Markov chain specified by Step 2. The probability of evolution to each next-period state ξ described in words above is described mathematically by Eq. (27), although it is worth pointing out that the right-hand side of this equation can easily be modified to reflect more complex transition models. Note that by convention for computer programming purposes, $\xi = 1$ is assigned

---

[*] The assumption implicit in this assurance is that the decision-maker will not by accident develop a configuration other than Config. 7, which is generally reasonable. However, if this assumption is not reasonable and the distribution of probabilities that other configurations will be developed is known, this information may be easily incorporated.

to refer to {S = 1, y = 1}, ξ = 2 refers to {S = 2, y = 1}, ξ = N+1 refers to {S = 1, y = 2}, and so on through ξ = N×K referring to {S = N, y = K}. Once the next-period demand environment materializes, the decision-maker finds himself or herself at one particular total state and makes another decision about which of the N configurations to select for the following period.

$$P\left(\xi_{\tau+1} = j \mid \xi_\tau = i, a = k\right) = \begin{cases} P_{\lceil i/N \rceil \lceil j/N \rceil}, & \mathrm{mod}(j,N) = k \\ 0, & \text{otherwise} \end{cases} \tag{27}$$

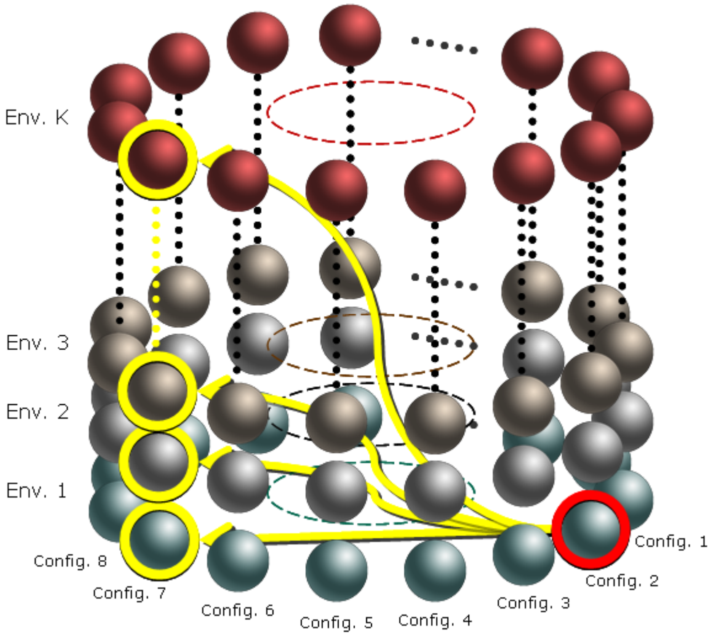

Figure 30. "Spindle" of $N \times K$ Total States. *Each layer corresponds to one demand environment and each vertical column corresponds to one configuration. Arrows illustrate that, due to demand environment uncertainty, multiple possible total states are possible in the next period if a decision is made to transition from one configuration to another (e.g., Config. 2 to Config. 7).*

4.4.2.2.2. Objective Function Aggregation

Second, in order to apply the MDP dynamic programming solution technique implied by Eq. (26), the multi-objective problem illustrated in Step 4 must be carefully converted to a single-objective problem. To do this, the present framework proposes to use the interpretation of the Pareto frontier as the set of optima for a weighted aggregate objective function over all possible weights. Thus, it is proposed that the Pareto frontier be found by forming an aggregate weighted objective function, solving the MDP problem as usual using this single objective, and repeating the process for a wide range of weights. While a simple additive weighting function is an appealing aggregate function, it suffers from an inability to detect concave segments of Pareto frontiers. To partially overcome this limitation, a heuristic technique (detailed in Appendix A) using the variable-power per-period aggregate objective function in Eq. (28) is used. In this equation, Ω is the number of per-period objectives, w_i is the weight on the i^{th} objective, T is the total number of time periods in the time horizon, γ_i is per-period performance of the system in terms of the i^{th} objective (normalized such that the sum of γ_i over all time periods cannot exceed unity or become negative, and such that higher values of γ_i are preferred), and n is the objective function power. The conversions used in this application to normalize element c_{ij} of the cost matrix into $\gamma_{1,ij}$ and element u_{ij} of the performance matrix into $\gamma_{2,ij}$ are shown in Eq. (29).

$$\eta(\xi,a) = -\sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T} - \gamma_i(\xi,a)\right)^n \tag{28}$$

$$\begin{gathered}
\gamma_{1,ij} = \frac{-c_{ij} - \min\limits_{1\le i,j\le N}\left(-c_{ij}\right)}{\vartheta\left(\max\limits_{1\le i,j\le N}\left(-c_{ij}\right) - \min\limits_{1\le i,j\le N}\left(-c_{ij}\right)\right)} \\
\gamma_{2,ij} = \frac{u_{ij} - \min\limits_{\substack{1\le i\le N\\ 1\le j\le K}}\left(u_{ij}\right)}{\vartheta\left(\max\limits_{\substack{1\le i\le N\\ 1\le j\le K}}\left(u_{ij}\right) - \min\limits_{\substack{1\le i\le N\\ 1\le j\le K}}\left(u_{ij}\right)\right)} \\
\vartheta = \begin{cases} T & ,\beta = 1 \\ \frac{1-\beta^T}{1-\beta} & ,\beta \ne 1 \end{cases}
\end{gathered} \tag{29}$$

4.4.2.3. Sample Results for Satellite Application

Applied to the satellite example, this unification of flexibility and MDP frameworks can be applied to find optimal policies efficiently for a range of decision-maker cost or performance preferences. These policy solutions take the form of a matrix with $N \times K$ rows and T columns, where each element (ξ,τ) indicates which of N possible actions or decisions should be made given the system is in state ξ at time τ. In the satellite example application, each policy matrix has dimensions 54 (total states) × 4 (time periods), and 9 options exist for each element of the matrix. If a full-factorial analysis of all possible policies were to be conducted (as was done for the simple case of paths in Step 4A), $9^{216} = 1.31\times10^{206}$ simulations would need to be executed! However, use of the structure of the problem as posed by Eq. (26) and scanning over weights and powers as suggested in Eq. (28) permits optimal policy solutions to be found within minutes on a standard desktop computer.

Expected cost and performance results for policy solutions to the satellite example are shown by each blue “x” in Figure 31. Among these, the nondominated (Pareto-optimal) solutions are highlighted and connected in red. Note that the minimum-cost and maximum-performance endpoints of the Pareto frontier are identical to those of the open-loop full factorial analysis of Figure 27, and the shape of the frontier largely mirrors that of Figure 27. However, several of the solutions on the frontier (particularly those on the convex portions of the frontier) outperform any that were possible via an open-loop policy, the reason for which is clear in viewing the example optimal policy solution in Figure 32. Note here that the evolution of the configuration state no longer follows a deterministic path through time but rather changes to respond to the changing demand environment, per the optimal policy specified in Table 16 that provides an expected total performance of 0.57 satellites available for a low expected total cost of $663 million. For example, Table 16 suggests that the decision-maker develop no satellites for the next period in the event that the current conflict environment is quiescent with full commercial capability (Environment 6) since this environment demands no government capability for satellites and is very likely to continue into the next period. On the other hand, if the current conflict environment is hostile and no commercial capacity is available (Environment 1), the underlying Markov chain reflects a 65% chance that the next-period demand environment will have use for at least one communications satellite, and Config. 3 (the one-communications-satellite configuration) is suggested for development from three configurations within this environment. Rule- or policy-based results such as this are impossible to capture using the fixed configuration paths of Step 4A.

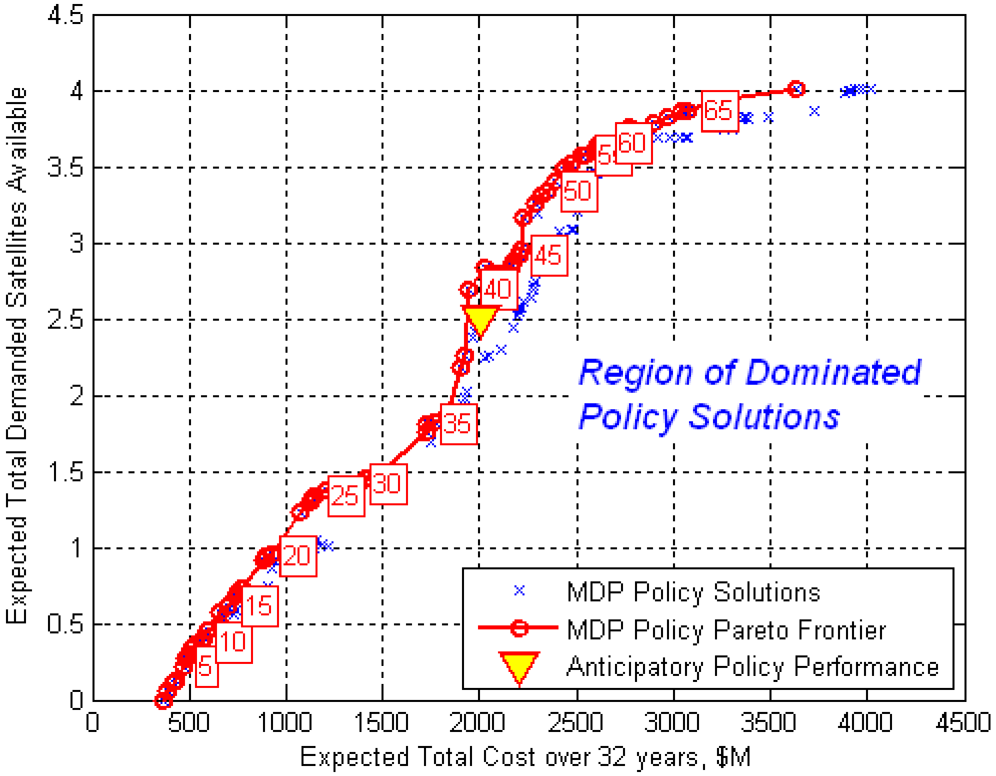

Figure 31. Trade between total demanded satellites available and total cost for MDP policy solutions. Policy identification numbers are indicated next to every fifth policy.

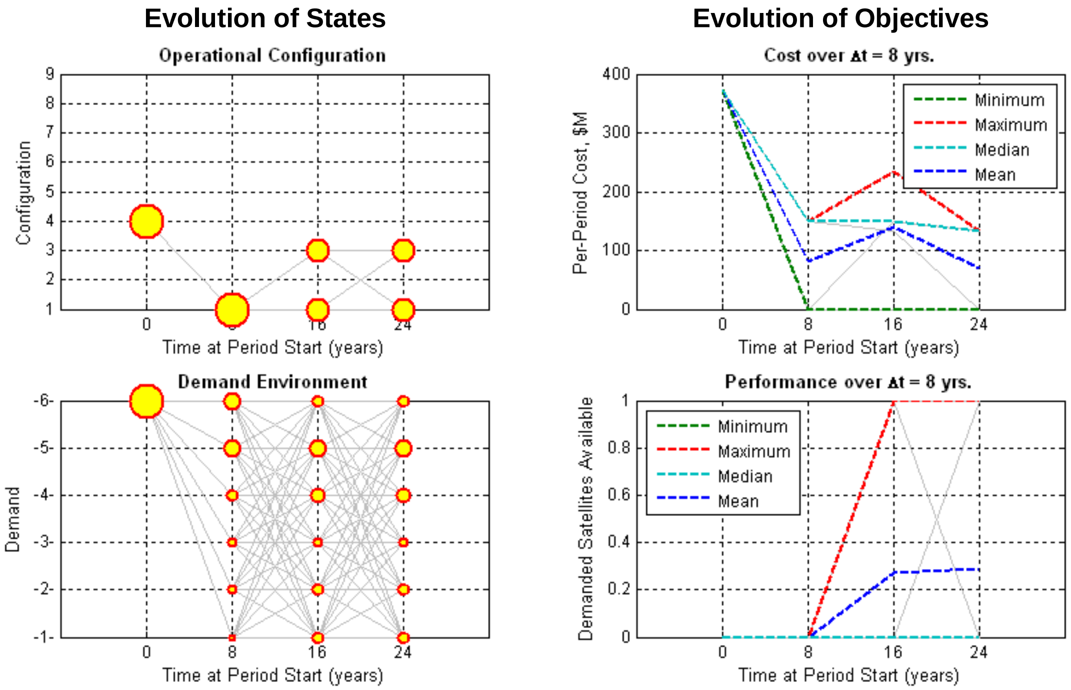

Figure 32. Evolution of states and objectives for Pareto-optimal policy #14 (defined in Table 16). *In the plots on the left, the size of circles indicates the relative number of Monte Carlo simulation cases that exist in a given configuration or demand state (on the y-axes) at a given time (on the x-axes). The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation.*

Table 16. Pareto-Optimal Policy #14 for Satellite Example.

Current State, *s*			**Time at Period Start (years), *t***			
Total State	**Env.**	**Config.**	**0**	**8**	**16**	**24**
1	-1-	1	3	3	3	1
2	-1-	2	1	1	1	1
3	-1-	3	3	3	3	1
4	-1-	4	1	1	1	1
5	-1-	5	1	1	1	1
6	-1-	6	3	1	3	1
7	-1-	7	1	1	1	1
8	-1-	8	1	1	1	1
9	-1-	9	1	1	1	1
10	-2-	1	3	1	3	1
11	-2-	2	1	1	1	1
12	-2-	3	1	1	1	1
13	-2-	4	1	1	1	1
14	-2-	5	1	1	1	1
15	-2-	6	1	1	1	1
16	-2-	7	1	1	1	1
17	-2-	8	1	1	1	1
18	-2-	9	1	1	1	1
19	-3-	1	1	1	1	1
20	-3-	2	1	1	1	1
21	-3-	3	1	1	1	1
22	-3-	4	1	1	1	1
23	-3-	5	1	1	1	1
24	-3-	6	1	1	1	1
25	-3-	7	1	1	1	1
26	-3-	8	1	1	1	1
27	-3-	9	1	1	1	1
28	-4-	1	3	3	3	1
29	-4-	2	1	1	1	1
30	-4-	3	3	3	3	1
31	-4-	4	1	1	1	1
32	-4-	5	1	1	1	1
33	-4-	6	1	1	3	1
34	-4-	7	1	1	1	1
35	-4-	8	1	1	1	1
36	-4-	9	1	1	1	1
37	-5-	1	3	3	3	1
38	-5-	2	1	1	1	1
39	-5-	3	1	1	1	1
40	-5-	4	1	1	1	1
41	-5-	5	1	1	1	1
42	-5-	6	1	1	1	1
43	-5-	7	1	1	1	1
44	-5-	8	1	1	1	1
45	-5-	9	1	1	1	1
46	-6-	1	1	1	1	1
47	-6-	2	1	1	1	1
48	-6-	3	1	1	1	1
49	-6-	4	1	1	1	1
50	-6-	5	1	1	1	1
51	-6-	6	1	1	1	1
52	-6-	7	1	1	1	1
53	-6-	8	1	1	1	1
54	-6-	9	1	1	1	1

Figure 31 also permits comparisons to be made with policies that might be brainstormed or proposed outside of the MDP solution procedure. For example, one reasonable policy that might be proposed is to always develop and field the configuration that least expensively maximizes performance in the most likely next-period demand environment.[*] The policy implied by this statement is provided in Table 17; for instance, if Config. 2 is currently operational in the Environment 1 (i.e., if the system is in total state 2), the most likely next-period demand environment according to Table 10 is Environment 4. To least expensively fulfill the demands of this environment, one reconnaissance and one communications satellite would be developed and launched, which places the system into Config. 4. Thus, as Table 17 shows, Config. 4 is the decision made from total state 2 at all except the final time period.[†]

The performance of this next-period anticipatory policy is summarized by the yellow triangle in Figure 31, and two important points can be noted. First, this policy is dominated by another discovered in the optimization process: Pareto-optimal policy #40 achieves a higher expected performance at a lower expected cost. Second, even if this anticipatory were Pareto-optimal (as it nearly is), note that it is just one of a multitude of policy options; it might be tempting for a decision-maker to adopt this intuitive policy, but Figure 31 illustrates that doing so automatically fixes the long-term cost and performance and ignores a wide variety of options that can reduce cost by a factor of 5 (with a certain trade in performance) or increase performance by a factor of 1.6 (with a certain trade in cost). Thus, the search throughout the policy design space permitted by Step 4B allows the decision-maker to understand the cost and performance trades available and select a policy tailored to his or her preferences.

---

[*] In the event that multiple demand environments have the same probability of materializing next, the environment with the demand for more satellites is used.

[†] The reason for the difference in the final time period decision is the same as discussed earlier in Section 4.4.1.

Table 17. Anticipatory Policy for Satellite Example.

Current State, *s*			Time at Period Start (years), *t*			
Total State	**Env.**	**Config.**	**0**	**8**	**16**	**24**
1	-1-	1	4	4	4	1
2	-1-	2	4	4	4	1
3	-1-	3	4	4	4	1
4	-1-	4	4	4	4	1
5	-1-	5	4	4	4	1
6	-1-	6	4	4	4	1
7	-1-	7	4	4	4	1
8	-1-	8	4	4	4	1
9	-1-	9	4	4	4	1
10	-2-	1	2	2	2	1
11	-2-	2	2	2	2	1
12	-2-	3	2	2	2	1
13	-2-	4	2	2	2	1
14	-2-	5	2	2	2	1
15	-2-	6	2	2	2	1
16	-2-	7	2	2	2	1
17	-2-	8	2	2	2	1
18	-2-	9	2	2	2	1
19	-3-	1	2	2	2	1
20	-3-	2	2	2	2	1
21	-3-	3	2	2	2	1
22	-3-	4	2	2	2	1
23	-3-	5	2	2	2	1
24	-3-	6	2	2	2	1
25	-3-	7	2	2	2	1
26	-3-	8	2	2	2	1
27	-3-	9	2	2	2	1
28	-4-	1	4	4	4	1
29	-4-	2	4	4	4	1
30	-4-	3	4	4	4	1
31	-4-	4	4	4	4	1
32	-4-	5	4	4	4	1
33	-4-	6	4	4	4	1
34	-4-	7	4	4	4	1
35	-4-	8	4	4	4	1
36	-4-	9	4	4	4	1
37	-5-	1	2	2	2	1
38	-5-	2	2	2	2	1
39	-5-	3	2	2	2	1
40	-5-	4	2	2	2	1
41	-5-	5	2	2	2	1
42	-5-	6	2	2	2	1
43	-5-	7	2	2	2	1
44	-5-	8	2	2	2	1
45	-5-	9	2	2	2	1
46	-6-	1	2	2	2	1
47	-6-	2	2	2	2	1
48	-6-	3	2	2	2	1
49	-6-	4	2	2	2	1
50	-6-	5	2	2	2	1
51	-6-	6	2	2	2	1
52	-6-	7	2	2	2	1
53	-6-	8	2	2	2	1
54	-6-	9	2	2	2	1

4.5. Step 5: Implications for Initial System Selection

Recall that a major purpose of this framework is to inform initial system selection. The analysis of Step 4 has produced a large set of data on optimal paths and policies to follow for the *entire* system time horizon, and it is easy to lose track of the implications this has for the *initial* system decision. This final step of the framework builds upon the analysis results of Step 4 to provide implications for this decision. Covered first are implications based on the expected-value Pareto frontiers of Step 4, followed by advanced topics that consider variations on these objectives and on the initial demand environment assumption.

4.5.1. Implications based on the Expected-Value Pareto Frontier

In the case of an open-loop path as discussed in Step 4, the initial decision is simply the first configuration in its associated configuration sequence. In the case of a policy, the initial decision is found by locating the initial condition in the row of the policy matrix and examining the element in the first column. To facilitate this, the initial configurations specified by the Pareto-optimal paths and policies found in Step 4 may be identified explicitly and plotted as a function of decision-maker cost (or performance) preference.

In the case of the satellite example, the initial configurations implied by the Pareto-optimal paths and policies of Figure 27 and Figure 31 are identified in Figure 33. Here, the Pareto frontier solutions of Figure 27 and Figure 31 are identified by their expected total cost on the *x*-axis. On the *y*-axis are the initial configuration decisions called for by each Pareto-optimal path (yellow circles) or policy (blue squares). Two particular observations can be made: First, only six of the nine configurations appear among the optimal initial decisions. All paths and policies with other initial decisions are dominated by paths and policies using these six configurations. Second, the number of satellites involved in the initial configuration tends to increase as the expected total cost of the system increases. For example, the optimal initial configuration tends to progress from no satellites (for a low budget) to one communications satellite (Config. 3) to a communications satellite and a reconnaissance satellite (Config. 4) to eventually two communications satellites and one reconnaissance satellite (Config. 8) and to two of each type of satellite (Config. 9). The primary exception to this occurs in the medium-cost region in which the initial decision to develop the one-reconnaissance-satellite configuration (Config. 2) is associated with Pareto-optimal paths and policies. Thus, for example, a decision-maker interested in minimizing cost without regard for performance would elect to develop Config. 1 or 3 initially, while a decision-maker interested in maximizing performance without regard for cost would opt for Config. 9 initially. A decision-maker seeking a compromise between these extremes should opt for Configs. 2, 4, or 8, but any other selection would result in a suboptimal long-term cost and performance result.

Also noted next to several paths and policies in Figure 33 and explicitly plotted in Figure 34 are the number of transitions Φ available from each initial configuration for the average per-period cost associated with each total cost. As discussed in Step 1, this number Φ can be considered an indicator of flexibility, and it can be seen that more flexible initial configurations (e.g., Φ = 7) are selected at higher cost and performance preferences. Thus, there exists some correlation between flexibility and performance on the Pareto frontier, which paves the way for an important discussion after the following section to conclude this theoretical discussion on the present framework for integrating flexibility into system design decisions.

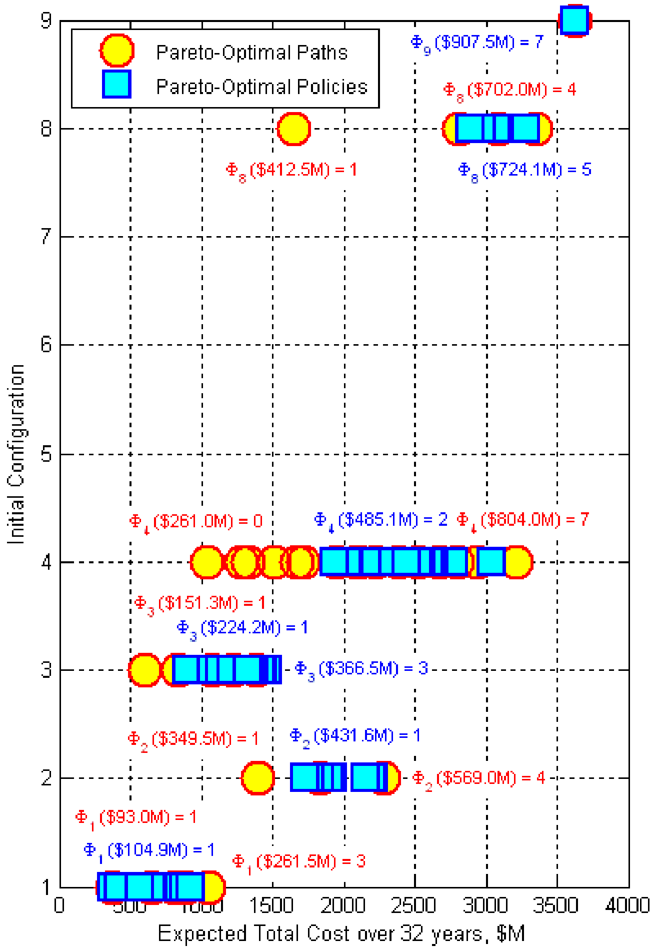

Figure 33. Initial configurations for Pareto-optimal paths and policies as a function of expected path or policy total cost. *Also noted are the numbers of transitions available for several initial configurations at their path or policy's average per-period budget requirements.*

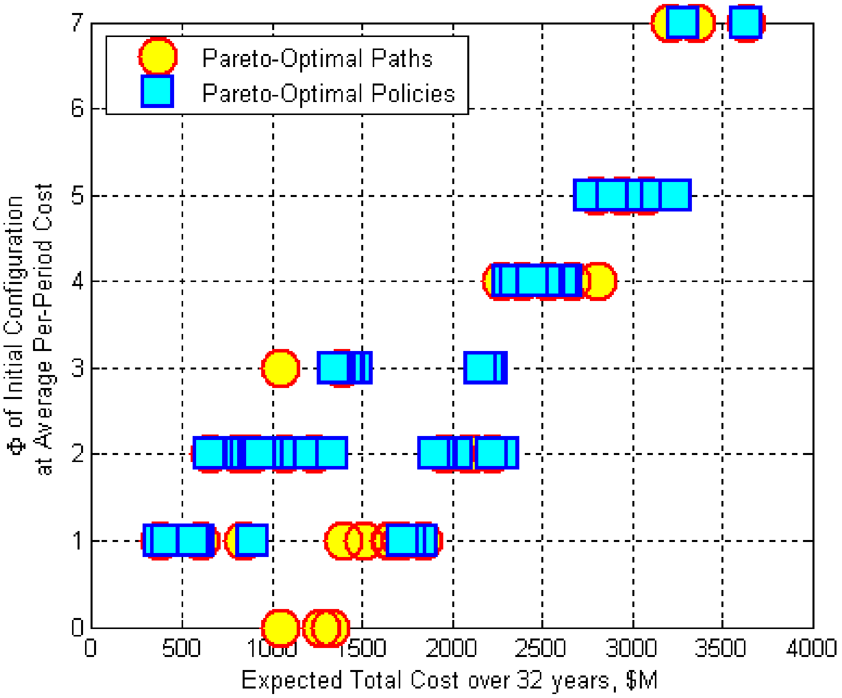

Figure 34. Number of Available Transitions for Pareto-Optimal Initial Configurations.

4.5.2. Accounting for Non-Expected-Value Objectives

An important consideration for initial system selection is the fact that expected-value objective functions of cumulative cost and performance metrics may not fully capture a decision-maker’s true objectives. For example, in the case of one-of-a-kind engineering programs in which the large sample sizes do not exist for which expected-value-based decisions would be most relevant, a decision-maker may have some interest in minimizing risks or deviations away from a central tendency measure of cost or performance.

In the event that the decision-maker’s true objectives are not expected values of cumulative costs and performance metrics, all is not lost in the approach of Steps 1-4. In fact, much is gained. Recall that the policy trade-space (e.g., the 1.31×10^{206} ways that the policy matrices of Table 16 and Table 17 could be populated) can be astronomically large, but that MDP dynamic programming techniques can permit optimal policies to be found quickly and efficiently over a wide range of decision-maker cost and performance preferences. The dynamic programming technique's limitation is that the objective function must take the form of a cumulative expected-value objective, but this can be valuable if the decision-maker is able to identify such an objective as an acceptable surrogate or starting point for an analysis such as one that is illustrated next.

In the following analysis, a customized multi-objective genetic algorithm is employed to perturb each of the policies identified in Figure 31, simulate each new hybrid policy, and search for non-dominated solutions in terms of any combination of metrics that can be accounted for via simulation. This genetic algorithm is customized in the sense that it is real-valued and not binary in order to avoid the need to represent each of the 9^{216} possible policies of the example application via a 685-bit binary string; instead, each member of the population is represented by its full policy matrix. At each iteration, each member has a 10% probability of mutation, which is associated with a random change of approximately 10% of the member's matrix elements. In addition, at each iteration a 75% probability of crossover exists, which occurs via the two-point splicing of the matrix rows of each member with a randomly selected other member to form one new member of the population. Elitism is employed to ensure the highest-performing member of the population is retained from one generation to the next, and the initial guess is also retained in the population throughout. The algorithm employs the infinity norm aggregate objective function (for rationale, see Appendix A) and sweeps across the range of weights to identify optimum solutions over a variety of decision-maker preferences. Note, however, that this particular selection of a genetic algorithm is meant only to illustrate how existing optimization and design space exploration techniques can be used to further explore the policy space introduced and efficiently solved for in Step 4B; a great deal of expansion and exploration of other algorithms is possible and warranted in the future.

The results of Figure 35 are produced by applying this policy exploration technique to the new metrics of 90th percentile (near-worst-case) total cost and 10th percentile (near-worst-case) total number of demanded satellites available, in addition to the expected-value versions of these metrics. Of particular note in the Figure 35 multivariate plot are four subplots: First, the data in the subplot of the second row and first column shows the familiar expected-value cost and performance trade, with slightly better Pareto frontier performance due to the genetic algorithm's search. Second, the data in the subplot of the last row and second column shows the 10th percentile performance vs. the 90th percentile cost; the performance data in this subplot is noticeably more discrete since fractional numbers of available satellites are not possible in a simulation. Finally, the upper left and bottom right subplots show the correlations between the new percentile-based metrics and their expected-value counterparts. In the cases of the cost subplot, linear correlation is particularly high ($R^2 = 0.88$) and strongly supports the use of expected value as a surrogate for optimizing the percentile-based metrics.

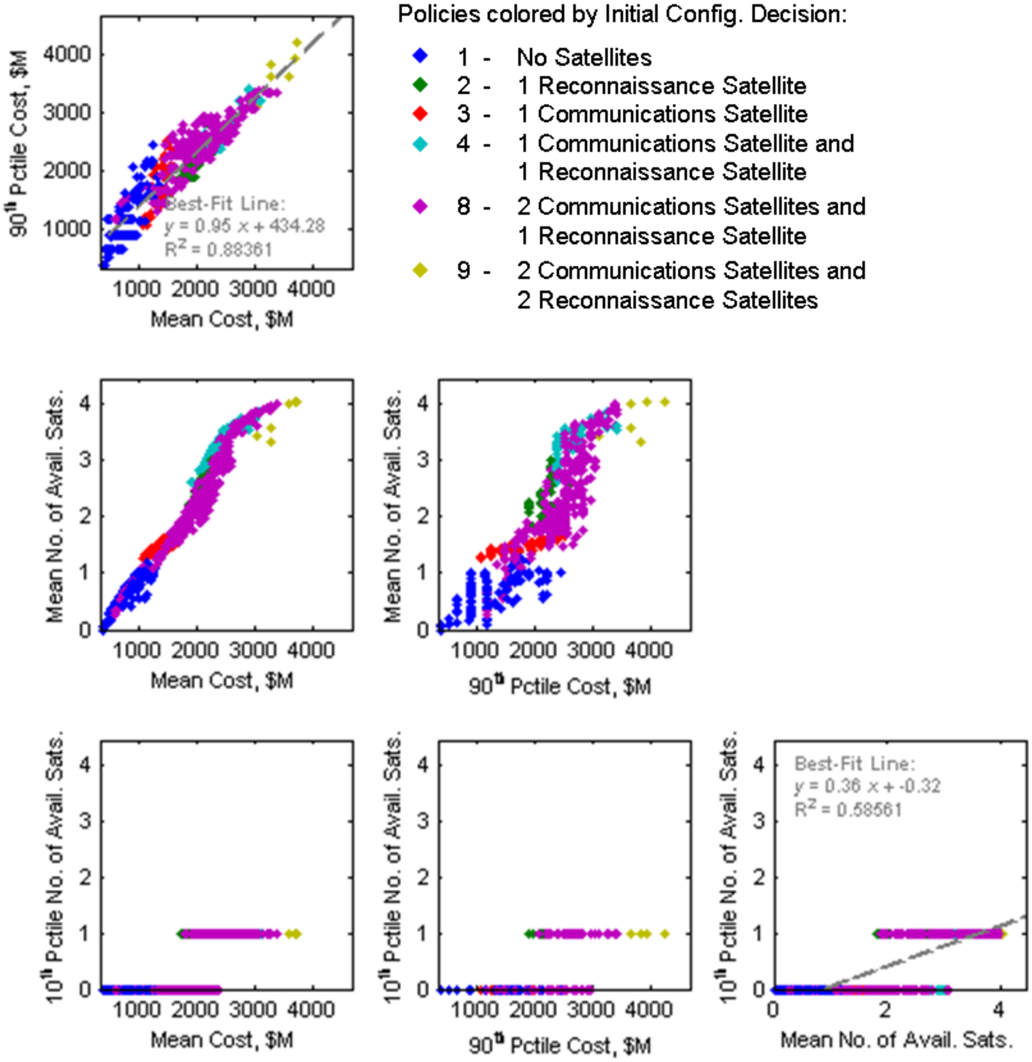

Figure 35. Multivariate plot of multi-objective genetic algorithm policy results. *Each data point indicates the performance of one policy result in terms of the four percentile-based and expected-value metrics of interest. Data points are colored by their corresponding policy's initial configuration decision.*

The usefulness of the multivariate plot of Figure 35 becomes more evident if cost or performance constraints are imposed by the decision-maker. For example, suppose that this decision-maker has a $3 billion limit on the funds available for supporting this system over its time horizon. If the decision-maker wishes to be 90% sure that this budget will not be breached, a $3 billion constraint may be imposed on the 90th percentile total cost metric. This constraint eliminates many high-cost (and also high-performance) options that formerly fell into the high 90th percentile cost regions of the multivariate plot that are now gray in Figure 36. Similarly, the decision-maker may wish to have 90% confidence that at least one service will be performed over the system's lifetime. In this case an additional constraint may be imposed on the last row of subplots in Figure 36. Combined, these two constraints eliminate a large number of the policy options available. Furthermore, they limit the decision-maker's options for which configuration to select initially. In this case, options are limited to just Configs. 2, 4, and 8 (associated with policies colored cyan, magenta, and green). While Figure 36 shows that all these initial configurations are associated with policies of equal 10th percentile performance, Config. 4 tends to be associated with policies of higher long-term cost than Config. 2, and Config. 8 is associated with policies across a range of long-term costs (both mean and 90th percentile). Note that if a decision-maker wished to have more insight into the behavior of any given policy, the policies associated with each data point in Figure 36 could easily be simulated and visualized in a manner identical to the example of Figure 32. In the end, in this notional scenario a recommendation for Configs. 2, 4, or 8 could be justified, depending on the cost and performance preferences of the decision-maker. Importantly, through the use of this thesis' framework, the majority of possible initial configurations have been eliminated, either due to long-term cost or performance constraints, or due to the Pareto sub-optimality of the policies with which they are associated.

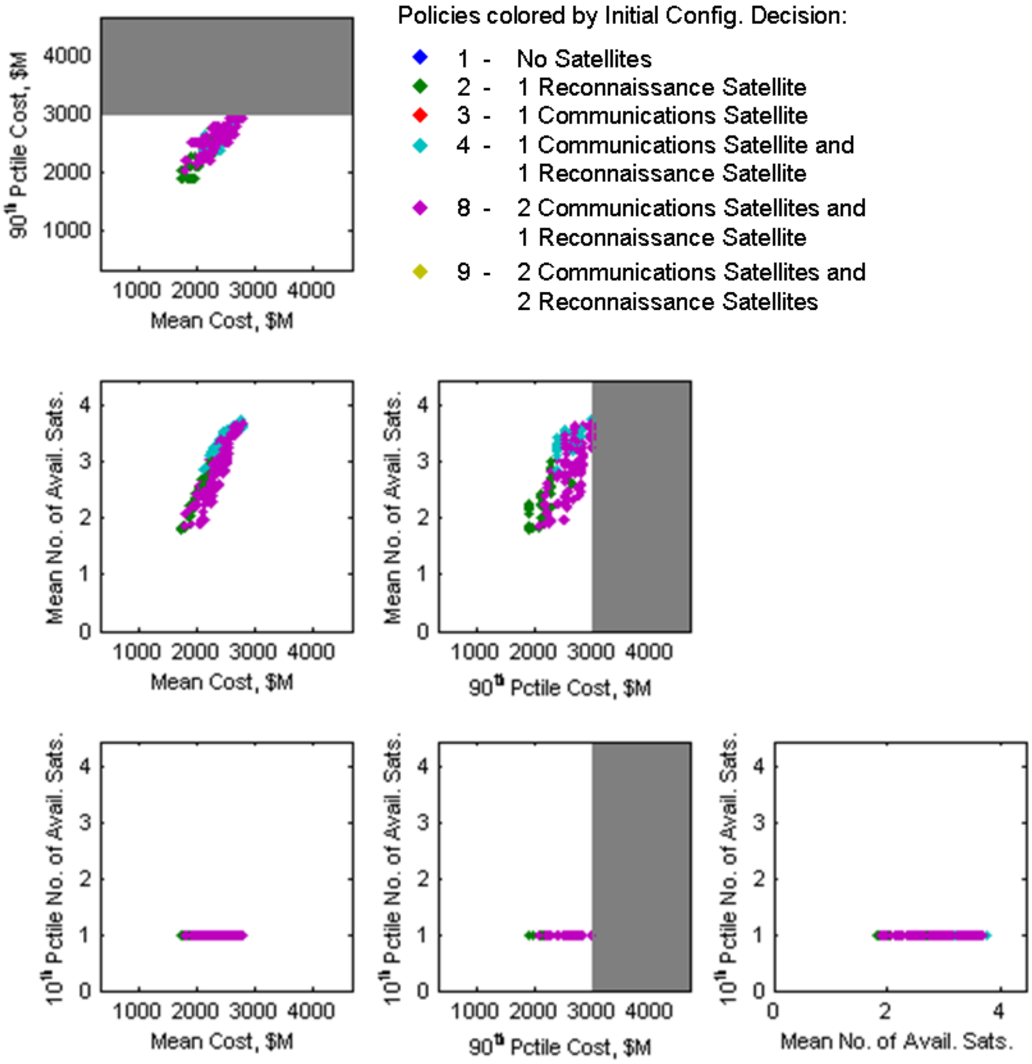

Figure 36. Multivariate plot of multi-objective genetic algorithm policy results with cost and performance constraints imposed. *Each data point indicates the performance of one policy result in terms of the four percentile-based and expected-value metrics of interest. Data points are colored by their corresponding policy's initial configuration decision. Gray areas indicate regions of the space eliminated due to cost and performance constraints.*

4.5.3. Flexibility, Entropy, and Policy

An interesting observation was made earlier regarding the fact that in the satellite program application there existed some correlation between the long-term performance and the flexibility $\Phi(b)$ of the initial configuration decision associated with the Pareto-optimal policies. This is natural for two reasons: First, higher-performing configurations are themselves often correlated with higher cost, and higher-cost configurations are themselves often correlated with lower switching costs (and thus higher numbers of available transitions, Φ). Second, because higher-performing configurations are correlated with higher cost, the average per-period budget *b* used to calculate Φ(*b*) of the initial configurations is higher for these configurations. Thus, this seemingly interesting question becomes somewhat less so in short order; however, a useful and very interesting question does present itself upon reflection of this previous question: How does the optimal behavior with respect to flexibility vary as certainty about the demand environment changes?

To address this empirically, consider the satellite example used throughout this chapter but with two new models for the demand environment. The first, the low-entropy-rate environment, is shown in Table 18 and has an entropy rate of 0.67 bits (for details on entropy rate, see the discussion in Step 2 and Eq. (16)). The second, the medium-entropy-rate environment, is shown in Table 19 and has an entropy rate of 1.26 bits. The reference model used throughout the demonstration in Steps 1-5 is provided in Table 10 and has a relatively high 2.36 bit entropy rate. Note that the lower the entropy rate of a Markov chain model is, the more deterministic it becomes.

Table 18. Low-Entropy-Rate Demand Markov Chain Transition Matrix (H′=0.67 bits).

		To Demand Environment, Y_j					
		- 1 - Hostile, None	- 2 - Hostile, Some	- 3 - Hostile, Full	- 4 - Quiescent, None	- 5 - Quiescent, Some	- 6 - Quiescent, Full
From Demand Environment, Y_i	- 1 - Hostile, None	0.00	0.00	0.00	0.98	0.02	0.00
	- 2 - Hostile, Some	0.07	0.00	0.00	0.00	0.85	0.07
	- 3 - Hostile, Full	0.00	0.00	0.00	0.00	1.00	0.00
	- 4 - Quiescent, None	0.04	0.00	0.00	0.48	0.48	0.00
	- 5 - Quiescent, Some	0.00	0.08	0.00	0.01	0.89	0.01
	- 6 - Quiescent, Full	0.00	0.00	0.00	0.06	0.70	0.23

Table 19. Medium-Entropy-Rate Demand Markov Chain Transition Matrix (H′=1.26 bits).

		To Demand Environment, Y_j					
		- 1 - Hostile, None	- 2 - Hostile, Some	- 3 - Hostile, Full	- 4 - Quiescent, None	- 5 - Quiescent, Some	- 6 - Quiescent, Full
From Demand Environment, Y_i	- 1 - Hostile, None	0.00	0.02	0.00	0.92	0.06	0.00
	- 2 - Hostile, Some	0.14	0.01	0.01	0.01	0.70	0.14
	- 3 - Hostile, Full	0.00	0.02	0.00	0.00	0.97	0.00
	- 4 - Quiescent, None	0.09	0.01	0.00	0.45	0.45	0.00
	- 5 - Quiescent, Some	0.01	0.15	0.01	0.05	0.74	0.05
	- 6 - Quiescent, Full	0.00	0.01	0.01	0.12	0.59	0.28

Executing the MDP optimization procedure of Step 4B for a range of decision-maker cost and performance preferences results in a Pareto frontier and a set of suggested initial configuration decisions for each of the two new demand environment Markov chains. Just as was done for the nominal (high entropy rate) case in Figure 34, for each of these chains the number of available transitions from each optimal policy's initial configuration for the average per-period cost can be plotted. These results are shown in Figure 37. Surprisingly, despite the large differences in entropy rate among the three cases, the three plots look remarkably alike: The number of available transitions spans from one at an expected total cost somewhat under $500 million to seven or eight at a cost between $3.5 and $4.0 billion, with a significant clustering of low-cost, low-flexibility options.

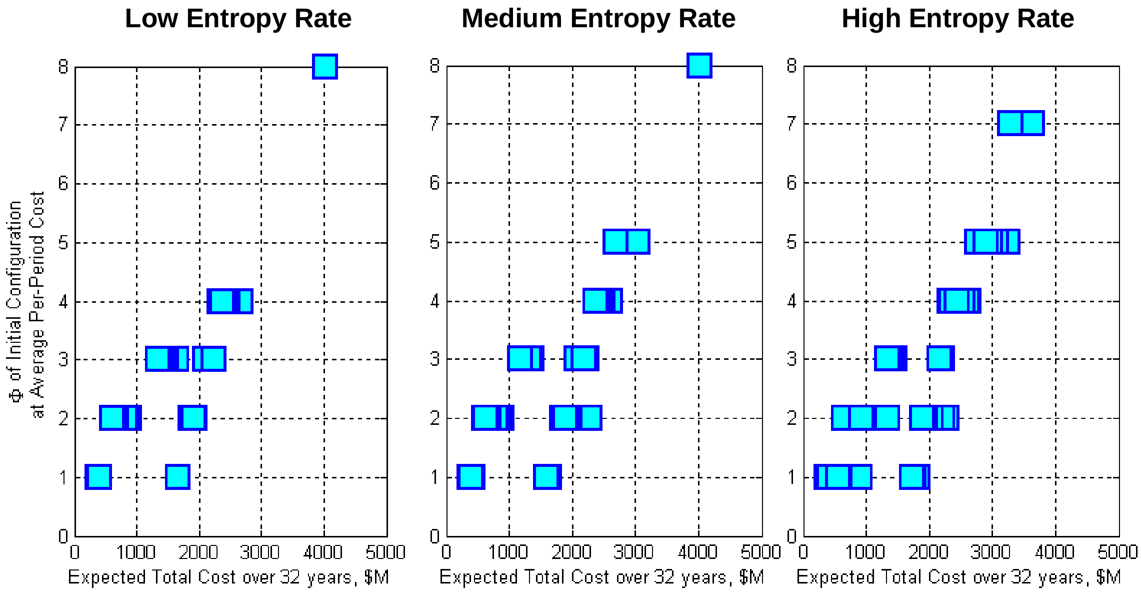

Figure 37. Φ of the Initial Configurations of Pareto-Optimal Policies for varying Demand Environment Entropy Rates.

If flexibility of the optimum initial configuration does not distinguish high- from medium- and low-entropy-rate (i.e., more from less random) demand environments, then what does? As the following discussion will highlight, the simple answer is *policy*. Policy specifies *how* a system's flexibility is exercised, which changes significantly with the stochastic nature of the demand environment. To this end, it is instructive to examine two special cases:

Consider first a case in which demand evolves deterministically according to the transition matrix in Table 20. In this matrix, Environments 4 and 5 are absorbing states, and Environment 5 is the state to which the system in the present simulation is absorbed (since the initial condition for this simulation is Env. 6). Intuitively, one would expect that Pareto-optimal policy solutions in the presence of no uncertainty are paths, and this is exactly the case. Figure 38 shows the Pareto frontier and the time-histories of configurations for the Pareto-optimal points for this deterministic case. Depending on the l performance desired, Config 2 is fielded for a greater or fewer number of time periods. The optimal policy in a deterministic demand environment thus degenerates to a path.

Table 20. Special Deterministic Markov Chain Transition Matrix (H′=0 bits).

		To Demand Environment, Y_j					
		- 1 - Hostile, None	- 2 - Hostile, Some	- 3 - Hostile, Full	- 4 - Quiescent, None	- 5 - Quiescent, Some	- 6 - Quiescent, Full
From Demand Environment, Y_i	- 1 - Hostile, None	0	0	0	1	0	0
	- 2 - Hostile, Some	0	0	0	0	1	0
	- 3 - Hostile, Full	0	0	0	0	1	0
	- 4 - Quiescent, None	0	0	0	1	0	0
	- 5 - Quiescent, Some	0	0	0	0	1	0
	- 6 - Quiescent, Full	0	0	0	0	1	0

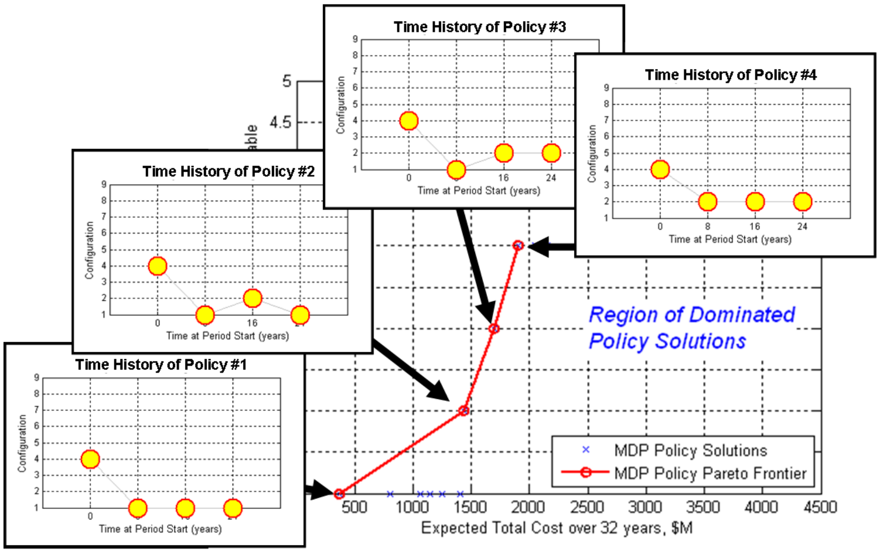

Figure 38. Trade between total demanded satellites available and total cost for MDP policy solutions subject to a deterministic (minimum entropy rate) demand environment Markov chain. *Configuration time histories for the optimal policies are overlaid.*

Consider second the case in which demand evolves uniformly randomly according to the transition matrix in Table 21. In this extreme case, a decision-maker has decided to assume that he or she has no knowledge about the likelihood of any future demand environment. Each environment is equally likely to occur in the future, and no knowledge about the current demand environment will improve knowledge about the likelihood of the next period's environment. In this case, since no information about demand evolution can be gained from the probability transition matrix, it may not be immediately obvious what optimal policies may exist. However, some suboptimal policies may be clear. For example, a policy that guides the decision-maker to alternate between developing Config. 1 and 9 (the most and least capable configurations) would be particularly wasteful, since Config. 9 would need to be re-developed in each time period. This simple example thus illustrates, and Figure 39 confirms, the important point that *Pareto-optimal policies exist even in a uniformly random (maximum entropy rate) demand environment.* This carries with it important implications for the usefulness of this thesis' framework even in the presence of complete uncertainty in the transition probabilities of the demand environment.

More importantly, what is the character of the Pareto-optimal policies in the case of the uniformly random demand environment? In the same format as Figure 38, Figure 39 shows the Pareto frontier as well as the time-histories of configurations for Pareto-optimal points in the maximum-entropy-rate case. Strikingly, most of these points (and in fact all of the points that fall on the convex portion of the frontier, which is known to contain the set of global optimum solutions) are actually paths. Thus, the optimal policy in response to total uncertainty in the demand environment also tends to degenerate to a path. In such an environment, the present contributes no information about the future, and no useful policy exists to specify how to adapt a system to future environment changes.

Table 21. Special Uniform Random Markov Chain Transition Matrix (H'=2.58 bits).

		To Demand Environment, Y_j					
		- 1 - Hostile, None	- 2 - Hostile, Some	- 3 - Hostile, Full	- 4 - Quiescent, None	- 5 - Quiescent, Some	- 6 - Quiescent, Full
From Demand Environment, Y_i	- 1 - Hostile, None	0.1667	0.1667	0.1667	0.1667	0.1667	0.1667
	- 2 - Hostile, Some	0.1667	0.1667	0.1667	0.1667	0.1667	0.1667
	- 3 - Hostile, Full	0.1667	0.1667	0.1667	0.1667	0.1667	0.1667
	- 4 - Quiescent, None	0.1667	0.1667	0.1667	0.1667	0.1667	0.1667
	- 5 - Quiescent, Some	0.1667	0.1667	0.1667	0.1667	0.1667	0.1667
	- 6 - Quiescent, Full	0.1667	0.1667	0.1667	0.1667	0.1667	0.1667

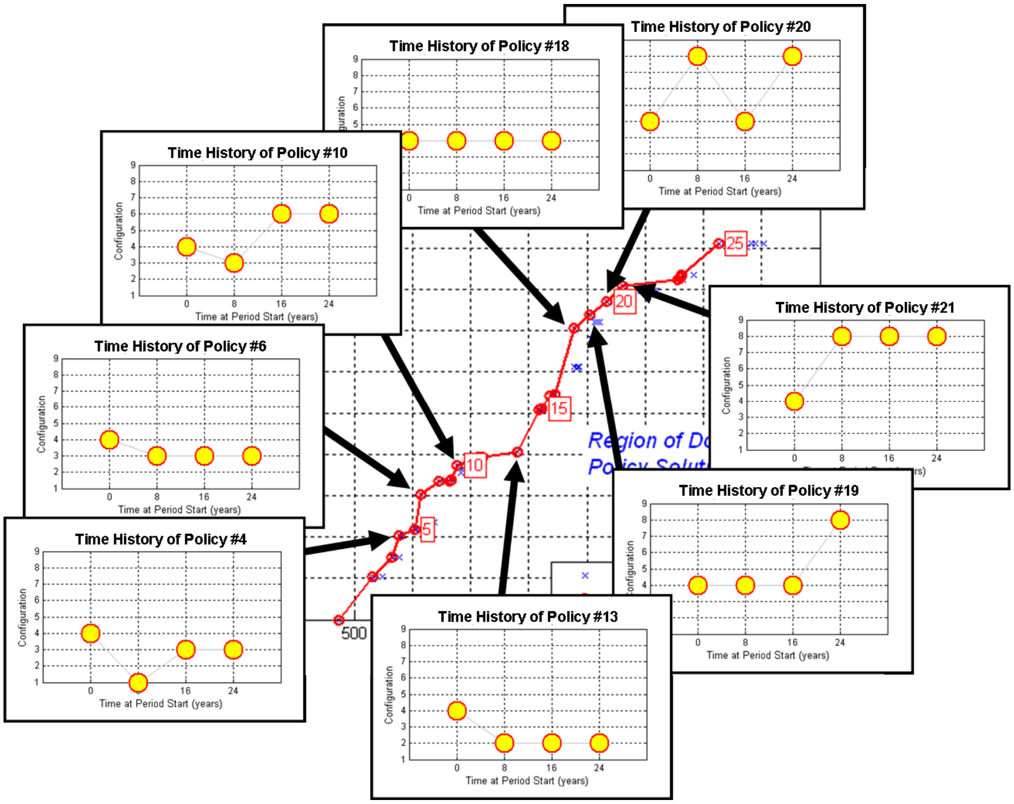

Figure 39. Trade between total demanded satellites available and total cost for MDP policy solutions subject to a uniformly random (maximum entropy rate) demand environment Markov chain. *Configuration time histories for several optimal policies are overlaid.*

In summary, despite the fact that the Pareto-optimal initial configurations for high-, medium-, and low-entropy-rate demand environments possessed similar flexibility Φ, this flexibility was used in very different ways:

- In the *low-entropy-rate* case, this ability to change configurations was used to respond to environment changes known to the decision-maker in advance, making this case akin to a traditional *optimization* problem in the sense that the influence of uncertainties over time need not be accounted for. More precisely, in this case the *path* (not necessarily the configuration) was optimized.

- In the *high-entropy-rate* case, a similar path was set in advance since, although the future was evolving stochastically, minimal or no information about the future could be gleaned from the present environment. This case is thus more akin to the concept of *robustness* in the sense that the selected sequence of configurations is able to perform well in whatever future environments materialize, but that no actions are intended to be taken to respond to these future environments. To be precise, the *path* (as opposed to the configuration) was made robust.

- In the *medium-entropy-rate* case, *flexibility* can be utilized to its fullest extent. In such a case, the Pareto-optimal policies are no longer paths but rather are true policies or "playbooks" that indicate to the decision-maker what action to take (i.e., configuration to develop) given future environment and configuration states. These policies take advantage of information gained about the future evolution of demand from knowledge of the present state.

These observations thus highlight the fundamental fact that although flexibility may exist in many systems, it is only when a mix of certainty and uncertainty in future environments exists that this flexibility will result in non-path policies. Only between the extremes of a deterministic and uniformly random demand environments will the ability of a system to be modified in response to changes in environments or requirements be exercised. That is, only *between* these entropy rate extremes is flexibility usefully exercised.

The observations of this section also emphasize where flexibility exists in the analyses of Step 4 that produced Pareto frontiers in terms of performance and cost objectives: While flexibility itself is not an explicit objective (cf. Section 2.5), it exists implicitly in the search for policies, particularly non-deterministic (non-path) policies. The ability of a system to be modified in response to changing environments or requirements (i.e., flexibility) is inherently linked to the rules that govern its modification response (i.e., policy), whether those rules are themselves explicit or implicit. This thesis advances work on flexibility by making explicit these system modification policies in a unification of traditional two-period state-centric concepts of flexibility with multi-period decision analysis techniques. As a result, policies governing *how system flexibility is used* are not only explicit and transparent, but can be traded and optimized to match a decision-maker's cost and performance preferences.

4.6. Summary

In a methodical manner, this chapter began from a foundational two-period state-centric concept of flexibility and showed how, through interpretation of this concept for space systems and linkage to the environments in which these systems may be required to operate, it can be unified with existing formulations and optimization techniques for Markov decision processes. Throughout the five-step framework developed here, several insightful analyses were developed. For example, in Step 1 the number of available transitions from a given configuration state at a given budget $\Phi_i(b)$ was developed as a surrogate metric for flexibility. Step 4 made the important distinction between paths and policies; while paths are a more traditional method of planning (e.g., laying out a set of actions to execute in future years), they preclude a decision-maker from considering the full "playbook" of if-then possibilities when making his or her decisions. Step 5 illustrated how the complicated policy (and, to some extent, path) results of Step 4 can be distilled into information that a decision-maker can use to make an initial system selection. For the more advanced practitioner, Step 5 addressed how the expected-value optima of Step 4 can be used as reasonable initial guesses for more local design space searches in the case that decision-makers have non-expected-value or non-cumulative objectives in mind. Finally, Step 5 also addressed the intriguing point that flexibility has a particular niche in environments of neither very high nor very low uncertainty, but rather in environments in which the present gives just *some* information about future demand. Emphasized was the inherent link between flexibility and policy, which specifies the conditions under which a system's flexibility is exercised.

This chapter has thus established the theoretical foundations of the present thesis. The following two chapters illustrate how this framework can be applied to problems of current interest to the space industry. First, in Chapter 5 comes a direct application to a relevant defense-related problem motivated by recent fractionation efforts of the Defense Advanced Research Projects Agency (DARPA). Second, in Chapter 6 is a significant extension of the basic framework presented in Chapter 4 to address decision-making for NASA human space exploration architecture planning.

CHAPTER 5

APPLICATION TO DISTRIBUTED- OR MULTI-PAYLOAD SATELLITE DESIGN DECISIONS

In July 2007, the U.S. Defense Advanced Research Projects Agency (DARPA) released a Broad Agency Announcement soliciting proposals for development of System F6 (Future Fast, Flexible, Fractionated, Free-Flying Spacecraft united by Information eXchange) [92]. DARPA's goal for F6 is ultimately a flight demonstration of an architecture in which the functionality of a traditional "monolithic" satellite is fulfilled with a fractionated cluster of free-flying, wirelessly interconnected modules. One special reference case defined in the context of fractionated spacecraft studies is that of a distributed-payload monolith satellite, in which payloads but not subsystems are distributed among free-flying modules (e.g., see Refs. [95]-[97]). To illustrate this thesis' framework in a step-by-step manner for a realistic application of intra-mission flexibility [28], this chapter poses an example in which design decisions must be made for a hypothetical multi-payload Department of Defense satellite system motivated by such a distributed-payload monolith concept. Of particular interest is the answer to the following question: How can a systems engineer or analyst select the design of the satellite system initially such that it can optimally (or Pareto-optimally) respond to the uncertain future demands that may be placed upon it?

Recall that this thesis' framework consists of five basic steps, outlined in Figure 12. First, system configuration options are identified and costs of switching from one configuration to another are compiled into a cost transition matrix. Second, probabilities that demand on the system will transition from one mission to another are compiled into a mission demand Markov chain. Third, one performance matrix for each design objective is populated to describe how well the identified system configurations perform in each of the identified mission demand environments. Fourth, possible future sequences of system configurations are simulated and sequences that are Pareto-optimal in terms of the decision-maker's objectives are identified. In a complementary approach, the system decision problem is formulated as a multi-objective variant of a Markov decision process, and Pareto-optimal decision policies are identified. Finally, the paths and policies from the latter step are synthesized into a set of data to inform initial system selection.

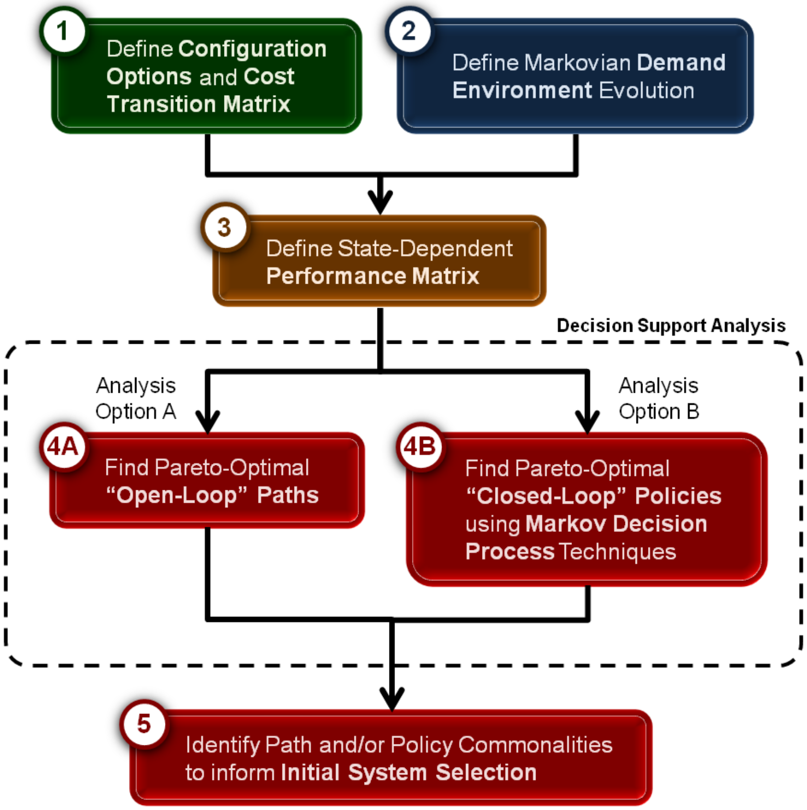
1
Define Configuration Options and Cost Transition Matrix
2
Define Markovian Demand Environment Evolution
3
Define State-Dependent Performance Matrix
Decision Support Analysis
Analysis Option A
Analysis Option B
4A
Find Pareto-Optimal "Open-Loop" Paths
4B
Find Pareto-Optimal "Closed-Loop" Policies using Markov Decision Process Techniques
5
Identify Path and/or Policy Commonalities to inform Initial System Selection

Figure 12. Five major steps of this thesis' framework.

5.1. Step 1: Define Configuration Options and the Cost Transition Matrix

As noted in Section 2.2, in 1984 economists Jones and Ostroy [51] suggested, "Flexibility is a property of initial positions. It refers to the cost, or possibility, of moving to various second period positions." Step 1 of this proposed framework begins by defining: What are the possible "positions", or engineering configurations, of this multi- or distributed-payload satellite system?

5.1.1. Defining the Configuration Space

In this application, suppose that the decision-maker has the option of inserting up to three specific payloads in any current or future system designs. One payload (PL1) provides detection of distress transmissions, another (PL2) provides high-bandwidth communications, and a third (PL3) provides high-resolution imagery. Assumptions for mass, power, and pointing requirements for these payloads are shown in Table 22.[*] Considering that these three payloads can be distributed among up to three on-orbit modules and that not all three payloads need be included in the system design (i.e., that omitting payloads is a valid consideration), there exist 15 distinct configuration options. These configurations are represented graphically in Figure 40 and, as noted in previous work [95], can be decomposed into subsets of configurations described by Bell numbers. Starting from the bottom, configurations 11-15 represent all possible ways of distributing three payloads among between one and three modules (i.e., from monolithic to fully fractionated). Configurations 5-10 cover all possible ways of distributing combinations of two payloads among up to two modules. Configurations 2-4 are the single-payload satellite system options, and Configuration 1 indicates the option to field no system at all.

Table 22. Assumed payload characteristics for example design. [98]-[102]

Payload No.	Payload Description	Flight Heritage	Mass (kg)	Power Requirement (W)	Pointing Requirement (deg.)
1	Search & Rescue Repeater	NOAA-N	24.0	53	1.00
2	LEO Transponders	Orbcomm	8.4	10	5.00
3	High Resolution Imager	NigeriaSat-2	41.0	55	0.01

---

[*] This list of payloads is limited to three for demonstration purposes only and can easily be increased if a decision-maker wishes to consider additional candidate payloads.

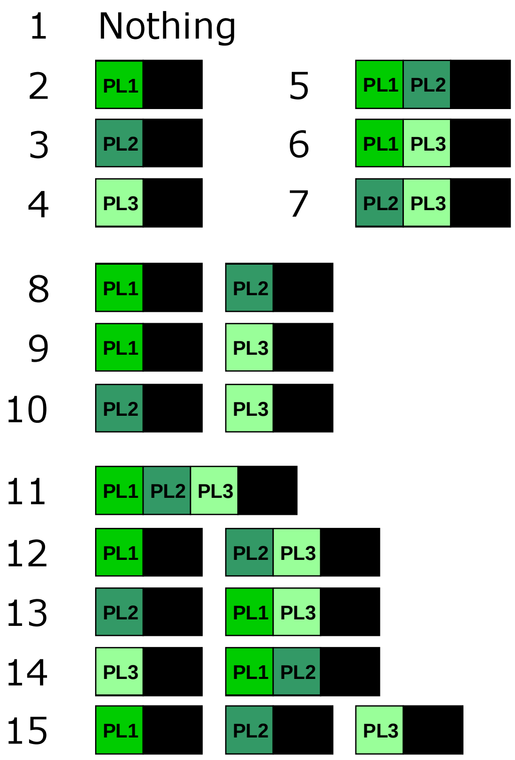

Figure 40. Possible system configurations. *Each distinct rectangular block represents a free-flying module. The payloads inside each module are indicated in green.*

Even at this early point in the process, enumeration of the designs within the configuration space reveals two extremes in approaches for evolving the system to meet future needs: The most modular (but in the long term, potentially costly) approach would be to launch new single-payload modules as new payloads are needed. A robust (but in the short term, potentially wasteful) approach would be to launch a single spacecraft with all three payloads, betting that all capabilities will eventually be required. A number of approaches fall between these extremes, and an important goal is to find the best possible sequence of configurations over the system's time horizon, given the uncertainty in future demand or requirements. One of the most important results of this search is eventual identification of the best possible initial design (i.e., what the decision-maker should build at the start of the program).

5.1.2. Defining the Cost Transition Matrix

With possible system configurations defined, it is next necessary to compile development and operations transition cost information. In this application, operations costs refer to the total costs required to operate the currently-fielded configuration over the coming time period. Development costs refer to the total costs required to design, develop, produce, and launch the components needed to transition from the current configuration to a new configuration over the coming time period.

In this application, suppose the decision-maker encounters a decision point every 30 months. At these points, a decision must be made regarding which of the 15 system configurations to develop and then field 30 months later. Demand for payload services in each 30-month operations period is uncertain a priori and materializes after development, with the possibility that it will then change in subsequent period (see Figure 41).

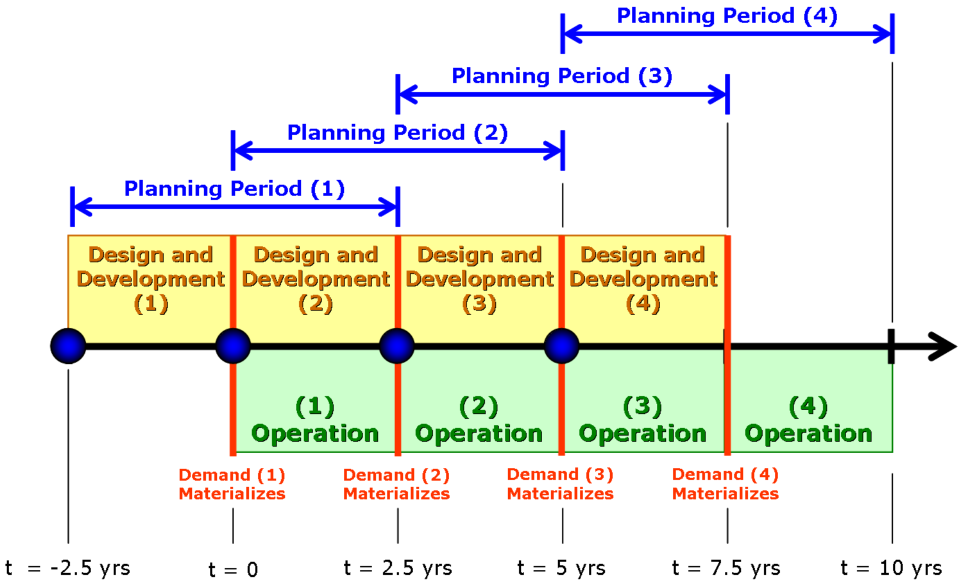

Figure 41. Planning periods and decision points for the distributed-payload satellite example.

Thus, the decision-maker has control over the system configuration but not the demand environment at each time step. However, at each decision point, the control that the decision-maker chooses to exercise comes at a certain cost. For example, if the decision-maker is at the second decision point and has Config. 2 already on-orbit, in order to transition to Config. 8 he/she would need to expend the appropriate resources to develop and launch a new module. In addition, he/she must simultaneously pay for the operation of the current on-orbit system.

These transition costs can be represented in matrix form. First, a development (or nonrecurring) cost matrix C_{dev} accounts for the one-time costs required to develop and produce one system given that another system already exists. This cost, which can also be considered a switching cost, is the cost most central to the notion of flexibility and may be computed through application-specific cost estimating relationships. In this case, application of the GT-FAST fractionated architecture synthesis tool [95],[103] using the payload assumptions of Table 22 for a 10-year design lifetime in a 410 km circular orbit produces the transition cost estimates in Table 23. These costs include appropriate spacecraft subsystem development and first-unit production, program management and systems engineering, software, ground segment development, launch, and assembly, test, and launch operations (ATLO).

Table 23. Development cost transition matrix, C_{dev} (data in \$FY08M).

		To Configuration														
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
From Configuration	1	0	169	131	184	175	200	189	197	252	212	204	257	228	258	280
	2	0	0	36	89	80	105	94	36	89	117	109	94	134	163	117
	3	0	75	0	89	80	105	94	75	158	89	109	163	105	163	158
	4	0	75	36	0	80	105	94	103	75	36	109	163	134	80	103
	5	0	75	36	89	0	105	94	103	158	117	109	163	134	89	186
	6	0	75	36	89	80	0	94	103	158	117	109	163	36	163	186
	7	0	75	36	89	80	105	0	103	158	117	109	75	134	163	186
	8	0	0	0	89	80	105	94	0	89	89	109	94	105	163	89
	9	0	0	36	0	80	105	94	36	0	36	109	94	134	80	36
	10	0	75	0	0	80	105	94	75	75	0	109	163	105	80	75
	11	0	75	36	89	80	105	94	103	158	117	0	163	134	163	186
	12	0	0	36	89	80	105	0	36	89	117	109	0	134	163	117
	13	0	75	0	89	80	0	94	75	158	89	109	163	0	163	158
	14	0	75	36	0	0	105	94	103	75	36	109	163	134	0	103
	15	0	0	0	0	80	105	94	0	0	0	109	94	105	80	0

Importantly, note that Table 23 accounts for the fact that free-flying modules for the next-period architecture need not be developed or produced if they exist already within the on-orbit cluster. The most obvious manifestation of this is that the diagonal of matrix C_{dev} consists entirely of zeros; this signifies the intuitive fact that it costs nothing to develop configuration i given that configuration i already exists. Similarly, note that no development costs are required to downgrade a configuration, such as a transition from Config. 15 (which, as shown in Figure 40, includes three single-payload modules) to Config. 2 (which consists of only the PL1 single-payload module). This highlights a simplifying assumption within the data of this particular matrix that the cost to shut down or decommission a module is zero; however, given proper decommissioning cost models, this information could easily be included in C_{dev}.

Second, a recurring cost matrix C_{rec} shown in Table 24 accounts for operations and any production beyond the first unit.[*] In this example application, first-unit production costs are the only applicable production costs, so the costs within this matrix are functions only of the row, i.e., the configuration that is operational over the length of the coming 30-month time period. These costs are also estimated using the GT-FAST tool, which draws upon a publicly-available NASA mission operations cost model [104].

---

[*] In some instances, the analyst may wish to account for all of production within the recurring cost matrix, since even one-time production for a unique flight unit is traditionally bookkept as a recurring cost. In the present application, one-time module production costs are considered more closely related to the one-time development costs and are accounted for in the development cost matrix.

Table 24. Recurring cost transition matrix, C_{rec} (data in $FY08M).

		To Configuration														
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
From Configuration	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	2	20	20	20	20	20	20	20	20	20	20	20	20	20	20	20
	3	16	16	16	16	16	16	16	16	16	16	16	16	16	16	16
	4	22	22	22	22	22	22	22	22	22	22	22	22	22	22	22
	5	21	21	21	21	21	21	21	21	21	21	21	21	21	21	21
	6	23	23	23	23	23	23	23	23	23	23	23	23	23	23	23
	7	22	22	22	22	22	22	22	22	22	22	22	22	22	22	22
	8	23	23	23	23	23	23	23	23	23	23	23	23	23	23	23
	9	29	29	29	29	29	29	29	29	29	29	29	29	29	29	29
	10	25	25	25	25	25	25	25	25	25	25	25	25	25	25	25
	11	24	24	24	24	24	24	24	24	24	24	24	24	24	24	24
	12	29	29	29	29	29	29	29	29	29	29	29	29	29	29	29
	13	26	26	26	26	26	26	26	26	26	26	26	26	26	26	26
	14	29	29	29	29	29	29	29	29	29	29	29	29	29	29	29
	15	31	31	31	31	31	31	31	31	31	31	31	31	31	31	31

Summing C_{dev} and C_{rec} from Table 23 and Table 24 yields the total cost transition matrix C in Table 25 Each element $c_{i,j}$ of this matrix specifies the total cost incurred over a subsequent 30-month period as the result of the decision to transition from developing configuration i to developing configuration j. For example, to transition from Config. 2 to Config. 8 requires developing, producing, and launching the module containing PL2 as well as operating the current Config. 2, for a total transition cost $c_{2,8}$ = $56 million.

Table 25. Total cost transition matrix, C (data in $FY08M).

		To Configuration														
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
From Configuration	1	0	169	131	184	175	200	189	197	252	212	204	257	228	258	280
	2	20	20	56	110	101	126	115	56	110	138	130	115	154	184	138
	3	16	91	16	106	97	122	111	91	174	106	125	179	122	180	174
	4	22	96	58	22	102	127	116	124	96	58	131	184	155	102	124
	5	21	96	57	110	21	126	115	123	178	138	130	184	155	110	207
	6	23	98	59	113	104	23	118	126	181	141	133	186	59	187	209
	7	22	97	58	112	103	128	22	125	180	140	131	97	156	186	208
	8	23	23	23	113	104	129	118	23	113	113	132	118	129	187	113
	9	29	29	65	29	109	134	123	65	29	65	138	123	162	109	65
	10	25	99	25	25	105	130	119	99	99	25	134	187	130	105	99
	11	24	98	60	113	104	129	118	126	181	141	24	186	157	187	209
	12	29	29	65	118	109	135	29	65	118	147	138	29	163	193	147
	13	26	101	26	116	107	26	121	101	184	116	135	189	26	190	184
	14	29	104	65	29	29	135	124	132	104	65	138	192	163	29	132
	15	31	31	31	31	112	137	126	31	31	31	141	126	137	112	31

5.1.3. Analyzing the Cost Transition Matrices

The data represented by the cost transition matrices can be analyzed, visualized, and related to flexibility in several useful ways. First, the relative trade between system initial costs and the switching costs (or one-time development costs) of Table 23 can be visualized as in Figure 42. In this figure, each vertical line indicates the range of switching costs from a given configuration, defined by the rows of Table 23. Solid dots indicate minimum and maximum values, and triangles indicate median values. Each vertical line is located horizontally at the cost needed to develop the configuration from scratch (in this case, Config. 1). For example, if no system currently exists and a decision-maker chooses to develop Config. 5 (involving a single module with PL1 and PL2 on board), a cost of \$175 million is incurred (on the *x*-axis), and the cost to switch configurations in the future varies from \$0 to \$186 million, depending on which future configuration is chosen. In contrast, if the decision-maker instead chooses to develop Config. 15 (involving three payloads among three modules), a cost of \$280 million is initially incurred, and the cost to switch configurations in the future varies from \$0 to \$109 million. Thus, to some extent Figure 42 empirically confirms the intuitive trend that future switching costs can often be reduced by earlier investments.

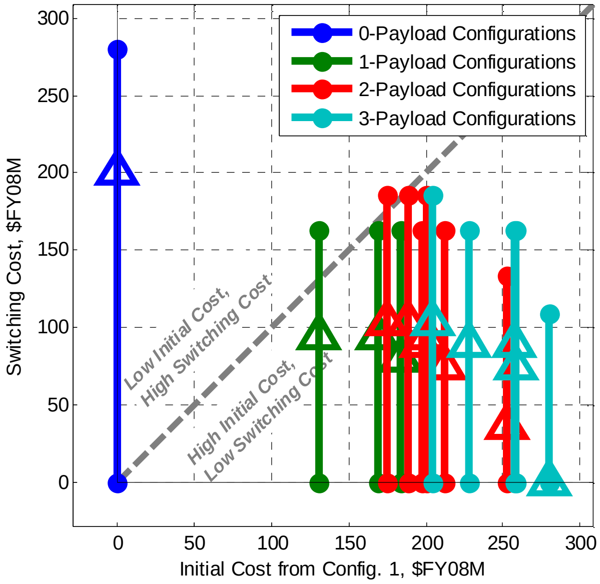

Figure 42. Switching cost vs. initial cost from Config. 1. *Vertical lines indicate ranges of switching costs from each configuration; some overlap. Solid dots indicate minima and maxima, and triangles indicate median values.*

Second, the data from the total cost transition matrix (Table 25) can be visualized directly in the context of the two-period state-centric notion of flexibility mentioned earlier. For this visualization see Figure 43. Here, each node in each of the three plots represents one of the configurations considered in the design space. Each node is named S_X, where X is the configuration number from Figure 40, and has a color indicative of the number of on-board payloads (consistent with the colors of Figure 42). Above each of the three plots is a budget, and for every element of the total cost transition matrix less than or equal to the given budget, a directed link is drawn. In cases where the total cost on the diagonal of the matrix is less than or equal to the budget, a dark circle is drawn around the appropriate node. For example, the middle plot of Figure 43 shows that, if the currently-fielded architecture is Config. 12, a $50 million budget for a given 30-month period would allow the decision-maker to transition to Configs. 1, 2, or 7, or to remain in Config. 12. In cases where no links or dark circles are associated with a configuration, the available budget is insufficient even to support operation of the current configuration into the next period.

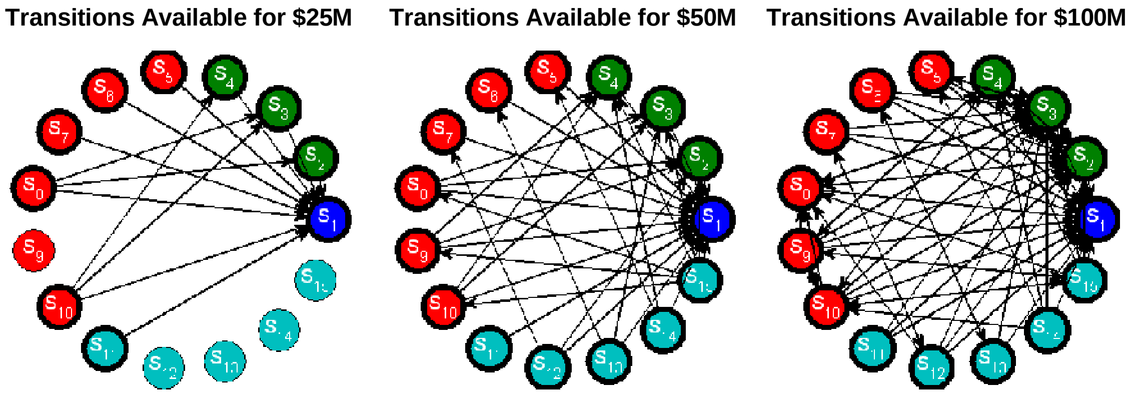
Transitions Available for $25M
Transitions Available for $50M
Transitions Available for $100M

Figure 43. Available transitions for three example 30-month budgets. *Self-transitions are available if a dark ring circles a given node. Colors indicate each configuration's number of payloads (0 = blue, 1 = green, 2 = red, 3 = cyan, consistent with Figure 42).*

A natural observation from Figure 43 is that, as budget is increased, more links become available. The total number of links in the graphs of Figure 43 increases from 23 at the $20 million budget to 47 at the $50 million budget and 78 at the $100 million budget. Eventually, at a large enough budget, all 225 links would appear. Linking this to the two-period state-centric concept of flexibility, a clear indicator of the flexibility of a given configuration *i* is the number of links or transitions available to it for a given budget *b* (the number of "outs" available, denoted $\Phi_i(b)$).

This indicator is plotted in Figure 44. The figure shows the number of available transitions as a function of available budget, where data for each configuration is represented by a single line. For example, the figure illustrates that for a per-period budget of $50 million, Config. 1 (the "nothing" configuration) has $\Phi = 1$ transition available, Configs. 2-7 and 11 each have $\Phi = 2$ available transitions, Configs. 8-10 and 12-14 each have $\Phi = 4$ available transitions, and Config. 15 has $\Phi = 8$ available transitions. It also shows that by a budget of $300 million, any configuration can be reached from any other configuration since all configurations have 15 available transitions.

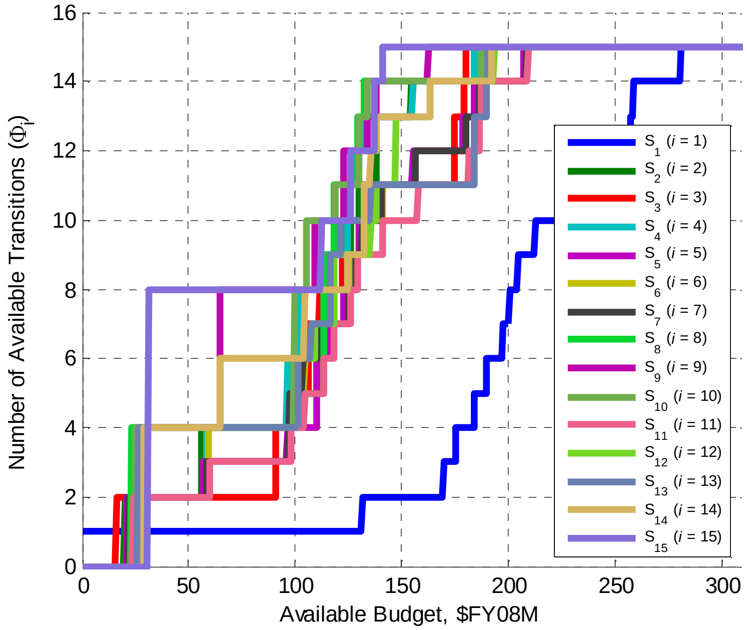

Figure 44. Available configuration transitions as a function of the available 30-month budget.

Figure 44 highlights a few interesting transition characteristics for the configurations in the design space defined in Figure 40. If the quantity Φ is interpreted as a surrogate measure of flexibility, then it is easily seen that Config. 1 is significantly less flexible than any other configuration over most of the budget range plotted in Fig. 9. For Config. 1, the first available transition to another configuration occurs at $131 million; for the same budget, other configurations can already make between 9 and 13 transitions. This occurs because Config. 1 has no capabilities that it can leverage to easily transition to other configurations, and all capabilities must be developed from scratch. It is also relevant to note that the three-payload monolith, Config. 11, which has no modules in common with other configurations, tends to have fewer transitions available than most other configurations at most budget levels. On the other hand, Config. 15 (the fully fractionated design) very quickly attains a large number of available transitions as budget increases; this configuration is the first to reach 8 transitions and the first to attain the ability to make all 15 available transitions. This occurs because Config. 15 consists of three single-payload modules that can easily be used as pieces of other configurations;

from Config. 15, the only modules that must be developed to reach other configurations are the two- or three-payload modules.

In terms of number of transitions, the other configurations within the design space generally fall between the bounds of Configs. 1 and 15. All illustrate that Φ is a monotonically increasing function of budget, which implies that any given configuration's flexibility increases with available budget. However, examples also can be found to illustrate that the *relative* flexibility between configurations is also a function of available budget. For example, at a budget of \$25 million, Config. 8 has four available transitions while Config. 15 has none. In other words, at a budget of \$25 million, it is reasonable to make the statement that Config. 8 is more flexible than Config. 15. However, at a budget of \$50 million, Config. 8 still has four available transitions while Config. 15 can make eight transitions. At this budget level, Config. 15 is more flexible than Config. 8, and the relative flexibility of these configurations has reversed. The reason for this "flexibility reversal" becomes evident when it is recalled that the cost transition matrix accounts for both development and recurring operations costs: When budget resources are scarce, operating a high-capability configuration (like Config. 15) consumes funds that would otherwise be available for developing the components needed to transition to another configuration. However, as financial resources become more abundant, more capable configurations become more flexible because they already possess capabilities transferrable toward the development of other configurations.

5.2. Step 2: Define Markovian Demand Environment Evolution

While Step 1 focused on defining the available configuration states for the distributed-payload satellite application of interest, the environment in which the system will operate has not yet been discussed. Step 2 fills this gap by proposing a model for the evolution of the environment. Unlike the configuration state, which is under the control of the decision-maker, the environment state will characterize the demands placed on the system at any given time, which inherently is not under the control of the decision-maker and evolves stochastically.

As mentioned in Section 5.1.1, up to three specific payloads are available for any current or future designs of the distributed-payload satellite system. One payload (PL1) provides detection of distress transmissions, another (PL2) provides high-bandwidth communications, and a third (PL3) provides high-resolution imagery. In terms of defining the demand environment, it is reasonable to expect that future demand may exist for the satellite system to provide any combination of these three services. For example, in one time period, only high-bandwidth communications may be required, and in another, both high-resolution imagery and high-bandwidth communications may be needed. Thus, there exist eight distinct demand environment states, indicated by the axes in Table 26. Note that these environment states are mutually exclusive and, for example, “1” should be interpreted as “1 only” and “1+2” should be interpreted as “1+2 only”.

It is also reasonable to expect that the evolution of demand for these services through time is unlikely to be properly modeled by a time series of independent random demand environments. Rather, a subsequent period’s demand likely depends at least in part upon the current demand, a dependence that can be captured using a Markov chain. stochastic model with an associated probability transition matrix. The particular probability transition matrix assumed for this example is shown in Table 26. Ideally, this matrix would be populated using a set of expert judgements regarding future demand behavior or, if they exist, probabilities based on historical data. In this notional example, the author’s judgement was used to select values that reflected a high likelihood that a current demand would be maintained (e.g., if high-resolution imagery is demanded in the current period, it would be likely to also be demanded in the next period) and tended to place lower probabilities on the need for dedicated distress transmission detection services. The probabilities in Table 26 also reflect an assumed conditional independence in the evolution of demand for each individual service; for example, the probability of demand evolving for all three services (1+2+3) in a subsequent period is equivalent to the product of three underlying probabilities that are conditional on demand in the current period and reflect the likelihood that demand evolves to each of the services individually. It is important to emphasize, however, that the particular probabilities in Table 26 are illustrative and can easily be substituted with if more data or other expert judgements become available.

Table 26. Assumed demand environment transition probability matrix.
Note that, in the demand environment naming convention, 1 indicates demand for distress transmission detection , 2 indicates demand for high-bandwidth communications, and 3 indicates demand for high-resolution imagery services.

		To Demand Environment							
		None	1	2	3	1+2	1+3	2+3	1+2+3
From Demand Environment	None	0.30	0.05	0.13	0.30	0.02	0.05	0.13	0.02
	1	0.20	0.15	0.09	0.20	0.06	0.15	0.09	0.06
	2	0.10	0.02	0.23	0.15	0.05	0.03	0.35	0.07
	3	0.10	0.08	0.07	0.23	0.05	0.19	0.16	0.12
	1+2	0.05	0.07	0.20	0.03	0.28	0.05	0.13	0.19
	1+3	0.05	0.05	0.05	0.20	0.05	0.20	0.20	0.20
	2+3	0.05	0.04	0.12	0.12	0.09	0.09	0.27	0.22
	1+2+3	0.02	0.02	0.08	0.08	0.08	0.08	0.32	0.32

The Markov chain of Table 26 can be visualized as a set of demand environment states as in Figure 45. In this figure, high-probability transitions are represented as thick dark links and low-probability transitions are represented as thin light links. The likelihood of self-transitions (along the diagonal in Table 26) are indicated by the darkness and thickness of rings around each state. Thus, for example, this figure immediately allows identification of the highest-probability and lowest-probability transitions in the Markov chain and demand environment evolution.

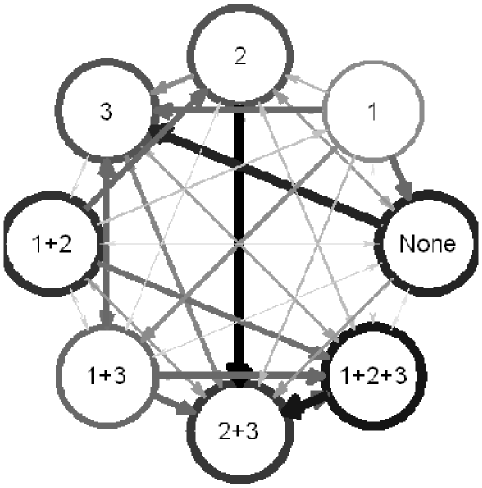

Figure 45. Visualization of the demand environment Markov chain described by Table 26.

5.3. Step 3: Define State-Dependent Performance Matrix

Linking the on-orbit configuration to the demand environment is a matrix that specifies the amount of reward (e.g., revenue or accumulated performance measure) earned in each time period as a function of the demand environment and system configuration in that period. The application here uses the matrix in Table 27, which specifies the number of demanded services that are performed given a particular configuration operating in a particular demand environment. For example, if the demand in one time period is for imagery and communications (column 7) and the vehicle on-orbit is in Config. 15 (the 3-payload fully-fractionated option, row 15), the decision-maker accumulates the successful performance of two demanded services. As a result, the decision-maker is incentivized to place payloads in orbit that will meet demand for services.

It is also worth noting that, although the present application adopts just one performance metric (and thus one performance matrix), multiple such matrices can be defined for any cumulative performance metrics of interest to the decision-maker. For example, a decision-maker may also be interested in a cumulative binary metric that indicates a 1 or 0 in each time period depending on whether performance demands were fully met; over the long term, such a metric would indicate the percentage of time that the system fully meets the demands placed upon it.

Table 27. Performance matrix quantifying the number of demanded services performed in a given time period.

	Demand Environment State							
Configuration State	None	1	2	3	1+2	1+3	2+3	1+2+3
1	0	0	0	0	0	0	0	0
2	0	1	0	0	1	1	0	1
3	0	0	1	0	1	0	1	1
4	0	0	0	1	0	1	1	1
5	0	1	1	0	2	1	1	2
6	0	1	0	1	1	2	1	2
7	0	0	1	1	1	1	2	2
8	0	1	1	0	2	1	1	2
9	0	1	0	1	1	2	1	2
10	0	0	1	1	1	1	2	2
11	0	1	1	1	2	2	2	3
12	0	1	1	1	2	2	2	3
13	0	1	1	1	2	2	2	3
14	0	1	1	1	2	2	2	3
15	0	1	1	1	2	2	2	3

5.4. Step 4: Decision Support Analysis

With configuration transitions, demand environment transitions, and a performance matrix defined, there now exists enough information to begin to answer the question of what is the “best” initial configuration the decision-maker can choose. Using Figure 41 as a framework for a simulation timeline, one time period before a configuration is fielded (in this distributed-payload satellite example, at *t* = -2.5 years), a decision-maker must choose which system configuration to initially design, develop, and produce. At *t* = 0, the system that had been developed over the previous time period is fielded, and a demand environment materializes. At this point, the system operator must make use of the currently operational system in attempting to fulfill the current demand. Meanwhile, the decision-maker must choose which configuration to design, develop, and produce over the coming period. The cycle then repeats for as many periods as fills the time horizon under consideration. In this case, the time horizon of interest is 10 years of operation.

The decision support analysis in this step is divided into two complementary analysis options. The first option, in which Pareto-optimal paths are identified, is simpler to implement and conceptually similar to long-term scheduling and roadmapping analysis. The second option, in which Pareto-optimal policies are identified, is a more complete consideration of the problem and is akin to developing an optimal "playbook" of what actions to take given all possible future evolutions of the environment.

5.4.1. Find Pareto-Optimal "Open-Loop" Paths

One question that Figure 41 prompts is: What configuration should the decision-maker choose to develop at each time increment? In other words, what configuration should be selected for each of the yellow design and development blocks in Figure 41? The answer is not obvious, especially since the demand environment evolves stochastically. For example, the decision-maker who wishes to be able to fulfill whatever demand the next period may bring would choose to build the most capable system possible, but this may come at substantial initial expense. The decision-maker who would gamble that tomorrow's demand will be the same as today's would develop few or no new architectural components and in doing so save significant resources; however, this may come with the inability to perform if the next period's demand materializes to require greater capability. Furthermore, whether one period's decision is best (e.g., high-reward or low-cost in the long run) is likely to be dependent on other decisions throughout the system lifetime. In the flexibility problem, it is in general necessary to consider all future decisions within a given time horizon in order to judge the appropriateness of any single decision. While this presents a unique difficulty within the realm of space system conceptual design, once complete it presents an automatic solution to the question of which configuration to select initially: The appropriate configuration to select initially is the first configuration decision from the "best" time-ordered sequence of decisions.

In this example, posing the problem such that we wish to find the optimal sequence of the four development decisions (each decision of which implies a selection among the 15 configuration options) means that there exist $15^4 = 50{,}625$ possible sequences (or paths). Since the configurations on these paths are identified by the time on the clock at which they are chosen, this type of specification will be referred to as an open-loop path.

Assuming an initial condition at $t = -2.5$ years in which the operational configuration is nothing (Config. 1) and there is demand for none of the services (the "None" environment), one approach to solving this problem is to simulate all 50,625 paths subject to the stochastically-changing demand environment and identify which produces the "best" combination of performance and cost. Thus, for each of the paths, 1000 Monte Carlo simulations are run. At each time step of a simulation, the following events and computations occur:

1. Mission demand evolves stochastically according to the Markov chain estimate of Table 26.
2. The operator of the currently operational configuration attempts to use this system to fulfill the new mission demand, earning credit according to the performance matrix.
3. The decision-maker chooses which configuration to develop in the current time period and field in the next time period, incurring a cost according to the cost transition matrix. An available choice in any time period is to retain the current configuration, which requires no additional development resources.

A sample set of Monte Carlo simulation results is shown in Figure 46. This figure shows the result of adopting a path representing an incremental buildup of capability in which Config. 4 (the PL3-only configuration) is fielded initially. In the next time period, a new module containing PL1 is launched, and PL2 is added in the third time period. The cluster of three modules operates until the end of the 10-year time horizon. Due to the simulation setup, a configuration decision must still be made in the final operational time period; since the cost of developing this final configuration will be incurred but no reward will be earned, Config. 1 (the "Nothing" configuration) is selected. As the bottom left portion of Figure 46 shows, this particular path (denoted as [4 9 15 15 1], by the configuration decisions made at each step) is subject to a stochastically changing demand environment. The size of each yellow dot indicates the likelihood of demand being in a particular state (on the *y*-axis) at a given time (on the *x*-axis); note that all simulations begin in the "None" demand environment at *t* = -2.5 years, as specified by the initial condition. The right-hand portion of Figure 46 indicates how per-period cost and performance vary over time. Note that the per-period cost decreases from $184 million for the initial investment to $31 million in the final operations period, and number of demanded services performed per period increases from zero to a mean of 1.67 in the final period. The total expected cost for this path over the time horizon is $407 million[*], and the total expected number of demanded services performed is 4.69.

---

[*] Note that once a path is chosen, cost is fixed. As a result, the expected cost is equivalent to the minimum, maximum, and median costs across all path-based Monte Carlo simulations.

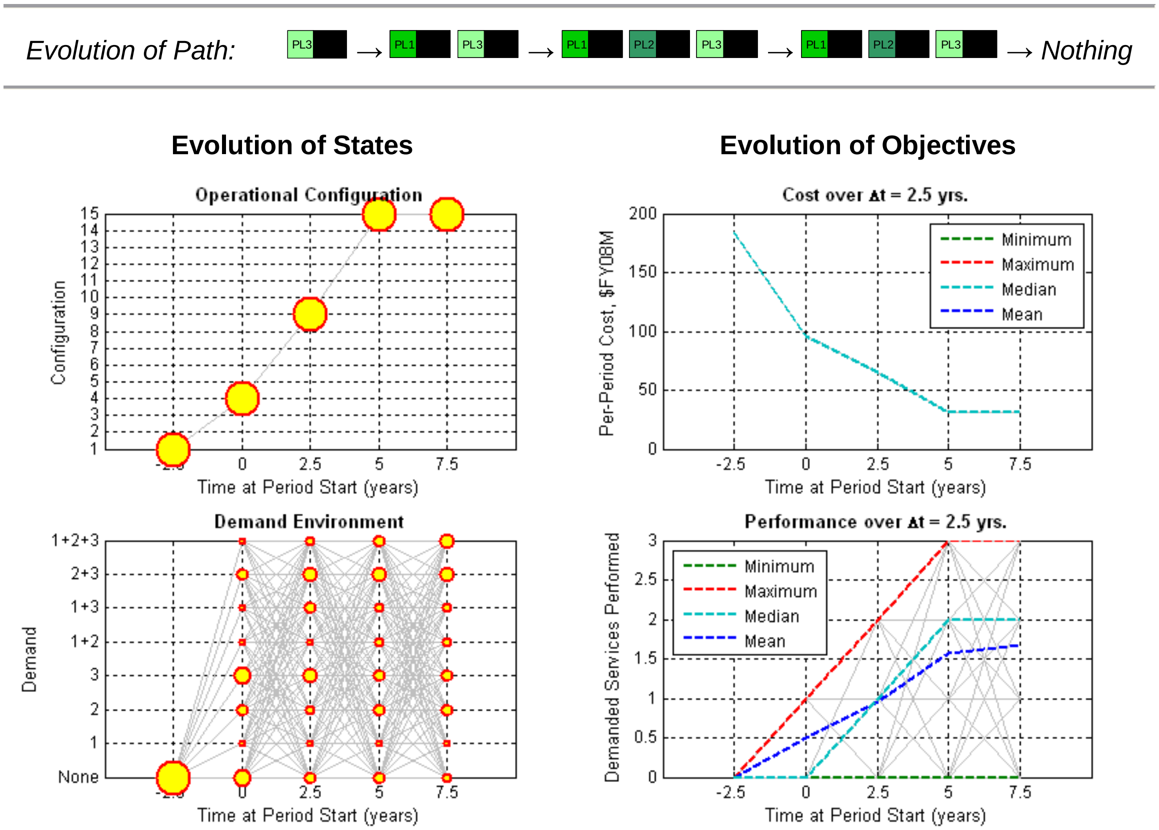

Figure 46. Evolution of configuration path [4 9 15 15 1], representative of an incremental capability buildup. *In the plots on the left, the size of circles indicate the relative number of Monte Carlo simulation cases that exist in a given configuration or demand environment state (on the y-axes) at a given time (on the x-axes). The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation. Note configuration and cost are deterministic, since a path is specified.*

Obtaining results like those in Figure 46 for each of the 50,625 possible paths allows the total expected performance to be computed and plotted against total cost for each path as in Figure 47. In this figure, each blue "x" represents the total cost and performance of one path[*]. Notice that, for the population as a whole, there is a general trend that, as more funds are invested, higher performance is expected. However, it is important to recall that the decision-maker has a choice of which path to select. As a

---

[*] These totals are taken over the t = -2.5 year period (at which there is zero performance due to the initial condition) and the four subsequent periods.

result, if he or she cares primarily about total cost and expected total demanded services performed, it would make little sense to select a high-cost, low-performance point toward the lower right of the cluster. Rather, the decision-maker would prefer to choose among the set of nondominated points that comprise the Pareto frontier. This Pareto frontier, shown in red in Figure 47, is composed of the set of possible configuration sequences for which one objective cannot be improved without the sacrifice of another. In this application, the frontier is comprised of just 12 of the 50,625 possible paths and helps to narrow the options considerably.

Listed next to each of the Pareto-optimal points in Figure 47 is its associated configuration path. Note that at the bottom left of the figure is the "do nothing" option in which Config. 1 is fielded for all time periods; this is cost-optimal but also provides the lowest possible performance. At the other extreme is the Pareto-optimal highest-performance option of fielding Config. 11, the three-payload monolithic satellite, for all time periods. The Pareto-optimal solutions between these two extremes involve developing Configs. 3, 5, 7, or 11, either immediately or after a 1-2 period delay. Notably absent from the frontier are the higher-cost multiple-module configurations.

One use of the data in Figure 47 becomes evident when the sample path from Figure 46 is overlaid as the yellow square in Figure 47. Here it can be seen that the incremental path [4 9 15 15 1] is dominated by solutions on the Pareto frontier. In fact, one particular path, [1 11 11 11 1], accumulates near-identical performance for a total cost about $131 million (32%) lower. In this Pareto-optimal path, detailed in Figure 48, the three-payload monolithic satellite is fielded after a one-period wait, during which time demand evolves toward an environment in which multiple services are demanded. Unlike the incremental path in Figure 46, which exhibits a gradual decrease in per-period cost, the Pareto-optimal path in Figure 48 exhibits an initial $204 million spike followed by $24 million in operations costs for three periods. As a result, this cost profile results in significant savings, and the system still performs well since all three payloads are available to fulfill all requested services at times in the future in which the environment has evolved to one in which multiple services tend to be demanded.

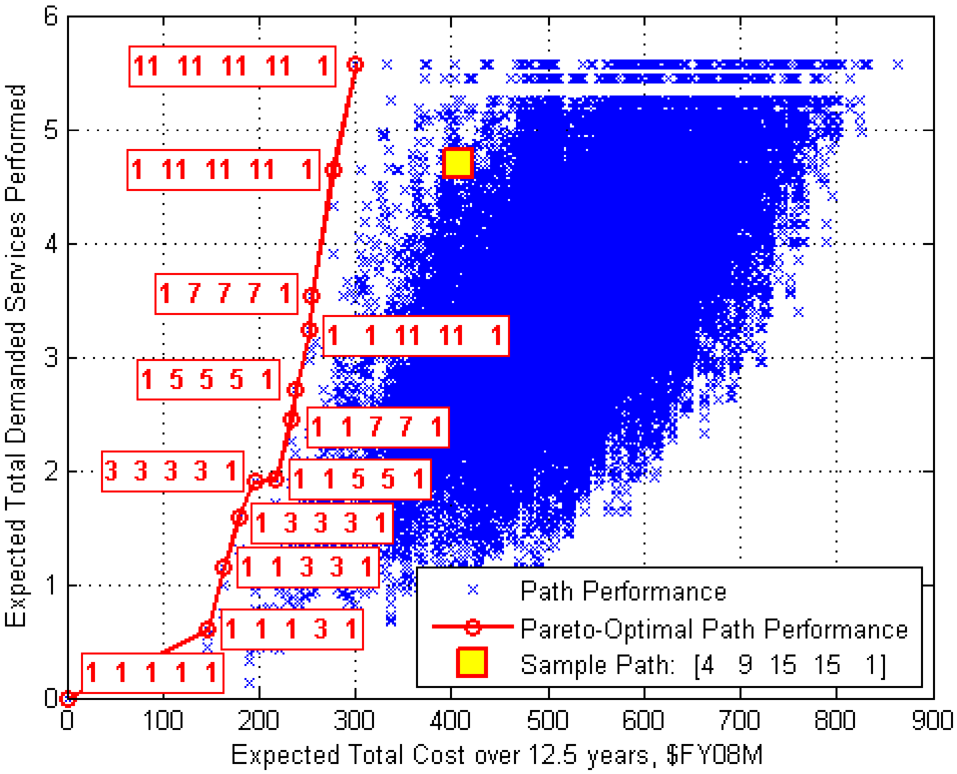

Figure 47. Trade between total demanded services performed and total cost for all open-loop paths. *Pareto-optimal paths are identified by 5-period configuration sequences listed next to red circles.*

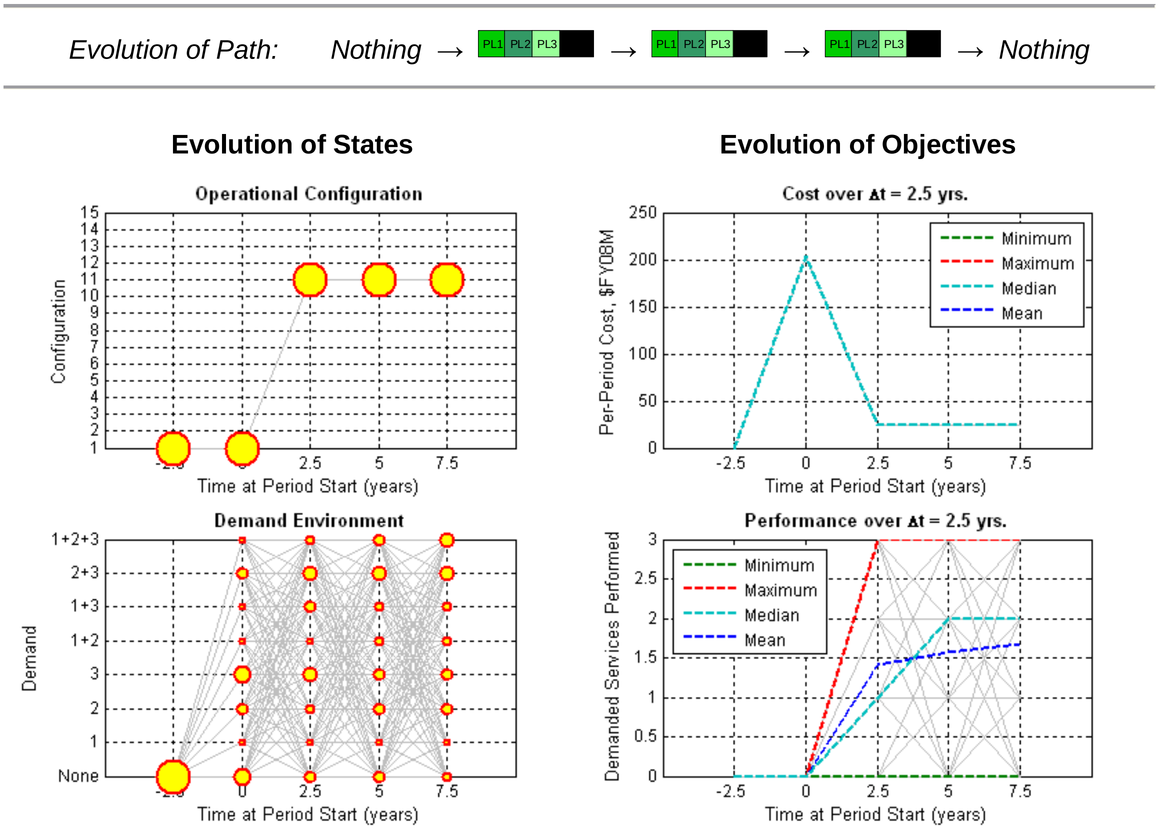

Figure 48. Evolution of configuration path [1 11 11 11 1], a Pareto-optimal path. *In the plots on the left, the size of circles indicate the relative number of Monte Carlo simulation cases that exist in a given configuration or demand environment state (on the y-axes) at a given time (on the x-axes). The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation. Note configuration and cost are deterministic, since a path is specified.*

5.4.2. Find Pareto-Optimal "Closed-Loop" Policies

While straightforward and conceptually similar to an optimization of typical long-term scheduling and roadmapping efforts, the analysis presented in Step 4A has two principal disadvantages. First, for applications with large numbers of configurations and long time horizons, it may not be practical to enumerate all possible paths. For example, if the number of time periods in the present application were doubled, the number of possible paths would increase from 50,625 to over 2.6 billion and take several years of run time on a standard desktop computer. Second, assuming a set path for the entirety of the system's lifetime neglects the ability of the decision-maker to make choices mid-program in response to the evolution of the demand environment.

To overcome these limitations, Step 4B presents a complementary analysis that draws on the unification of flexibility and Markov decision process (MDP) frameworks discussed in depth in Section 4.4.2.2. To proceed with this analysis, the configuration and environment state spaces of Steps 1 and 2 are combined into a single total state space (i.e., Total State = {Configuration State, Demand State}). In this distributed-payload satellite example, there are 15 configuration states × 8 environments = 120 total states, which Figure 49 illustrates graphically. In this three-dimensional "spindle" of total states, each vertical layer represents a particular demand environment and each column represents a particular configuration. Thus, it is possible for the fielded system to be in any configuration and operating in any demand environment at any particular point in time. Since configuration is under the control of the decision-maker, he or she can choose to move to any vertical column of the spindle at any point in time (recognizing that it takes one time step to make this move). However, the demand environment is not under the control of the decision-maker. Illustrated in Figure 49 is an instance where Config. 15 is operating in Demand Environment 1. If the decision-maker chooses to develop Config. 10 for the next time period, he or she is assured to move to the column corresponding to Config. 10*; however, since the demand environment evolution is stochastic, the layer to which he or she moves is uncertain and depends on the evolution of the Markov chain specified by Step 2. Once the demand environment materializes, the decision-maker finds himself or herself at one particular total state and makes another decision about which of the 15 configurations to select for the following period.

---

* The assumption implicit in this assurance is that the decision-maker will not by accident develop a configuration other than Config. 10, which is considered reasonable in this application.

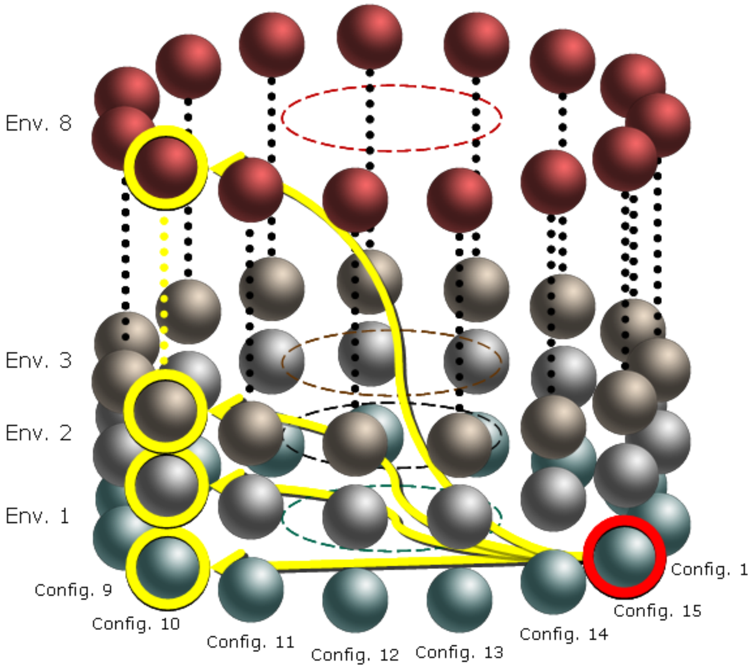

Figure 49. "Spindle" of Total States. *Each layer corresponds to one demand environment and each vertical column corresponds to one configuration. Environments 4-7 are not depicted. Arrows illustrate that, due to demand environment uncertainty, multiple possible total states are possible in the next period if a decision is made to transition from one configuration to another (e.g., Config. 15 to Config. 10).*

Accounting for probabilities of transitions within the total state space, configuration decisions from each state, and per-period cost and performance aggregation as detailed in Section 4.4.2.2, the solutions for Pareto-optimal decision policies take the form of a matrix with 120 rows and 5 columns, where each element (ξ,τ) indicates which of the 15 available configurations be developed next given the system is in state ξ at time period τ. If a full-factorial analysis of all possible policies were to be conducted (as was done for the simple case of paths in Step 4A), $15^{600} = 10^{706}$ simulations would need to be executed! However, use of the structure of the problem as posed in Section 4.4.2.2 permits optimal policy solutions to be found within hours or minutes on a standard desktop computer.

Expected cost and performance results for policy solutions to the distributed-payload satellite system application are shown by each blue "x" in Figure 50. Among these solutions, the nondominated (Pareto-optimal) solutions are highlighted and connected in red. Note that the minimum-cost and maximum-performance endpoints of the Pareto frontier are identical to those of the open-loop full factorial analysis of Figure 47, and the shape of the frontier largely mirrors that of Figure 47.[*] However, an interesting solution with performance superior to any available from an open-loop path is visible at an expected total cost of $40 million. Depicted in Figure 51 in the same format as the open-loop results earlier, it can be seen that this policy solution is nearly the same as the "do nothing" policy but with one exception: As the top left plot shows, at the $t = 0$ time period the policy occasionally (in 14.5% of cases) calls for a decision to develop and subsequently field the three-payload monolith. Whether decision is made is governed by the demand environment, as the policy indicates in Table 28. In this table, the policy solution itself is shown, and the action specified by the policy is provided for a system in any state s (the row) at any time t (the column). Looking only at the eight total states that are associated with Config. 1 (i.e., total states 1, 16, 31, 46, 61, 76, 91, and 106), it can be seen that the decision to develop Config. 11 rather than Config. 1 at $t = 0$ occurs only in total states 91 and 106, which correspond to a situation in which either the 2+3 or 1+2+3 demand environment exists. In other words, this policy achieves a low expected cost and an appreciable expected performance by only developing the three-payload monolith if a substantial demand for services materializes early during the program. Such a result is impossible to capture using the fixed configuration paths of Step 4A.

---

[*] The sparsity of points on this frontier is largely due its concavity: Only four of the frontier points could be found using $n = 1$ in Eq. (28). The heuristic method adopted for improving the frontier estimate by increasing n beyond unity was only partially successful in identifying the full frontier, and this is thus a clear area for future development.

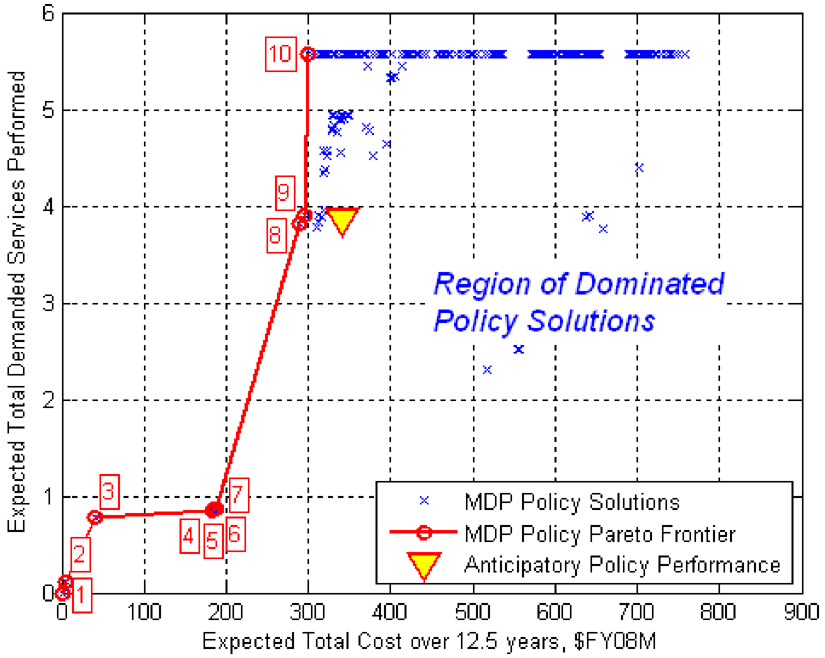

Figure 50. Trade between total demanded services performed and total cost for MDP policy solutions.

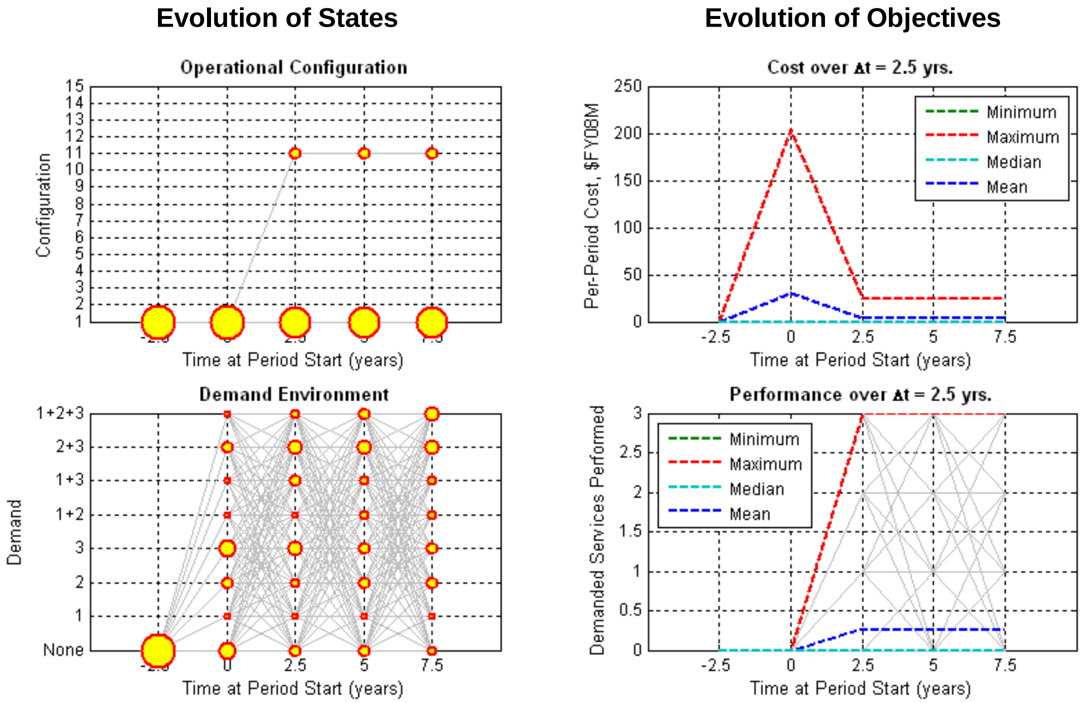

Figure 51. Evolution of states and objectives for Pareto-optimal policy #3 (defined in Table 28). *In the plots on the left, the sizes of circles indicate the relative number of Monte Carlo simulation cases that exist in a given configuration or demand state (on the y-axes) at a given time (on the x-axes). The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation.*

Table 28. Pareto-optimal policy #3. *Configuration decisions for a system in state s at time t are indicated by matrix elements shaded in gray.*

Current State, *s*			Time at Period Start (years), *t*					Current State, *s*			Time at Period Start (years), *t*				
Total State	**Env.**	**Config.**	**-2.5**	**0**	**2.5**	**5**	**7.5**	**Total State**	**Env.**	**Config.**	**-2.5**	**0**	**2.5**	**5**	**7.5**
1	None	1	1	1	1	1	1	61	1+2	1	11	1	1	1	1
2	None	2	11	12	1	1	2	62	1+2	2	11	12	8	8	1
3	None	3	11	11	3	1	3	63	1+2	3	11	11	3	3	3
4	None	4	14	10	4	4	1	64	1+2	4	14	14	14	10	4
5	None	5	14	5	5	5	1	65	1+2	5	5	5	5	5	5
6	None	6	13	13	6	6	6	66	1+2	6	13	13	13	13	6
7	None	7	7	7	7	7	1	67	1+2	7	7	7	7	7	7
8	None	8	11	8	8	8	1	68	1+2	8	11	8	8	8	2
9	None	9	15	15	9	9	4	69	1+2	9	15	15	15	15	9
10	None	10	10	10	10	10	10	70	1+2	10	14	10	10	10	3
11	None	11	11	11	11	11	11	71	1+2	11	11	11	11	11	1
12	None	12	12	12	12	12	2	72	1+2	12	12	12	12	12	1
13	None	13	13	13	13	13	6	73	1+2	13	13	13	13	13	6
14	None	14	14	14	14	14	1	74	1+2	14	14	14	14	14	4
15	None	15	15	15	15	15	2	75	1+2	15	15	15	15	15	15
16	1	1	11	1	1	1	1	76	1+3	1	11	1	1	1	1
17	1	2	11	11	2	2	2	77	1+3	2	11	12	12	2	2
18	1	3	11	11	3	1	3	78	1+3	3	11	11	11	3	1
19	1	4	14	14	4	4	4	79	1+3	4	14	14	10	4	1
20	1	5	11	5	5	5	1	80	1+3	5	14	14	5	5	5
21	1	6	13	13	6	6	1	81	1+3	6	13	13	13	6	6
22	1	7	7	7	7	7	7	82	1+3	7	7	7	7	7	7
23	1	8	11	8	8	8	2	83	1+3	8	11	11	8	8	8
24	1	9	15	15	15	9	1	84	1+3	9	15	15	15	9	9
25	1	10	10	10	10	10	10	85	1+3	10	10	10	10	10	1
26	1	11	11	11	11	11	1	86	1+3	11	11	11	11	11	11
27	1	12	12	12	12	12	2	87	1+3	12	12	12	12	12	12
28	1	13	13	13	13	13	3	88	1+3	13	13	13	13	13	3
29	1	14	14	14	14	14	1	89	1+3	14	14	14	14	14	5
30	1	15	15	15	15	15	9	90	1+3	15	15	15	15	15	4
31	2	1	11	1	1	1	1	91	2+3	1	11	11	1	1	1
32	2	2	11	11	8	1	1	92	2+3	2	11	11	12	2	2
33	2	3	11	11	3	3	1	93	2+3	3	11	11	3	3	3
34	2	4	14	10	10	4	1	94	2+3	4	14	14	10	4	4
35	2	5	11	5	5	5	1	95	2+3	5	14	5	5	5	5
36	2	6	13	13	13	6	1	96	2+3	6	13	13	13	6	1
37	2	7	7	7	7	7	1	97	2+3	7	7	7	7	7	7
38	2	8	11	8	8	8	8	98	2+3	8	11	8	8	8	2
39	2	9	15	15	15	9	1	99	2+3	9	15	15	15	9	4
40	2	10	10	10	10	10	4	100	2+3	10	10	10	10	10	4
41	2	11	11	11	11	11	11	101	2+3	11	11	11	11	11	1
42	2	12	12	12	12	12	2	102	2+3	12	12	12	12	12	12
43	2	13	13	13	13	13	1	103	2+3	13	13	13	13	13	13
44	2	14	14	14	14	14	4	104	2+3	14	14	14	14	14	4
45	2	15	15	15	15	15	8	105	2+3	15	15	15	15	15	15
46	3	1	11	1	1	1	1	106	1+2+3	1	11	11	1	1	1
47	3	2	11	11	12	2	2	107	1+2+3	2	11	11	12	8	2
48	3	3	11	11	3	3	1	108	1+2+3	3	11	11	11	3	3
49	3	4	14	14	10	4	1	109	1+2+3	4	14	14	10	10	1
50	3	5	14	5	5	5	1	110	1+2+3	5	11	14	5	5	5
51	3	6	13	13	13	6	6	111	1+2+3	6	13	13	13	13	6
52	3	7	7	7	7	7	7	112	1+2+3	7	7	7	7	7	1
53	3	8	11	8	8	8	3	113	1+2+3	8	11	11	8	8	8
54	3	9	15	15	15	9	9	114	1+2+3	9	15	15	15	15	4
55	3	10	10	10	10	10	4	115	1+2+3	10	10	10	10	10	3
56	3	11	11	11	11	11	1	116	1+2+3	11	11	11	11	11	11
57	3	12	12	12	12	12	1	117	1+2+3	12	12	12	12	12	2
58	3	13	13	13	13	13	6	118	1+2+3	13	13	13	13	13	3
59	3	14	14	14	14	14	5	119	1+2+3	14	14	14	14	14	1
60	3	15	15	15	15	15	10	120	1+2+3	15	15	15	15	15	15

Figure 50 also permits comparisons to be made with policies that might be brainstormed or proposed outside of the MDP solution procedure. For example, one reasonable policy that might be proposed is to always develop and field the configuration that least expensively maximizes performance in the most likely next-period demand environment.[*] The policy implied by this statement is provided in Table 29; for instance, if Config. 2 (the PL1-only configuration) is currently operational in the "1+2" demand environment (i.e., if the system is in total state 62), the most likely next-period demand environment according to Table 26 is also the "1+2" demand. To least expensively fulfill both the PL1 and PL2 functions demanded in this environment, a single PL2-only module would be developed and launched, which places the system into Config. 8. Thus, as Table 29 shows, Config. 8 is the decision made from total state 62 at all except the final time period.[†]

The performance of this next-period anticipatory policy is summarized by the yellow triangle in Figure 50 and detailed in Figure 52. Figure 50 in particular illustrates two interesting and important points regarding this anticipatory policy: First, this policy is dominated by others discovered in the optimization process: Both policies 9 and 10 on the Pareto frontier perform, on average, more demanded services at a lower cost. Second, this anticipatory policy is just one of many options; even if it were nondominated, selection of this particular policy carries with it no options regarding cost and performance preferences. In contrast, a search throughout the policy design space (as was completed in order to produce Figure 50) allows the decision-maker to understand the cost and performance trades available and select a policy according to his or her preferences.

---

[*] In the event that multiple demand environments have the same probability of materializing next, the environment with the demand for more services is used.

[†] The reason for the difference in the final time period decision is the same as discussed earlier in Section 5.4.1.

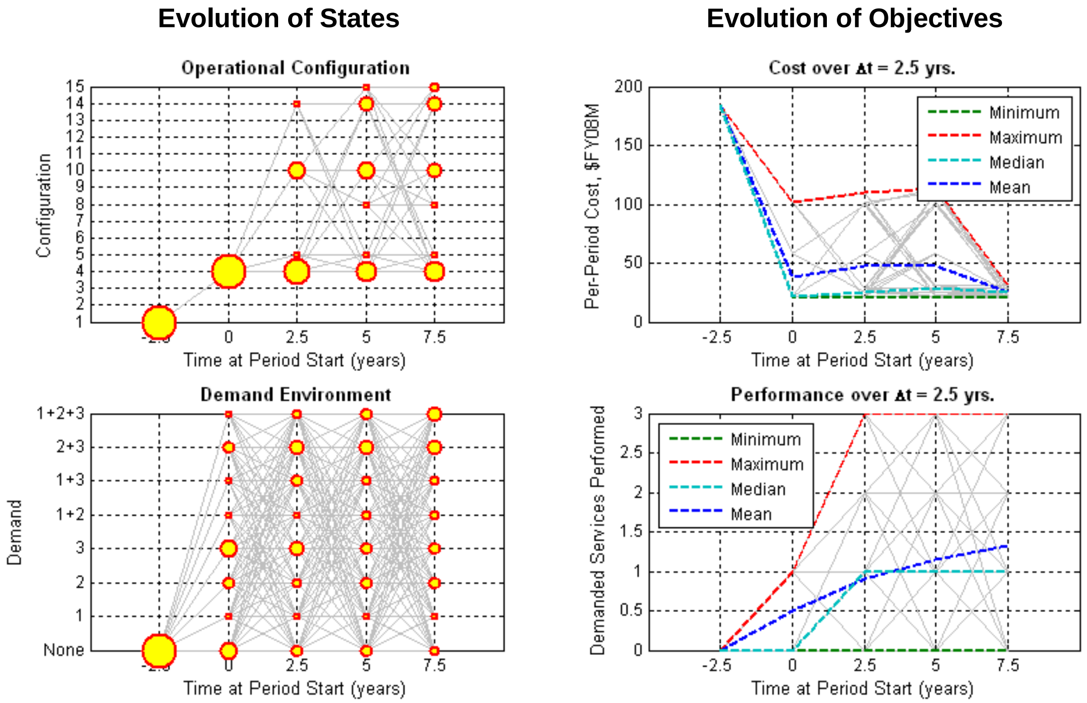

Figure 52. Evolution of states and objectives for an anticipatory policy (defined in Table 29). *In the plots on the left, the size of circles indicate the relative number of Monte Carlo simulation cases that exist in a given configuration or demand state (on the y-axes) at a given time (on the x-axes). The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation.*

Table 29. Anticipatory Policy. *Configuration decisions for a system in state s at time t are indicated by matrix elements shaded in gray.*

Current State, *s*			**Time at Period Start (years), *t***				
Total State	**Env.**	**Config.**	**-2.5**	**0**	**2.5**	**5**	**7.5**
1	None	1	4	4	4	4	1
2	None	2	4	4	4	4	1
3	None	3	4	4	4	4	1
4	None	4	4	4	4	4	1
5	None	5	4	4	4	4	1
6	None	6	6	6	6	6	1
7	None	7	7	7	7	7	1
8	None	8	4	4	4	4	1
9	None	9	4	4	4	4	1
10	None	10	4	4	4	4	1
11	None	11	11	11	11	11	1
12	None	12	7	7	7	7	1
13	None	13	6	6	6	6	1
14	None	14	4	4	4	4	1
15	None	15	4	4	4	4	1
16	1	1	4	4	4	4	1
17	1	2	4	4	4	4	1
18	1	3	4	4	4	4	1
19	1	4	4	4	4	4	1
20	1	5	4	4	4	4	1
21	1	6	6	6	6	6	1
22	1	7	7	7	7	7	1
23	1	8	4	4	4	4	1
24	1	9	4	4	4	4	1
25	1	10	4	4	4	4	1
26	1	11	11	11	11	11	1
27	1	12	7	7	7	7	1
28	1	13	6	6	6	6	1
29	1	14	4	4	4	4	1
30	1	15	4	4	4	4	1
31	2	1	7	7	7	7	1
32	2	2	7	7	7	7	1
33	2	3	10	10	10	10	1
34	2	4	10	10	10	10	1
35	2	5	14	14	14	14	1
36	2	6	13	13	13	13	1
37	2	7	7	7	7	7	1
38	2	8	10	10	10	10	1
39	2	9	10	10	10	10	1
40	2	10	10	10	10	10	1
41	2	11	11	11	11	11	1
42	2	12	7	7	7	7	1
43	2	13	13	13	13	13	1
44	2	14	14	14	14	14	1
45	2	15	10	10	10	10	1
46	3	1	4	4	4	4	1
47	3	2	4	4	4	4	1
48	3	3	4	4	4	4	1
49	3	4	4	4	4	4	1
50	3	5	4	4	4	4	1
51	3	6	6	6	6	6	1
52	3	7	7	7	7	7	1
53	3	8	4	4	4	4	1
54	3	9	4	4	4	4	1
55	3	10	4	4	4	4	1
56	3	11	11	11	11	11	1
57	3	12	7	7	7	7	1
58	3	13	6	6	6	6	1
59	3	14	4	4	4	4	1
60	3	15	4	4	4	4	1
61	1+2	1	5	5	5	5	1
62	1+2	2	8	8	8	8	1
63	1+2	3	8	8	8	8	1
64	1+2	4	5	5	5	5	1
65	1+2	5	5	5	5	5	1
66	1+2	6	13	13	13	13	1
67	1+2	7	12	12	12	12	1
68	1+2	8	8	8	8	8	1
69	1+2	9	8	8	8	8	1
70	1+2	10	8	8	8	8	1
71	1+2	11	11	11	11	11	1
72	1+2	12	12	12	12	12	1
73	1+2	13	13	13	13	13	1
74	1+2	14	5	5	5	5	1
75	1+2	15	8	8	8	8	1
76	1+3	1	11	11	11	11	1
77	1+3	2	12	12	12	12	1
78	1+3	3	13	13	13	13	1
79	1+3	4	14	14	14	14	1
80	1+3	5	14	14	14	14	1
81	1+3	6	13	13	13	13	1
82	1+3	7	12	12	12	12	1
83	1+3	8	15	15	15	15	1
84	1+3	9	15	15	15	15	1
85	1+3	10	15	15	15	15	1
86	1+3	11	11	11	11	11	1
87	1+3	12	12	12	12	12	1
88	1+3	13	13	13	13	13	1
89	1+3	14	14	14	14	14	1
90	1+3	15	15	15	15	15	1
91	2+3	1	7	7	7	7	1
92	2+3	2	7	7	7	7	1
93	2+3	3	10	10	10	10	1
94	2+3	4	10	10	10	10	1
95	2+3	5	14	14	14	14	1
96	2+3	6	13	13	13	13	1
97	2+3	7	7	7	7	7	1
98	2+3	8	10	10	10	10	1
99	2+3	9	10	10	10	10	1
100	2+3	10	10	10	10	10	1
101	2+3	11	11	11	11	11	1
102	2+3	12	7	7	7	7	1
103	2+3	13	13	13	13	13	1
104	2+3	14	14	14	14	14	1
105	2+3	15	10	10	10	10	1
106	1+2+3	1	11	11	11	11	1
107	1+2+3	2	12	12	12	12	1
108	1+2+3	3	13	13	13	13	1
109	1+2+3	4	14	14	14	14	1
110	1+2+3	5	14	14	14	14	1
111	1+2+3	6	13	13	13	13	1
112	1+2+3	7	12	12	12	12	1
113	1+2+3	8	15	15	15	15	1
114	1+2+3	9	15	15	15	15	1
115	1+2+3	10	15	15	15	15	1
116	1+2+3	11	11	11	11	11	1
117	1+2+3	12	12	12	12	12	1
118	1+2+3	13	13	13	13	13	1
119	1+2+3	14	14	14	14	14	1
120	1+2+3	15	15	15	15	15	1

5.5. Step 5: Implications for Initial System Selection

The analysis of Step 4 has produced a large set of data on optimal paths and policies to follow for the *entire* system time horizon, and it is easy to lose track of the implications this has for the *initial* system decision. This final step of the framework builds upon the analysis results of Step 4 to provide implications for this decision.

5.5.1. Implications based on the Expected-Value Pareto Frontier

In the case of a path, the initial decision is simply the first configuration in its associated configuration sequence. In the case of a policy, the initial decision is found by locating the initial condition in the row of the policy matrix (in this distributed-payload satellite application, at total state 1, which corresponds to the "nothing" configuration fielded and no services demanded) and examining the element in the first column (in this case, the t = -2.5 year column). To facilitate this, the initial configurations specified by the Pareto-optimal paths and policies found in Figure 47 and Figure 50 are identified in Figure 53. In this figure, the Pareto frontier solutions of Figure 47 and Figure 50 are identified by their expected total cost on the x-axis. On the y-axis are the initial configuration decisions called for by each Pareto-optimal path (yellow circles) or policy (blue squares). Two particular observations can be made: First, only three configurations (Configs. 1, 3, and 11) appear among the optimal initial decisions. All paths and policies with other initial decisions are dominated by paths and policies using these three configurations. Second, the size of the initial configuration tends to increase as the expected total cost of the system increases. For example, only the "Nothing" configuration (Config. 1) appears as an optimal initial decision for total expected budgets under $195 million; these solutions tend to be either policies that wait until sufficient demand materializes to justify the expenditure of funds or paths that tend to delay initial operational capability until demand evolves substantially beyond the initial "None" environment. At the highest expected total cost is the decision to initially develop the three-payload monolith (Config. 11), which is the least expensive method to ensure complete capture of all possible future demand for services.

Also noted next to several paths and policies in Figure 53 are the number of transitions Φ available from each initial configuration (1, 3, or 11) for the average per-period cost associated with each total cost. As discussed in Step 1, this number Φ is an indicator of flexibility, and it can be seen that more flexible initial configurations (Φ = 2 or Φ = 3) are selected at higher cost and performance preferences. Thus, there exists some correlation between flexibility and performance. However, the maximum-performance (and maximum-cost) Config. 11 initial decision is far from the most flexible for its average $60 million per-period budget; Figure 44 illustrates that the fully-fractionated three-payload configuration (Config. 15) has Φ = 8 transitions available for the same budget. Thus, this example illustrates that maximization of performance does not necessarily translate into maximization of the flexibility of a system's configuration.

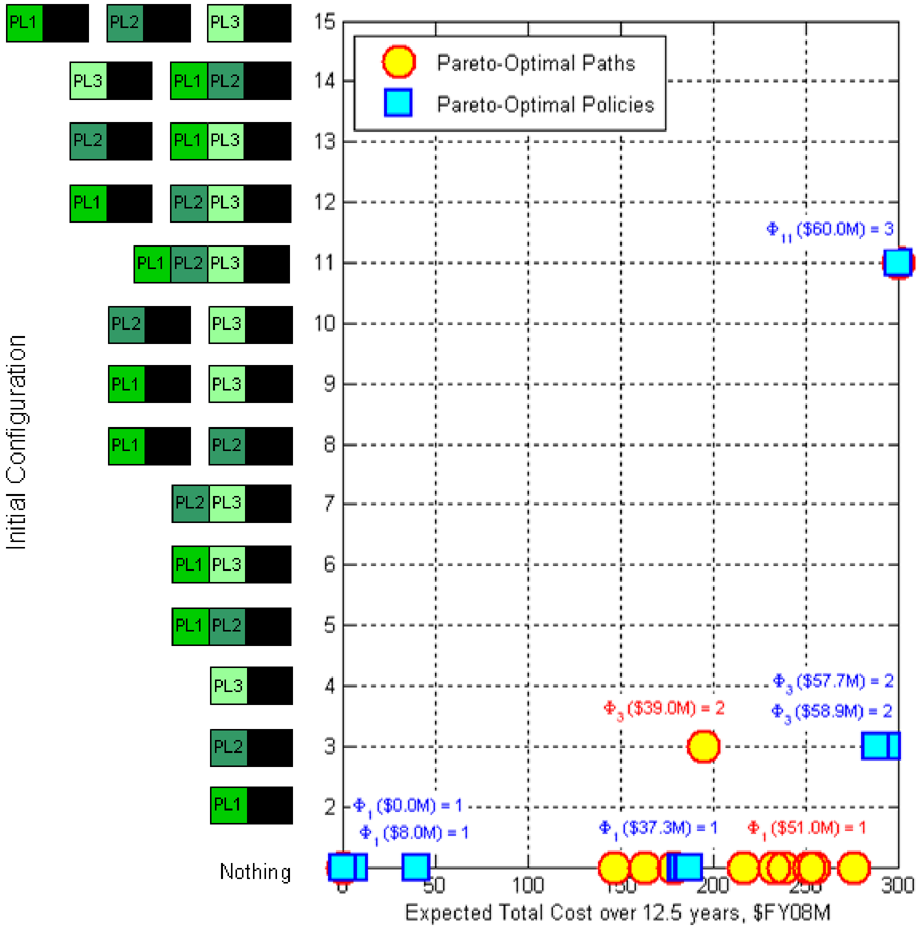

Figure 53. Initial configurations for Pareto-optimal paths and policies as a function of expected path or policy total cost. *Also noted are the numbers of transitions available for several initial configurations at their path or policy's average per-period budget requirements.*

5.5.2. Accounting for Non-Expected-Value Objectives

A final relevant consideration for initial system selection is the fact that expected-value objective functions for the cumulative cost and performance metrics may not fully capture a decision-maker's true objectives. Use of these expected-value objectives enables the use of MDP dynamic programming techniques to efficiently explore the astronomically large policy trade-space; however, in the case of one-of-a-kind satellite programs a decision-maker may also be interested in minimizing risks associated with a given expected level of cost or performance.

Operating under the assumption that the expected-value optima discovered in Step 4 are reasonable initial guesses for desirable policies, a multi-objective genetic algorithm may be employed to perturb each of the policies identified in Figure 50, simulate each new hybrid policy, and search for non-dominated solutions in terms of any combination of metrics that can be accounted for via simulation. The results of Figure 54 are produced by applying this technique to the new metrics of 90th percentile (near-worst-case) total cost and 10th percentile (near-worst-case) total number of demanded services performed, in addition to the expected-value versions of these metrics. Of particular note in the Figure 54 multivariate plot are four subplots: First, the data in the subplot of the second row and first column shows the familiar expected-value cost and performance trade, with slightly better Pareto frontier performance due to the genetic algorithm's search. Second, the data in the subplot of the last row and second column shows the 10th percentile performance vs. the 90th percentile cost; the performance data in this subplot is noticeably more discrete since fractional numbers of services performed are not possible in a simulation. Finally, the upper left and bottom right subplots show the correlations between the new percentile-based metrics and their expected-value counterparts. In the cases of both subplots, linear correlation is quite strong ($R^2 = 0.85$ and 0.88) and supports the use of expected value as a surrogate for optimizing the percentile-based metrics.

Also of note in Figure 54 is that each data point, which represents a particular policy result, has a color that corresponds to the initial configuration decision implied by its associated policy. Of particular note is that these initial decisions differ little from those implied by the original path- and MDP-policy-based results in Figure 53. Use of Config. 1 initially is still associated with low cost and performance; use of Config. 3 is associated with medium values for both objectives; and Config. 11 is associated with the highest levels of cost and performance. The primary difference is the introduction of Config. 13 as an initial decisions, which has performance and cost levels that are generally competitive with Config. 11.

The usefulness of the multivariate plot of Figure 54 becomes more evident if cost or performance constraints are imposed by the decision-maker. For example, suppose that this decision-maker has a $500 million limit on the funds available for supporting this system over its time horizon. If the decision-maker wishes to be 90% sure that this budget will not be breached, a $500 million constraint may be imposed on the 90th percentile total cost metric. This constraint eliminates many high-cost (and also high-performance) options that formerly fell into the high 90th percentile cost regions of the multivariate plot that are now gray in Figure 55. Similarly, the decision-maker may wish to have 90% confidence that more than one service will be performed over the system's lifetime. In this case an additional constraint may be imposed, represented by the horizontal gray stripe in the subplots of the last row in Figure 55. Combined, these two constraints eliminate a large number of the policy options available. As Figure 55 shows, no policy options remain for which the "Nothing" configuration is acceptable. Furthermore, in both the expected-value-based and percentile-based performance vs. cost subplots, policies involving the three-payload monolith (Config. 11) as an initial configuration exhibit lower cost for the same (or better) performance as those that involve Config. 13. As a result, the decision is narrowed to one of whether to select a policy that suggests Config. 11 as an initial decision (at an expected and 90th percentile total cost of $300 million, with 5.6 expected services performed and 3 services performed in the 10th percentile) or, instead, Config. 3 (at an expected $285 million and 90th percentile $331 million total cost, with 3.8 expected services performed and 2 services performed in the 10th percentile). While no objectively correct decision exists, it is likely that the small ($15 million, or 5%) difference in expected cost and large (1.8 services, or 38%) difference in performance between the options would compel many decision-makers to accept the slightly higher budget for such a significant performance increase.

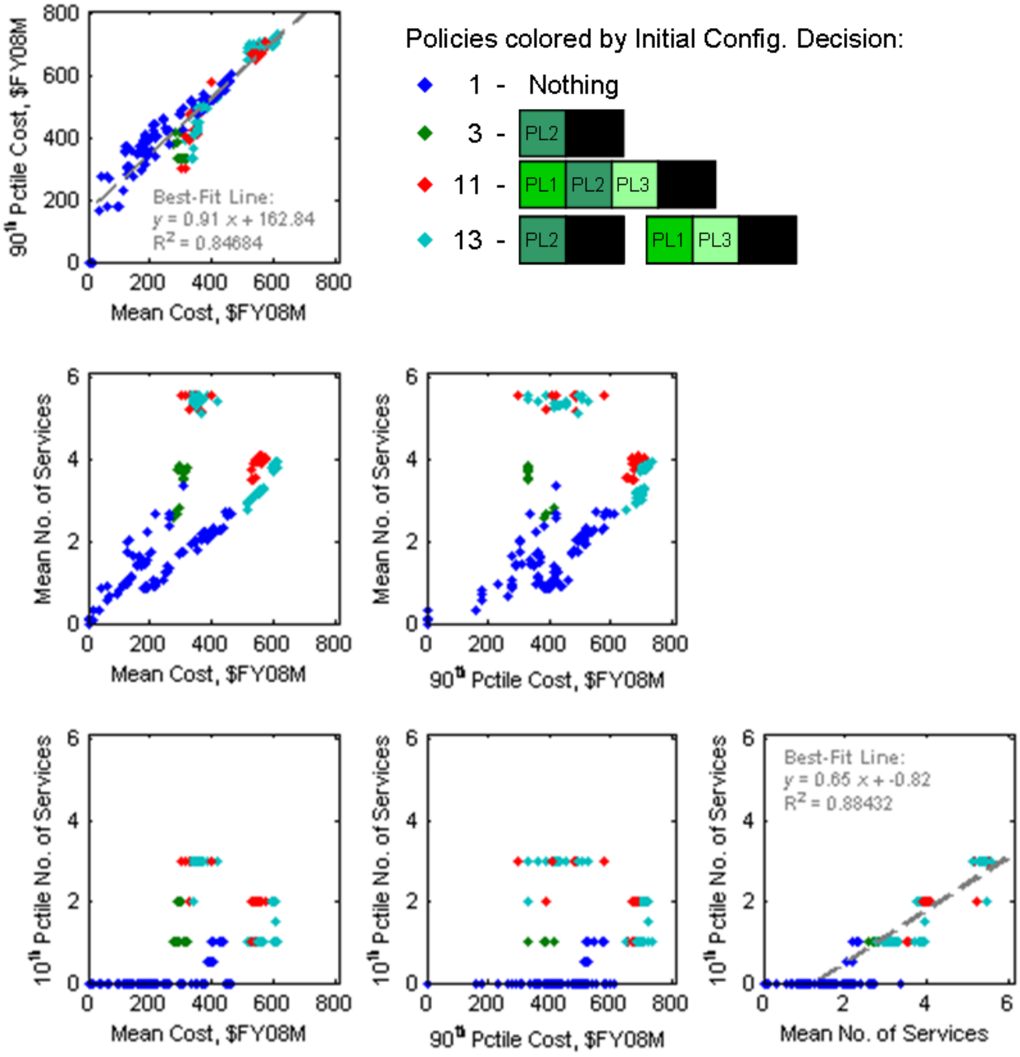

Figure 54. Multivariate plot of multi-objective genetic algorithm policy results. *Each data point indicates the performance of one policy result in terms of the four percentile-based and expected-value metrics of interest. Data points are colored by their corresponding policy's initial configuration decision.*

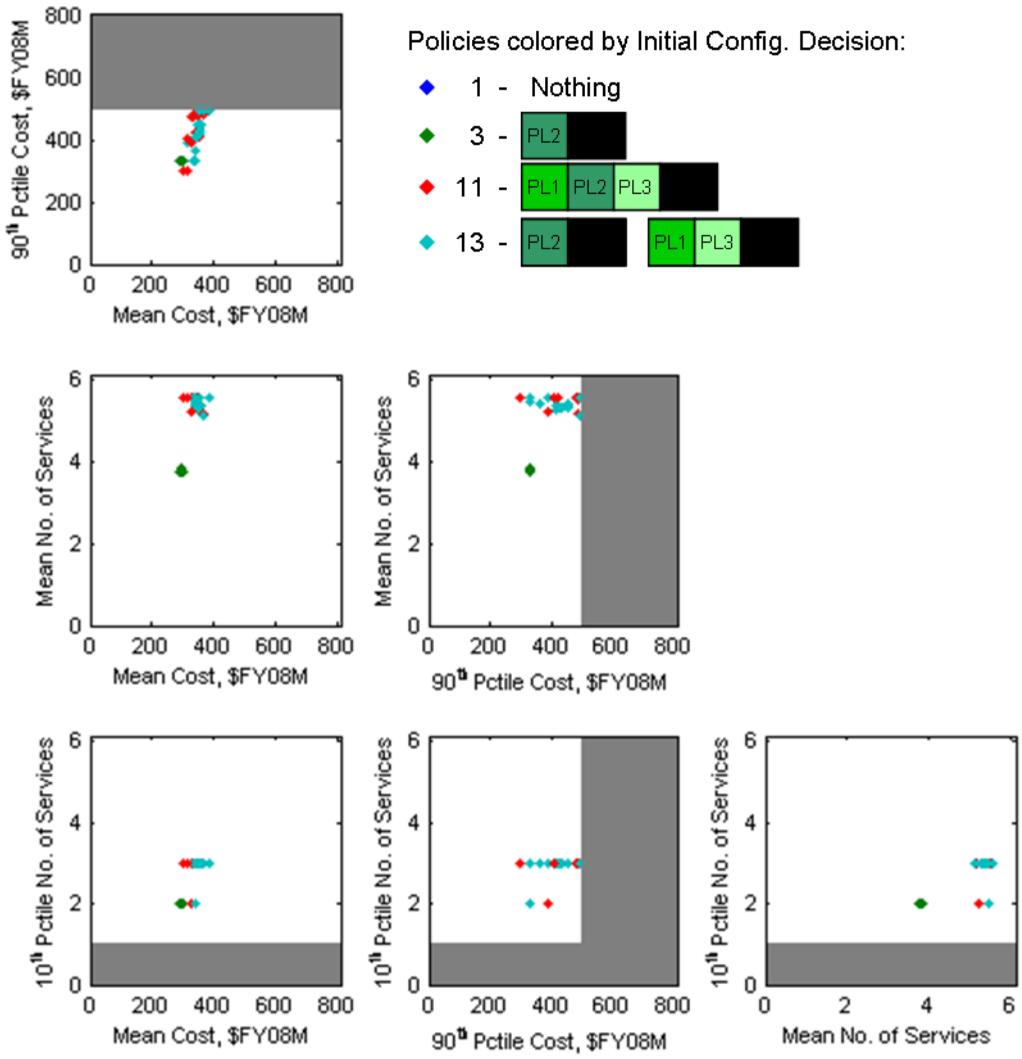

Figure 55. Multivariate plot of multi-objective genetic algorithm policy results with cost and performance constraints imposed. *Each data point indicates the performance of one policy result in terms of the four percentile-based and expected-value metrics of interest. Data points are colored by their corresponding policy's initial configuration decision. Gray areas indicate regions of the space eliminated due to cost and performance constraints.*

5.6. Summary

First and foremost, this chapter has demonstrated how the theoretical framework posed in Chapter 4 can be applied to a class of problem directly relevant to the space industry today. This chapter began with definition of a problem in which a hypothetical Department of Defense decision-maker was faced with a decision about what combination of payloads to launch upon potentially multiple distributed, free-flying satellites. Population and analysis of cost transition matrices revealed information about the flexibility of various configuration options. For example, the three-payload fully-fractionated configuration (Config. 15) has significantly more next-period transition options than any other configuration just above a 30-month budget of $31 million and retains a high number of available transitions for even higher budgets; below this budget it has no options because of its high operations costs. The three-payload monolith, on the other hand, tends to have fewer available transitions than most other configurations at most budget levels since it has no modules in common with other configurations.

Population of a Markov chain representing the evolution of the demand for payload services and population of a performance matrix representing the number of demanded services performed by a given configuration in a given demand environment enabled Steps 4A and 4B of the framework to find Pareto-optimal decision paths and policies. Step 4A returned the interesting result that just 12 of the 50,625 possible paths for this four-period decision problem were Pareto-optimal, and only four configurations (Configs. 3, 5, 7, and 11) appeared within these paths. Perhaps more interesting was the identification of a path with a one-period delay followed by the fielding of the three-payload monolith that dominated a strategy of incremental capability buildup. Step 4B illustrated how Markov decision process techniques were able to efficiently find a Pareto frontier of policies more intricate than a simple path. Illustrated was one case in which the MDP solution procedure found an optimal compromise between maximum performance and minimum cost by identifying a policy which only developed the three-payload monolith if an appropriate level of demand for particular payloads materialized early during the program. Furthermore, the MDP solution procedure was shown to identify policies that dominated an anticipatory policy that a human might have proposed as a reasonable policy.

In the application of Step 5 of this thesis' framework to the distributed- or multi-payload satellite problem, initial decisions were objectively narrowed to just four configurations: Configs. 1, 3, 11, and 13. Imposing 90th percentile budget and 10th percentile performance constraints narrowed the list down to three, and the existing trades would likely compel most decision-makers toward selection of Config. 11, the three-payload monolith, as the initial configuration. This selection is an interesting result, particularly since Config. 11 is one of the least flexible options as identified in Step 1.

Notably, the selection of this relatively inflexible configuration is not contradictory to this thesis' framework. As emphasized in Section 2.5, it is a tenet of this work that a decision-maker cares about flexibility principally because of cost and performance benefits it may enable in the future. Thus, this example highlights the fact that finding a minimum-cost, maximum-performance solution in a changing demand environment may not be equivalent to finding a solution with maximum flexibility. However, until the proper analysis and optimization is run to account for the ability of the system to change over time, this equivalence cannot be known. In fact, the particular result that favors the monolith can only be said to hold for the numerical inputs assumed in this chapter. Future investigations are encouraged to modify these inputs to explore under what circumstances monolithic and fractionated spacecraft are favored as solutions.

CHAPTER 6

ADVANCED APPLICATION: NASA HUMAN EXPLORATION ARCHITECTURE DECISION-MAKING

Richard Bellman prefaced his original 1957 book on dynamic programming [93] by observing that “in modern life, in economic, industrial, scientific and even political spheres, we are continually surrounded by multi-stage decision processes. Some of these we treat on the basis of experience, some we resolve by rule-of-thumb, and some are too complex for anything but an educated guess and a prayer.”

While realistic and reasonably complex, the space system planning applications in Chapters 4-5 have been intended mainly to demonstrate the core concepts of this thesis’ framework. In the present chapter, a major current systems planning challenge within the National Aeronautics and Space Administration is selected to illustrate the applicability and utility of this new framework for problems that might otherwise be well beyond the complexity threshold for even an educated guess or a prayer. This particular example will illustrate the framework for inter-mission flexibility applications [28].

As described in Chapter 1, the Review of U.S. Human Spaceflight Plans Committee (Augustine Committee) was formed in 2009 to assess the status and direction of NASA’s human spaceflight program. One of the viable exploration programs the committee proposed was a “flexible path” involving the development of systems to enable mission options for a variety of inner solar system destinations, including the lunar vicinity, Earth-Moon and Sun-Earth Lagrange points, near-Earth objects, Mars vicinity, and the moons of Mars. This new approach is distinct in that, instead of focusing on a single path to a single destination, it focuses on providing options to allow the human space program to adapt to changing expectations or demands as exploration progresses. Since 2009, the Flexible Path strategy has made a substantial impact and has been largely adopted in the formulation of architecture plans within NASA.

However, due perhaps in part to time constraints, the Augustine Committee's justification for naming this approach the Flexible Path was qualitative. No attempt appears to have been made to quantitatively compare the flexibility of the Flexible Path approach to the flexibility of the other approaches (Mars First and Moon First) that the committee presented. As demands upon NASA evolve over the coming few decades, the question remains: What architectural components should the agency develop today so that, in the long run, it is able to minimize the total cost of the human spaceflight program and maximize total return in an environment of changing mission expectations?

The aim of the present chapter is twofold: First, with the help of probability, schedule, and cost estimates obtained through extensive interaction with NASA personnel, it is intended that this advanced application will shed light on and inform decision-making for the present NASA architecture decision-making challenge. Second, in the process of realistically addressing this challenge, several advances to this thesis' methodology are introduced:

- ***Large Configuration State Space.*** In the previous applications in Chapters 4-5, the configuration state space consisted of on the order of ten candidate configurations. This NASA application illustrates how the present framework, when supplied appropriate computing power, can be used even for state spaces with thousands of configuration options. In order to handle this large state space, a computer code is developed to automatically create the necessary cost transition matrices based on the component-by-component configuration definitions. All relevant costs are accounted for, including development and first unit, production, mission and ground operations, program management and systems engineering, and program termination and shutdown costs.

- ***Intermediate Development Architectures, Operations Architectures, and Memory Architectures.*** In the previous applications, it was assumed that the demand environment would be relatively constant during the development of a new architecture. Further, it was taken for granted that the configuration developed during the one period would always become the configuration fielded (or operated) in the subsequent fielded. While convenient for demonstrating the fundamentals of this thesis' framework, these assumptions are not required. In the NASA example, the very real possibility of demand changing mid-development is modeled, as are the options to cancel a program in mid-development and to not utilize all components of an architecture just developed. This is accomplished by introducing intermediate architectures representing systems that are not operational but which are partially developed. The introduction of these intermediate architectures subsequently requires defining the configuration state by three elements: the development architecture, operations architecture, and memory architecture.

- ***Configuration-Dependent Demand.*** The previous applications have effectively assumed that the environment described by the mission demand Markov chain evolves independently of the configuration that the decision-maker selects to respond to this demand. In some situations, such as if there exist many actors responding to the same environment (as in a scenario of perfect competition in economics), this independence may be a realistic approximation. If there exist relatively few actors (in economics, for example, an oligarchy), it may be more appropriate to consider the influences of decision-maker choices on future demand. NASA falls into this latter category, as the demands that are placed upon it are at least partially dependent on whether it is achieving its currently demanded mission. The capability to handle this interaction between configuration and environment is demonstrated here.

- ***Elicitation of Expert-Opinion Markov Chain Probabilities.*** In the absence of sufficient historical data for inner solar system destination demand, this application takes the step of eliciting expert opinion for the probabilities of demand transition. To permit this data to be extensible to analyses of different timesteps, the data is elicited as a continuous-time Markov chain and then discretized to the proper step. Expert opinion also contributes to the selection of the mission return (or performance) figure of merit.

- ***Endogenous Schedule-Slippage Uncertainties.*** To account for the existence of endogenous uncertainty unrelated to exogenous mission demand changes, a basic probabilistic model is incorporated to model the probability of development program schedule slippage.

- ***Exploration of High-Performing Policies in terms of Non-Cumulative as well as Non-Expected-Value Objective Functions.*** In the final step of previous applications, exploration of non-expected-value objective functions focused on measures of dispersion of the original cumulative objective functions. In the NASA application, exploration among non-cumulative objectives is included.

Recall that this thesis' framework consists of five basic steps, outlined in Figure 12. First, system configuration options are identified and costs of switching from one configuration to another are compiled into a cost transition matrix. Second, probabilities that demand on the system will transition from one mission to another are compiled into a mission demand Markov chain. Third, one performance matrix for each design objective is populated to describe how well the identified system configurations perform in each of the identified mission demand environments. Fourth, possible future sequences of system configurations are simulated and sequences that are Pareto-optimal in terms of the decision-maker's objectives are identified. In a complementary approach, the system decision problem is formulated as a multi-objective variant of a Markov decision process, and Pareto-optimal decision policies are identified. It is worth noting that, due to the large configuration state space of the NASA problem, the traditional Step 4A full-factorial search paths becomes infeasible and the Markov decision process approach will be used exclusively. Finally, the paths and policies from the latter step are synthesized into a set of data to inform initial system selection.

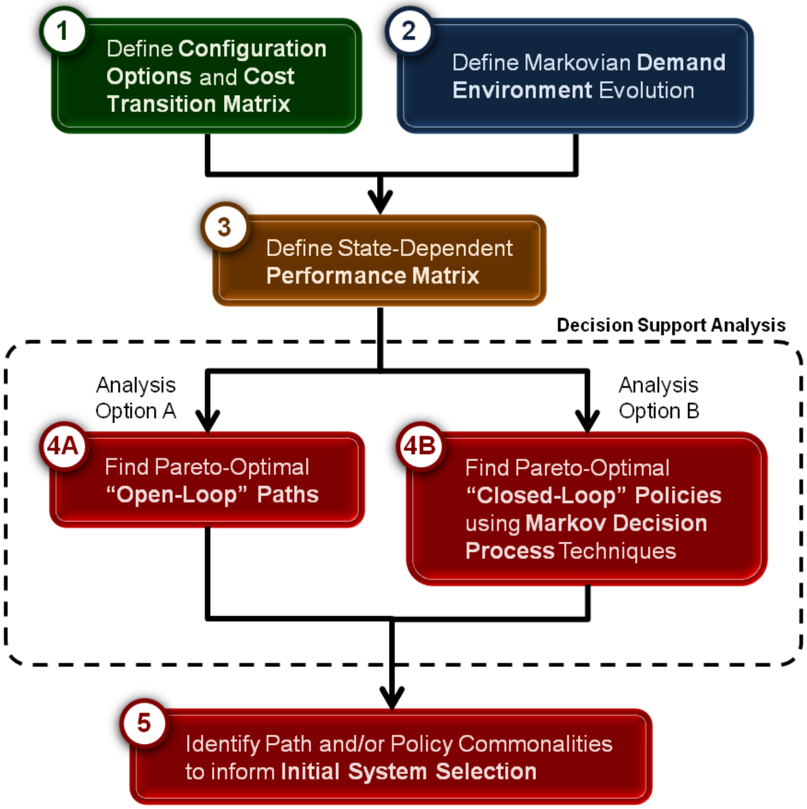

1
Define Configuration Options and Cost Transition Matrix
2
Define Markovian Demand Environment Evolution
3
Define State-Dependent Performance Matrix
Decision Support Analysis
Analysis Option A
Analysis Option B
4A
Find Pareto-Optimal "Open-Loop" Paths
4B
Find Pareto-Optimal "Closed-Loop" Policies using Markov Decision Process Techniques
5
Identify Path and/or Policy Commonalities to inform Initial System Selection

Figure 12. Five major steps of this thesis' framework.

6.1. Step 1: Define Configuration Options and the Cost Transition Matrix

As in the previous examples, the fundamental question for the first step of this framework is: What are the relevant “positions”, or the configurations, that the system in question can take? Subsequently, what are the costs of transitioning between any two of these configurations?

It is also important to define up-front the duration of the period (or time step) of interest. For the human spaceflight architecture decision problem, this period will be set at two years in duration. This duration corresponds to the U.S. Congressional election cycle as well as one-half the length of the U.S. presidential election cycle and historically one-half the median time between appointments of new NASA administrators [105]. The implication of this period selection is an assumption that decision points, at which architecture selections are either re-confirmed or changed in response to mission demand changes, occur once per new Congress and twice per Presidential and, on average, NASA administration.

6.1.1. Defining the Configuration Space

As suggested in Chapter 4, the specific system configurations relevant for this problem of interest originate from the definition of a morphological matrix as shown in Table 30. Each row denotes a particular defining attribute of an architecture, which in all cases is the number of architecture components (e.g., launch vehicles, crew vehicles, in-space propulsion stages, landers, and other systems) that will be developed and produced during a particular two-year increment. Possible values for each attribute are listed as options. For example, the second row indicates that an engineer might consider fielding architectures that require zero, three, four, six, eight, ten, or twelve heavy-lift launch vehicles over a two-year period.* Since an architecture is defined once one attribute

---

* In theory, any integer value for this attribute is acceptable; for brevity, the morphological matrix shown here only lists the values that will be used in later architecture definitions.

value is selected from each row, it can be easily computed that in this simple morphological matrix there exist over 6.58 billion possible architectures. However, some of these architectures make little practical sense. For example, one option within these 6.58 billion is to develop and produce 12 heavy-lift launch vehicles (HLVs) and no other systems or capabilities (i.e., select zero for all attributes other than HLVs); such an architecture would not be able to achieve any missions asked of it since no in-space elements or capabilities exist.

Table 30. Morphological Matrix for the Human Spaceflight Architecture Application.

Attribute	Relevant Options						
No. of Crew Launch Vehicles (CLVs)	0	4					
No. of Heavy Lift Launch Vehicles (HLVs)	0	3	4	6	8	10	12
No. of Commercial Cargo Launch Vehicles (CCLVs)	0	6					
No. of Multi-Purpose Crew Vehicles (MPCVs)	0	1	2	4			
No. of Commercial Cargo Logistics Modules (CCLMs)	0	6					
No. of Small Chemical Stages	0	1	2				
No. of Medium Chemical Stages	0	1	2	3	4	6	8
No. of Large Chemical Stages	0	2	4	6	8		
No. of Deep-Space Habitation Modules	0	1	2				
No. of Lunar Landers	0	8					
No. of Mars Landers	0	2					
No. of Multi-Mission Pressurized Rovers	0	1	2	8			
No. of Unpressurized Rovers	0	2	4				
No. of Science Rovers	0	2	4				
No. of Surface Habitats	0	1					
No. of Logistics Modules	0	2	4				
No. of Power Generation and Storage Units	0	1	2				
No. of ISRU Systems	0	2					
No. of Surface Extravehicular Activity (EVA) Suits	0	10	20				
No. of In-Space Extravehicular Activity (EVA) Suits	0	8	12				
No. of Supporting Communications/Navigation Satellites	0	1					

One way to overcome this limitation as well as restrict the architectures considered to a manageable number is to use the morphological matrix to assist in brainstorming themed configuration options [77]. In this application, a reasonable theme to select is the architecture's intended mission destination, which drives mission duration and spacecraft velocity change (ΔV) requirements that subsequently suggest certain numbers of stages and launch vehicles as well as habitat and excursion vehicles. With the assistance of NASA Johnson Space Center personnel, the ten themed architectures defined in Table 31 are selected, based largely on recent studies of the agency-wide Human Exploration Framework Team (HEFT) and Human Spaceflight Architecture Team (HAT). Note that eight of these architectures are directly themed upon destinations suggested by the Augustine Committee, and Architectures 1 and 10 in some respects bound the configuration space: Architecture 1 is the option to develop nothing, and the "Deep Space" themed Architecture 10 is the option to develop and produce the maximum of the number of components specified by Architectures 3, 4, 6, 7, and 8. Also note that, to avoid confusion with the demand environments in Step 2 and beyond, these architectures may in general be referred to by their architecture numbers (1-10). Finally, the reader may note in Table 31 that four engines are listed as architectural components but are not parts of the morphological matrix in Table 30. These numbers of engines are not user-defined but are directly dependent other components in an architecture; for example, there exist five RS-68-class engines for each heavy-lift launch vehicle.

Table 31. Architecture Definitions for the Human Space Exploration Application.

	Architecture: Number, Theme, and Icon									
	-1- Nothing	-2- LEO	-3- GEO Servicing	-4- Lunar Orbit	-5- Lunar Surface	-6- Sun-Earth L2	-7- Near-Earth Object	-8- Mars Moon	-9- Mars Surface	-10- Deep Space
Architecture Components	▭		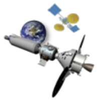					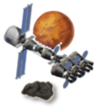	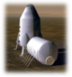	
1. Crew Launch Vehicles (CLVs)	-	4	-	-	-	-	-	-	-	-
2. Heavy Lift Launch Vehicles (HLVs)	-	-	3	4	12	8	6	4	10	8
3. Commercial Cargo Launch Vehicles (CCLVs)	-	6	-	-	-	-	-	-	-	-
4. Multi-Purpose Crew Vehicles (MPCVs)	-	4	2	4	4	2	2	1	1	4
5. Commercial Cargo Logistics Modules (CCLMs)	-	6	-	-	-	-	-	-	-	-
6. Small Chemical Stages	-	-	-	-	-	2	-	-	1	2
7. Medium Chemical Stages	-	-	3	4	8	4	6	2	1	6
8. Large Chemical Stages	-	-	-	-	-	-	-	2	6	4
9. Deep-Space Habitation Modules	-	-	1	-	-	2	2	1	1	2
10. Lunar Landers	-	-	-	-	8	-	-	-	-	-
11. Mars Landers	-	-	-	-	-	-	-	-	2	-
12. Multi-Mission Pressurized Rovers	-	-	1	-	8	2	2	2	2	2
13. Unpressurized Rovers	-	-	-	-	4	-	-	-	2	-
14. Science Rovers	-	-	-	-	4	-	-	-	2	-
15. Surface Habitats	-	-	-	-	-	-	-	-	1	-
16. Logistics Modules	-	-	2	-	4	2	2	-	-	2
17. Power Generation and Storage Units	-	-	-	-	2	-	-	-	1	-
18. ISRU Systems	-	-	-	-	2	-	-	-	2	-
19. Surface Extravehicular Activity (EVA) Suits	-	-	-	-	20	-	-	-	10	-
20. In-Space Extravehicular Activity (EVA) Suits	-	8	8	8	-	8	8	12	-	12
21. Supporting Communications/Navigation Satellites	-	-	-	-	1	-	-	-	1	-
22. RS-68-Class Engine	-	-	15	20	60	40	30	20	50	40
23. J-2X-Class Engine	-	4	6	8	24	16	12	8	20	16
24. RL-10B-2-Class Engine	-	-	6	8	56	12	12	14	44	36
25. AJ-10-Class Engine	-	4	2	4	4	2	2	1	1	4

Can this architecture space be visualized? In this application, each architecture in Table 31 is described by a vector of 25 variables and can be accurately visualized only in a 25-dimensional hyperspace. In this hyperspace, similar architectures would be separated by smaller distances and dissimilar architectures would be separated by larger distances. While it is not possible to display these 25 dimensions graphically on a two-dimensional page, it is possible to preserve much of this distance property by solving for two-dimensional coordinates for which pairwise Euclidean distances minimize the sum of the squared errors with the true 25-dimensional Euclidean distances. This is expressed in Eq. (30), where x_i and y_i indicate the abscissa and ordinate, respectively, of architecture i in the two-dimensional Euclidean space, A_i indicates the true coordinates of architecture i in the 25-dimensional space, and λ is a free scaling factor.

$$\begin{aligned} &\min_{\hat{x},\hat{y},\lambda} \sum_{\substack{\forall i \\ j\le i}} \left(\left\| \begin{bmatrix} \hat{x}_j \\ \hat{y}_j \end{bmatrix} - \begin{bmatrix} \hat{x}_i \\ \hat{y}_i \end{bmatrix} \right\| - \lambda \left\| \vec{A}_j - \vec{A}_i \right\| \right)^2 \\ &\text{subject to} \quad 0 \le \hat{x}_i \le 1, \forall i \\ &\qquad\qquad\quad\; 0 \le \hat{y}_i \le 1, \forall i \end{aligned} \tag{30}$$

For the architectures in Table 31, the two-dimensional architecture space in Figure 56, with architecture coordinates listed in Table 32, results.[*] These coordinates make some intuitive sense: Note that architectures 3, 4, 6, 7, 8, and 10, which are characterized by similar beyond-LEO but non-surface destination requirements, are grouped together. Architectures 1 and 2 (the “Nothing” and “LEO” themed architectures) are somewhat separated and toward the right of the graphic, while the architectures intended for lunar and Mars surface destinations are significantly separated toward the upper left of the graphic.

---

[*] Minor adjustments have been made to the true minimum solution to prevent state circles from overlapping. The set of solution coordinates have also been normalized to span the range of zero to unity.

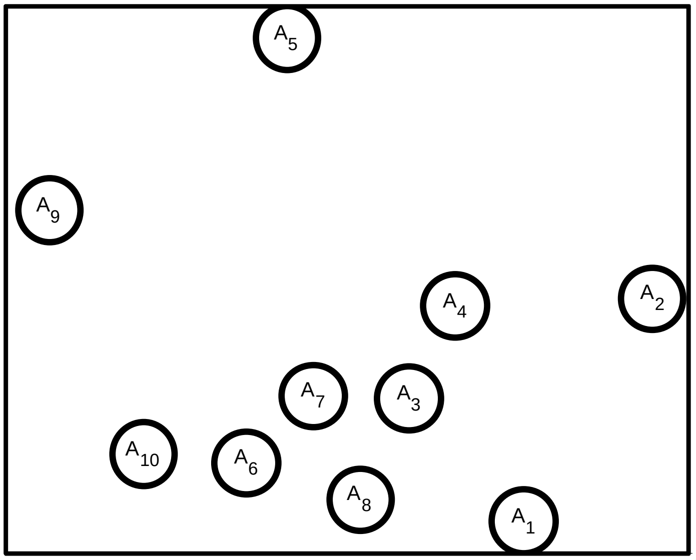

Figure 56. Architectural Projection for the Human Space Exploration Application

Table 32. Architectural Projection Two-Dimensional Coordinates.

Arch.	*x*-coordinate (abscissa)	*y*-coordinate (ordinate)
1	0.7871	0.0000
2	1.0000	0.3686
3	0.5962	0.2025
4	0.6728	0.3563
5	0.3942	0.8000
6	0.3271	0.0963
7	0.4385	0.2073
8	0.5156	0.0353
9	0.0000	0.5158
10	0.1560	0.1097

At this point, it may be noted that the word “architecture” has been used instead of “configuration” when describing the ten sets of components in Table 31. The primary reason for this is that the two-year timestep of the present application requires a distinction between an architecture and the full decision available to the decision-maker at any point in time (i.e., the “configuration”): Because development will span multiple two-year time periods,[*] it is inappropriate to assume the decision-maker will field the

---

[*] In this application, the development of each architecture (except for the “Nothing” architecture) is approximated as nominally taking eight years (four periods). This is in agreement with timelines for crew exploration vehicle, crew launch vehicle, lunar lander, heavy-lift launch

previously-developed architecture in the current period since in many cases development of that architecture is not yet finished. As a result, at any given decision point the decision-maker must choose what architecture to operate in addition to what architecture to develop or continue to develop. Furthermore, in a later development period it may be necessary for the decision-maker to keep in memory an architecture that had been current when development had started, since this would affect whether certain development costs need to be incurred.

Thus, this {Development, Operations, Memory} architecture triplet will define a configuration for the purposes of this application. A notional four-period sequence of configurations (architecture triplets) that a decision-maker could choose is depicted in Figure 57. Reading this sequence from left to right, in the first two-year period development for the near-Earth object (NEO) themed architecture is begun while operation of the low-Earth-orbit (LEO) themed architecture continues and the capabilities of a previous lunar surface themed architecture are in memory. In the second period, NEO architecture development continues into its second phase, as does operation of the LEO architecture. The third and fourth periods see continuation of the NEO architecture

---

vehicle, Earth departure stage, and lunar surface systems development from the 2005 ESAS report, which were there all baselined on 7-9 year schedules [29]. This is justified historically, for example, by the Space Shuttle Orbiter, for which authority to proceed was obtained in August 1972 and for which mating to its Solid Rocket Boosters and External Tank in preparation for its first flight occurred in November 1980, just over eight years later [106]. Even on the accelerated Apollo program, the time between selection of North American Aviation as the prime contractor for the Apollo Command and Service Module in November 1961 [107] and the first successful manned flight of the program in October 1968 was nearly seven years. (It deserves note that Apollo Spacecraft 012, assigned to the crew of Apollo 1, was received at Kennedy Space Center in August 1966 [108]. However, the 113 engineering orders not accomplished at the time of delivery and 623 engineering changes ordered subsequent to delivery [108], in addition to the fire that took the lives of astronauts Grissom, White, and Chaffee, suggest that the five-year development implied by the 1966 delivery date would not be appropriate.)

development to completion but discontinuation of LEO operation (with a substitution of operation of the Nothing architecture). Throughout all periods, the lunar surface architecture capability is retained in memory.

Two fundamental observations can be made from the example sequence of configurations in Figure 57: First, each configuration (or architecture triplet) represents a decision over which an appropriate decision-maker is assumed to have direct influence. It is within the decision-maker's prerogative whether to continue to develop a current development architecture or to develop a new architecture instead[*]; it is within the decision-maker's prerogative whether to continue or discontinue operations of an architecture; and it is within the decision-maker's prerogative to hold or not hold within institutional memory the capabilities associated with a previous architecture. Second, there are *many* possible configurations, even for only the ten architectures defined in Table 31. Considering that nine of these architectures have four-period developments (the exception is Architecture 1), there are in theory (9×4 + 1) × 10 × 10 = 3,700 configurations. Each may be assigned an identification number, as are the four configurations in Figure 57. In practice, however, there are somewhat fewer than 3,700 relevant configurations – only 3,286 – because of some practical considerations concerning relevant and allowable states and transitions detailed next.

---

[*] This highlights the present assumption that only one architecture is assumed to be under development at a time. If a future analyst wishes to apply this technique to options in which multiple architectures can be under development at once, he or she need only to define a new architecture for each multi-architecture option under consideration.

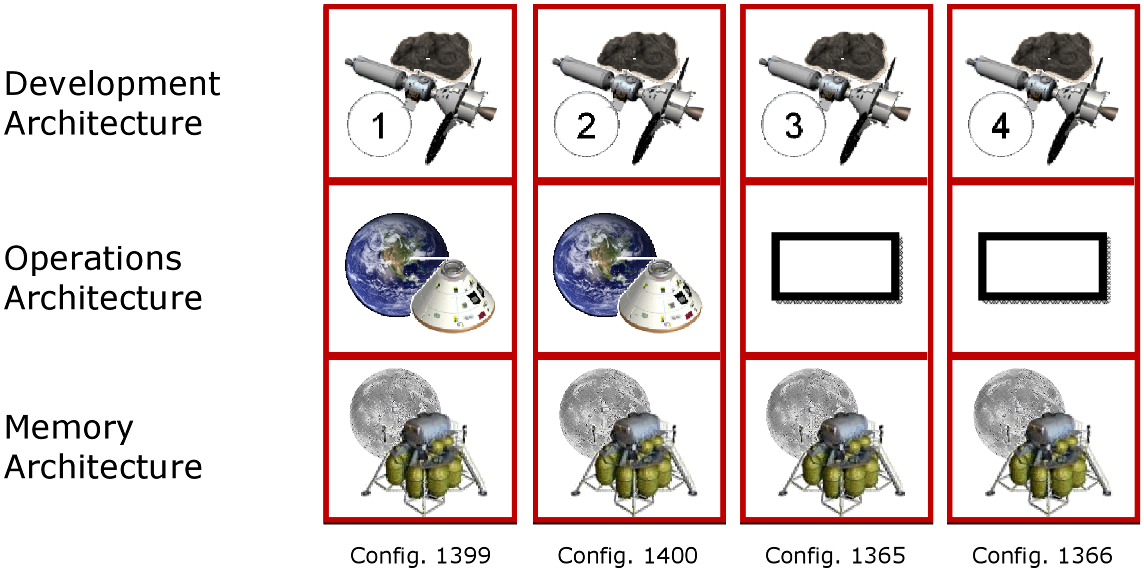

Figure 57. Notional Sequence of Configurations (Architecture Triplets).

6.1.2. Defining the Cost Transition Matrix

A key component of Step 1 of this framework is the definition of transition costs among configurations. To accomplish this, a cost model is necessary. While the cost transition matrices of the previous examples in Chapters 4 and 5 consisted of 81 and 225 elements, respectively, and in the absence of an automated cost model may have been estimable manually by an experienced cost analyst, the same cannot be said about the present application. With 3,286 configurations, the cost transition matrix for the NASA human space exploration application consists of nearly 10.8 million elements, without question *requiring* an automated model. This model, developed specifically for this application from publicly-available data, is discussed in detail in Appendix B and summarized in brief via Figure 58. The model takes as inputs the architecture definitions of Table 31 and, coupled with mass estimates for each architecture component (also documented in Appendix B) and the definitions of each configuration (or architecture triplet, as described in Section 6.1.1), combines estimates of development, production, management and systems engineering, operations, retirement, and termination liability costs to produce an estimate for the total cost required to transition from one configuration to another over a two-year time step. The model has components coded in MATLAB (outlined in orange in Figure 58) and in Microsoft Excel and Visual Basic (outlined in green in Figure 58); once executed with a given set of inputs, it is able to populate a full cost transition matrix within 25 minutes on a standard desktop computer.

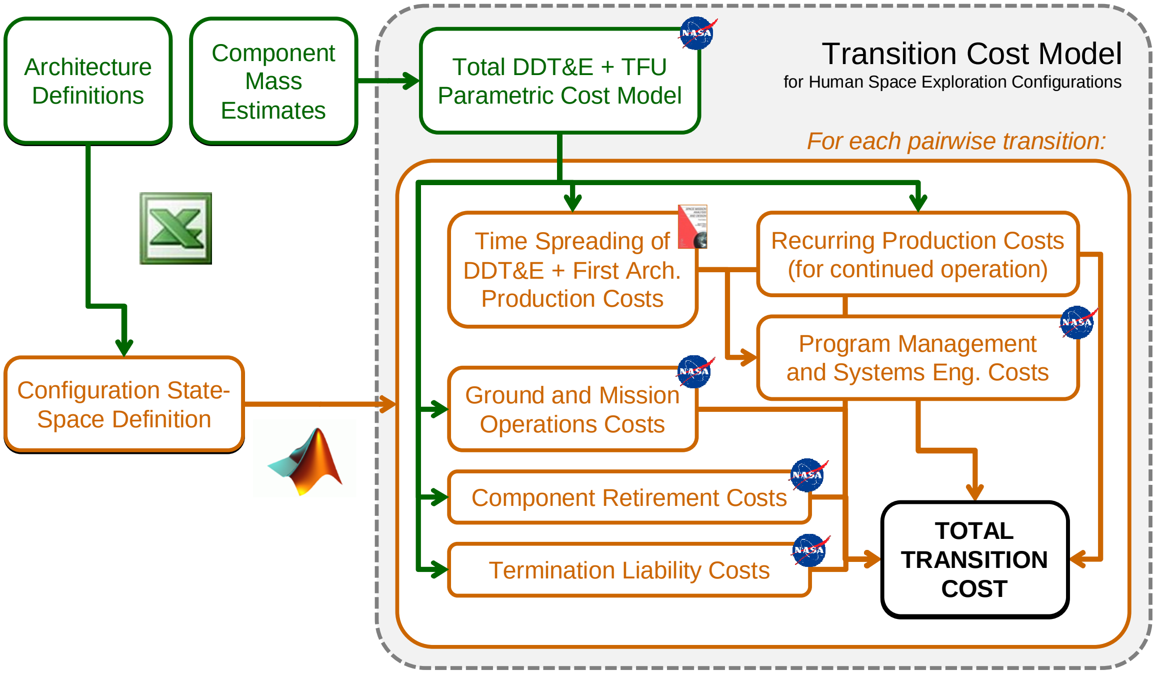
Architecture Definitions
Component Mass Estimates
Total DDT&E + TFU Parametric Cost Model
Transition Cost Model
for Human Space Exploration Configurations
For each pairwise transition:
Time Spreading of DDT&E + First Arch. Production Costs
Recurring Production Costs (for continued operation)
Program Management and Systems Eng. Costs
Configuration State-Space Definition
Ground and Mission Operations Costs
Component Retirement Costs
Termination Liability Costs
TOTAL TRANSITION COST

Figure 58. Transition Cost Model for Human Space Exploration Configurations.

In the previous examples of Chapters 4 and 5, any configuration could be reached from any other configuration given a sufficient expenditure of resources. However, with the introduction of intermediate architectures in this new configuration state space, it becomes be evident that some transitions will no longer make logical sense. For example, it should not be possible to skip phases of development (see Rule 3 to follow).

Other transitions violate cost model assumptions. For example, a four-period (eight-year) development is costed to include production of the flight units necessary for the first period of operation as listed in Table 31; for each period in which a configuration is *not* in the final phase of architecture development, production costs are estimated for the continuation of operations of the current operations architecture into the next period.

Applying this costing assumption, it would not be consistent to allow a current operations architecture to be used in a subsequent configuration if the current development architecture is in its final phase, since production has been accounted only for fielding of the new development architecture (see Rule 5 below).

These disallowed transitions are summarized by the six restrictions below, which are correlated with the illustrations in Figure 59:

1. ***Improper Program Initiation.*** Development of an architecture may only be initiated into phase 1.
2. ***Premature Operations Initiation.*** An architecture may not be placed into operation if it had not been in the final phase of development in the previous period (or, alternatively, if it or an architecture of which it is a subset had not been in operation in the previous period).
3. ***Premature Advancement.*** Development of an architecture may not advance more than one phase in one time period.
4. ***Unavailable Memory Architecture.*** An architecture that is not a subset of the current operations architecture, memory architecture, or just-completed development architecture may not be placed into memory.
5. ***Unavailable Operations Architecture.*** If development has just completed on a particular architecture, an architecture that is not the same as or a subset of this architecture can not be placed into operation. If an architecture is in mid-development, the next-period operations architecture must be the same as or a subset of the current operations architecture.
6. ***Stagnation and Partial Backtracking.*** If a particular architecture remains in development between periods, it must either progress in development phase or

restart in phase 1. Stagnation and partial backtracking (e.g., remaining in phase 4, or backtracking from phase 3 to 2) is not permitted.*

In concert with the recognition of disallowed configuration transitions, several configurations themselves become evident as unnecessary: For example, since the memory architecture is only useful in reducing costs only if it differs from the operations architecture or from a completed (phase 4) development architecture, configurations with the same architecture in memory as in operation (excluding those with both the "Nothing" architecture in memory and operation) can be removed from the state space. This first filter removes 333 configurations from the theoretical 3,700. Second, 81 configurations with the same architecture in memory as in completed (phase 4) development are removed from the state space. In total, 414 unnecessary configurations are removed, bringing the total number of configurations in this application's state space to 3,286.

---

* That this rule should be important is not intuitive and is related to the fact that the underlying cost models distribute both development and first-period production costs based on historical data [98] over the four-period development timeframe. Prior to its implementation, the dynamic programming optimization algorithm in Step 4 cleverly found solutions in which an architecture could be developed to phase 4, for example, and remain in phase 4 development indefinitely while continuing to field operational flight units at a fraction of the cost that would normally be required to produce the full set of flight units if not distributed over the four-period development. Note also that this rule only applies to the *decision* to proceed in developing an architecture. It will be possible, as will be described in Section 6.4, for stagnation to occur in the final period of development, but only as a result of probabilistic schedule slippage and not as a result of a choice on the part of the decision-maker.

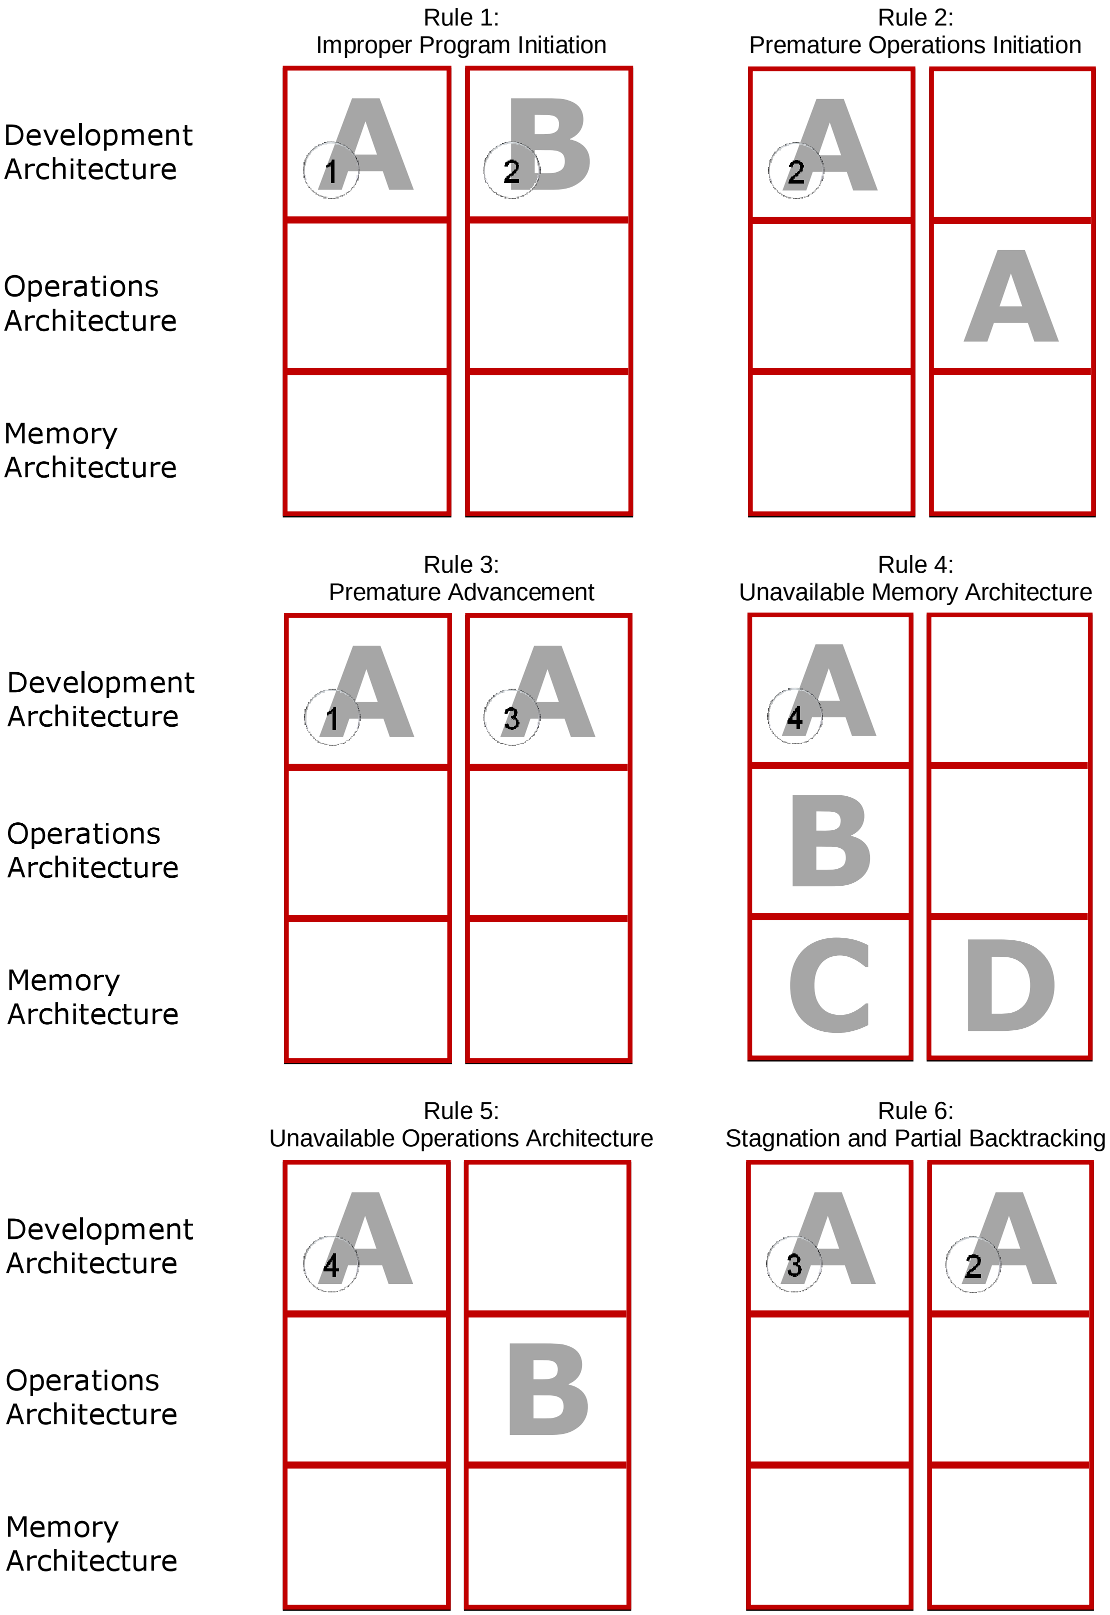
Rule 1:
Improper Program Initiation
Rule 2:
Premature Operations Initiation
Development Architecture
Operations Architecture
Memory Architecture
A
1
B
2
A
2
A
Rule 3:
Premature Advancement
Rule 4:
Unavailable Memory Architecture
A
1
A
3
A
4
B
C
D
Rule 5:
Unavailable Operations Architecture
Rule 6:
Stagnation and Partial Backtracking
A
4
B
A
3
A
2

Figure 59. Illustrations of Disallowed Transition Rules.

While the resulting 3,286 × 3,286 cost transition matrix is too large to reproduce here as it was in the examples of Chapters 4-5, it is visualized in part via Figure 60. On the left in this figure, all matrix elements representing transitions not excluded by the transition rules are marked in black or gray; the darker the element, the lower the cost of transition. In this matrix, 453,007 elements (4.2% of the 10.8 million total matrix elements) have costs associated with them; the remaining elements are excluded due to the transition rules. On the right in this figure is the distribution of these per-period costs. Note that many (approximately 41%) of these costs are at or below the $12.9 billion budget obtained when the NASA FY11 authorization for exploration plus non-International-Space-Station operations [109] is doubled to obtain an appropriate budget estimate for a two-year period length. Other transition costs are quite high; later in this analysis we will consider the ability of this framework to exclude consideration of transitions that are too costly in the short run. This per-period cost limitation will be implemented, in effect, as a seventh transition rule.

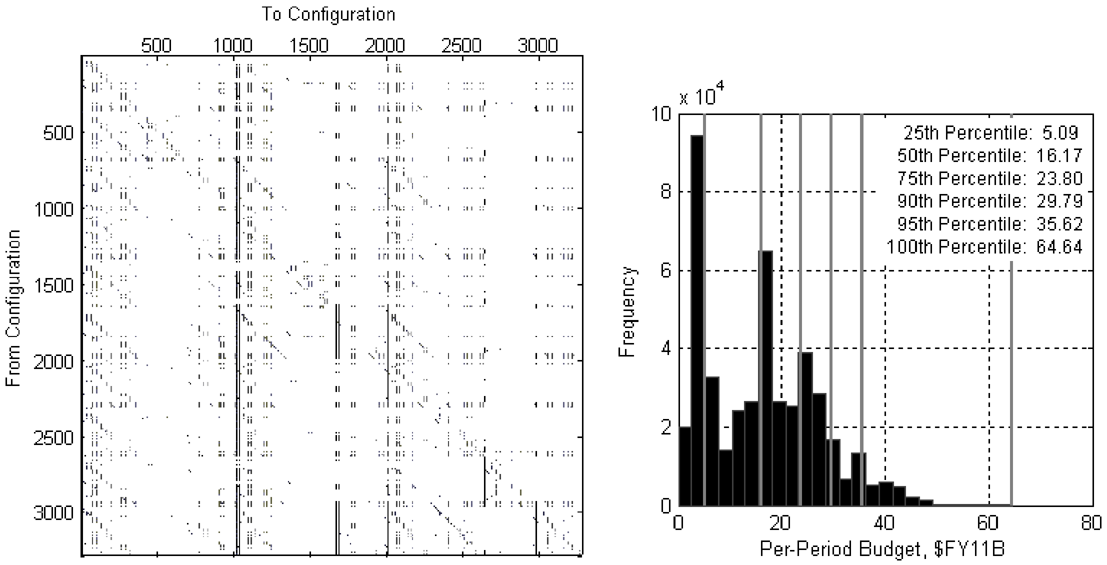

Figure 60. Cost Transition Matrix Visualization (left) and Distribution of Costs (right).

6.1.3. Analyzing the Cost Transition Matrix

As in the previous examples of Chapters 4-5, the data represented by the cost transition matrix can be analyzed, visualized, and related to flexibility in several useful ways. Covered for this application are (1) visualization of transitions available through the configuration and architecture state spaces and (2) analysis of available transitions from various configuration states as a function of budget (i.e., $\Phi_i(b)$).

6.1.3.1. Visualization of State Space Transitions

In the previous satellite examples of Chapters 4-5, the configuration state spaces were small enough to allow relatively uncomplicated visualizations of available transitions (as links) between configuration states (as nodes) for given budgets in the context of the two-period state-centric notion of flexibility. While in principle the same visualizations can be created for large configuration state spaces, some care must be taken to reduce the large volume of data (in this case, for a 3,286-configuration state space) to an interpretable form. In support of this goal, this subsection presents two views of configuration or architecture state space transitions: The first view, most analogous to those in Chapters 4-5, deals with the configuration state space and deals with transitions available among all 3,286 configurations in the configuration state space over the two-year time increment selected for the NASA application. The second view presents a less complicated view of transition options among architectures (not configurations), with the disadvantage that only time-independent costs can be considered.

6.1.3.1.1. Configuration State Space Transitions

With the cost transition matrix for a two-year time increment defined, visualization of available transitions in the configuration state-space also requires a definition of the arrangement of the configuration nodes. While there exist many arrangements (for example, the example of Chapter 4 used a euclidean space defined by the two design variables, and that of Chapter 5 chose a circular arrangement since a single cardinal or ordinal design variable was not as clearly defined), perhaps the most intuitive choice for visualization of the NASA example is to match the {Development, Operations, Memory} architecture triplet itself to three orthogonal axes in a euclidean space. Visualization based on this choice of a configuration node arrangement is shown in Figure 61 through Figure 65.

In Figure 61 through Figure 65, each of the 3,286 configurations exists as one of the gray points, plotted by its development, operations, and memory architectures. Recall that 10 architectures are considered in this example (see Table 31), and each (except for Architecture 1, the "Nothing" architecture) has four phases of development. Thus, the operations and memory architecture axes take values from one to ten, and the development architecture axes of the figures are labeled alphanumerically, where the initial number indicates the architecture number and the letter indicates the phase of development (e.g., development architecture 5c in these plots refers to the third phase of development of Architecture 5).

Each of the figures is labeled with a particular budget, and a link is drawn between two configurations in a figure for every element of the total cost transition matrix with a value less than or equal to the given budget. The color of the link indicates the cost of the transition: Blue indicates a low cost and bright orange indicates a high cost relative to the budget. As a consequence of this relative color selection, as budget is increased through the figures, the bright orange links are those that have been just enabled by the increased budget. Gray links throughout indicate transitions that "retreat" to the configuration with the "Nothing" architecture in development, operations, and memory.

The most clear observation from these figures is the dramatic increase in available transitions as budget increases from low levels and the apparent leveling off at higher budget levels (the difference in the number of transitions available between $5 billion and $12.9 billion is not as visually apparent as between $500 million and $5 billion, for example). Particularly at low budget levels, it may be noted that the available transitions have points near the origin (i.e., “Nothing” as the development, operations, and memory architecture) in common; if directed arrows were placed on each link it would be seen that most available transitions at these levels are “retreating” towards lower capabilities at low budget levels. However, at the lowest budget level, it is notable that many configurations have no options at all due, for example, to shutdown costs that exceed the allowed budget. At the $500 million budget level, only configurations with either the LEO-themed architecture or nothing in operation have any options at all. As in the previous examples, this visualization illustrates that the number of transitions available to a system can be a strong (and nonlinear) function of available resources or budget.

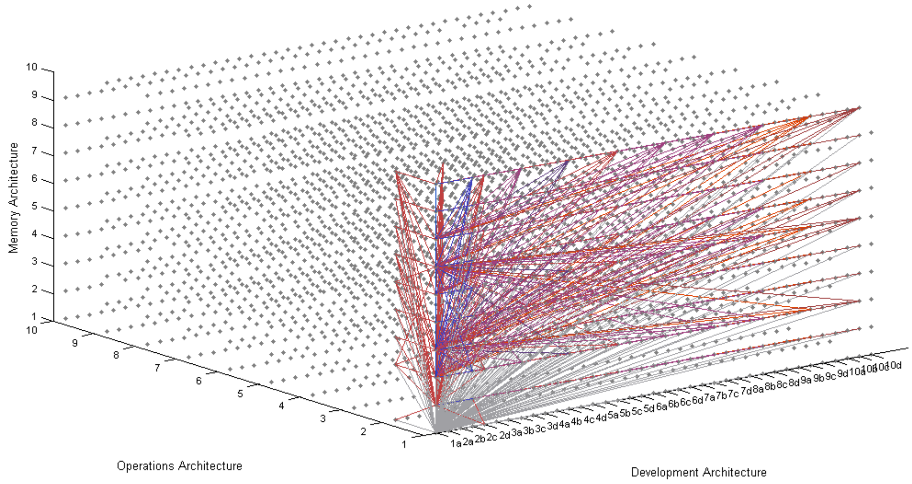

Figure 61. Available configuration transitions for a $500 million 2-year budget.

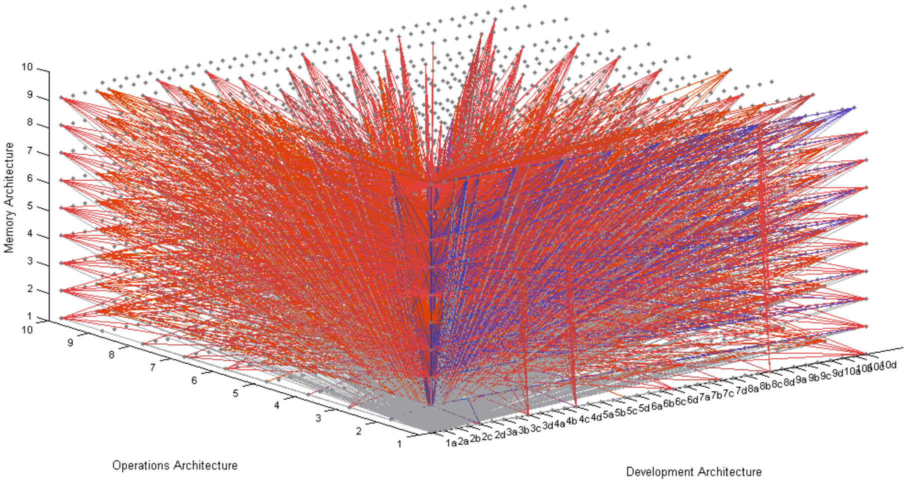

Figure 62. Available configuration transitions for a $1 billion 2-year budget.

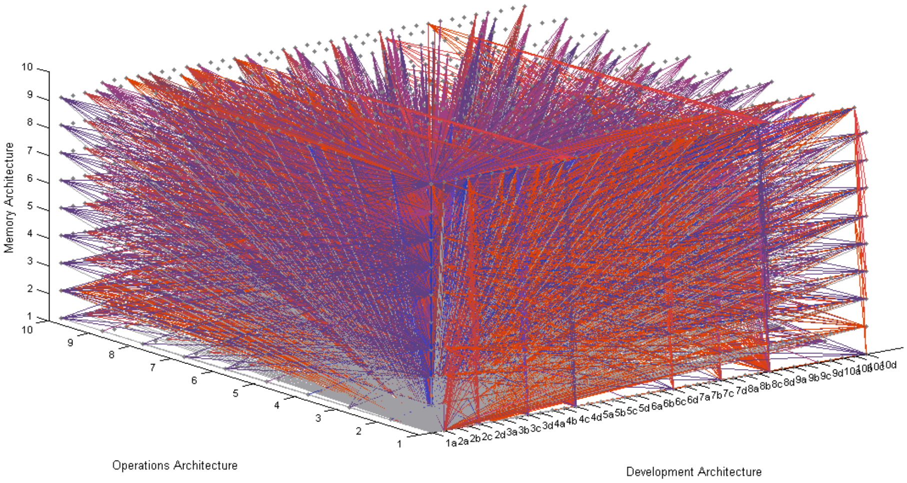

Figure 63. Available configuration transitions for a $2 billion 2-year budget.

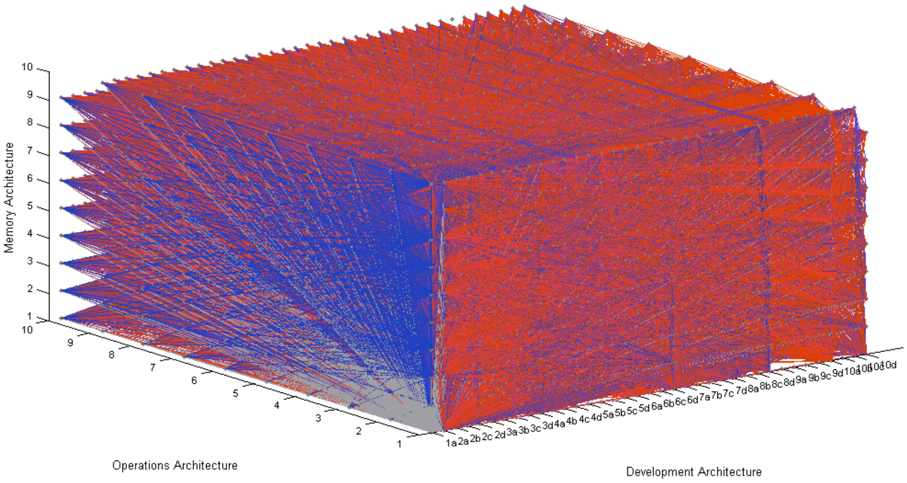

Figure 64. Available configuration transitions for a $5 billion 2-year budget.

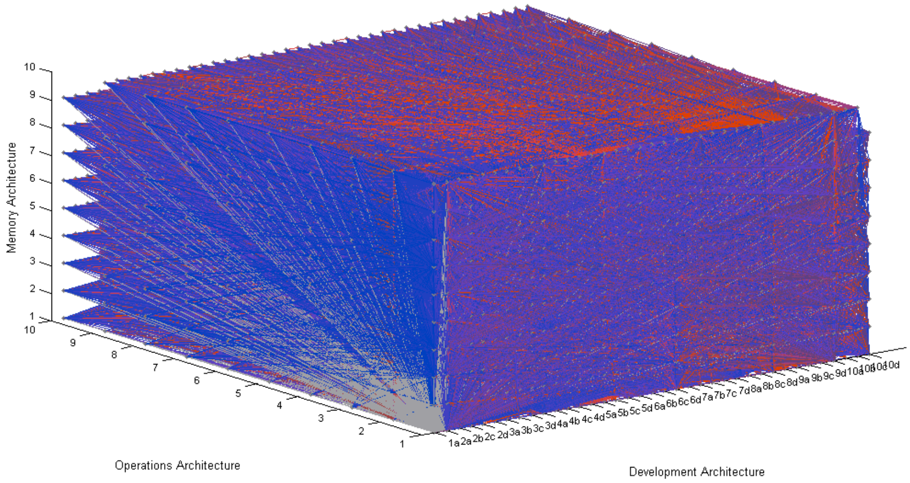

Figure 65. Available configuration transitions for a $12.9 billion 2-year budget.

6.1.3.1.2. Transitions for Time-Independent Costs in the Architecture Space

While Figure 61 through Figure 65 are comprehensive and precise in visually recording all transitions available for a given budget among all 3,286 available system configurations over the two-year time increment for the human space exploration application, they provide so much data that some trends and insight may be easily lost. To partially overcome this limitation, the following brief analysis reverts to the simpler visualization of the architecture space in Figure 56. In this visualization, the ten architectures under consideration from Table 31 are projected onto a two-dimensional plane such that similar architectures are grouped together and dissimilar architectures are placed farther apart. While the following analysis has the limitation that it does not consider certain costs, it lends some helpful insight into the set of architectures under consideration and the relative costs of switching between them.

In the following analysis, the transition cost model for human space exploration configurations described in Appendix B is applied to estimate only the "time-independent" costs of DDT&E, DDT&E-related program management and systems engineering, and retirement. Unlike production, operations, and termination liability costs, which depend directly on a decision-maker's choices regarding whether to extend or shorten the duration of a program, these time-independent costs define a minimum bound for total costs incurred to successfully transition from the development and operation of one architecture to another. These costs are shown in Table 33, which has properties similar to those of development cost matrices of the examples in Chapters 4-5. Notably, the diagonal consists entirely of zeros, signifying that it costs nothing to develop architecture *i* given that architecture *i* already exists; and since architecture *i* is desired next, neither are any retirement costs required. Similarly, the elements in the first row (from the "Nothing" architecture) are the highest costs in any given column; based directly on the data in this matrix, Figure 66 adds further empirical confirmation to the expectation that greater initial investment in an architecture tends to result in lower future switching costs. However, unlike the previous examples of Chapters 4-5, note that the first column of this matrix does not consist entirely of zeros. This is due to the retirement costs accounted for in the transition cost model (for details, see Appendix B), which in Step 4 of this framework will impose an additional barrier to changing architectures and in effect add inertia toward the continuation of current program plans as a consideration in the search for optimal policies.

Table 33. Time-Independent Architecture Transition Costs (in $FY11B).

		Architecture									
		-1- Nothing	-2- LEO	-3- GEO Servicing	-4- Lunar Orbit	-5- Lunar Surface	-6- Sun-Earth L2	-7- Near-Earth Object	-8- Mars Moon	-9- Mars Surface	-10- Deep Space
Architecture	-1- Nothing	0.00	26.15	60.93	44.15	70.12	64.09	60.93	61.73	95.88	68.17
	-2- LEO	0.41	0.00	47.84	31.07	57.44	51.01	47.84	48.64	83.20	55.08
	-3- GEO Servicing	0.79	13.49	0.00	0.29	19.21	3.16	0.00	4.16	38.72	7.24
	-4- Lunar Orbit	0.61	13.30	16.78	0.00	26.38	19.94	16.78	17.57	52.14	24.02
	-5- Lunar Surface	0.90	14.00	10.17	0.83	0.00	13.33	10.17	14.29	36.37	17.41
	-6- Sun-Earth L2	0.82	13.52	0.08	0.33	19.26	0.00	0.08	4.21	35.55	4.08
	-7- Near-Earth Object	0.79	13.49	0.00	0.29	19.21	3.16	0.00	4.16	38.72	7.24
	-8- Mars Moon	0.80	13.50	3.37	0.30	22.55	6.54	3.37	0.00	34.57	6.44
	-9- Mars Surface	1.14	14.26	4.25	1.12	10.94	4.21	4.25	0.92	0.00	4.17
	-10- Deep Space	0.86	13.56	0.15	0.38	19.32	0.09	0.15	0.13	31.48	0.00

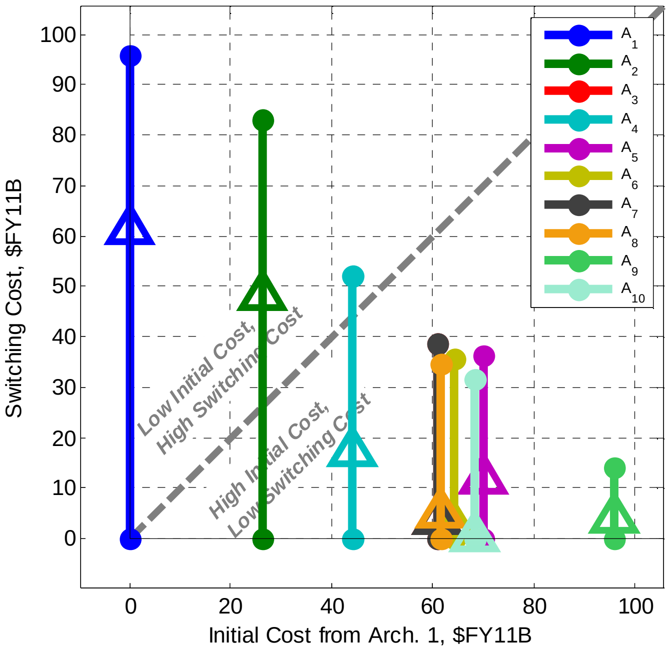

Figure 66. Switching cost vs. initial cost from Architecture 1. *Vertical lines indicate ranges of switching costs from each configuration; some overlap. Solid dots indicate minima and maxima, and triangles indicate median values.*

Using the cost matrix in Table 33 to draw available transitions as links between the architectures in the projection of Figure 56 results in the visualization of Figure 67. Each node in Figure 67 represents an architecture in the architecture space. Each node is named A_X, where X is the configuration number from Table 33, and has a color indicative of the initial cost to develop the architecture from the "Nothing" architecture (Archictecture 1); blue indicates an architecture with a low initial investment cost (e.g., the "Nothing" or LEO-themed architectures), while bright orange indicates an architecture with a high initial investment cost (e.g., the Mars-Surface-themed architecture). In these projections, the lower-initial-cost architectures tend to appear toward the bottom right, while the high-initial-cost architectures tend to appear toward the upper left. Above each of the plots is a budget, and a directed link is drawn for every element of the cost transition matrix less than or equal to the given budget. In this particular matrix, the diagonal consists of zeros, so a dark ring encircles every node to indicate that self-transitions are possible for any budget.

As usual, Figure 67 illustrates that higher budgets permit more transitions, and again illustrated is how the increase in number of transitions can be highly nonlinear with budget. For example, substantially more transitions become visible in the budget increase from $500 million to $5 billion than in the even greater budget interval from $5 billion to $10 billion.

More interesting, however, is *where* the transitions appear, which could not be easily ascertained from the large configuration-space plots in Section 6.1.3.1.1. Note that at the low $500 million budget level, many transitions are available; however, as Figure 67 shows, these are all *local* transitions. Recall that the architectures in Figure 67 are arranged such that physically similar architectures are located nearer to each other, while physically dissimilar architectures are located farther from each other. The $500 million budget plot shows that Architectures 3, 4, 6, 7, 8, and 10, which are characterized by similar sets of components intended for beyond-LEO but non-surface destinations, have some flexibility to transition to each other, but not to architectures outside their local group – not even to the "Nothing" architecture (Architecture 1), which would require retirement costs higher than the $500 million budget. As the budget is increased, the ability to transition between distant groups also increases, at first in the "retreating" or "shutdown" direction toward the lower-cost architectures at the bottom right and then in both directions. If the budget were raised to the $96 billion maximum of the matrix, all pairwise links would appear.

In summary, while this analysis was prefaced with the acknowledgement that it neglects some of the costs incurred to a decision-maker, it is illustrative in that it highlights the particular influence of shutdown costs and the relationship between the physical similarity of architectures and costs of transition. In considering the ten human space exploration architectures of Table 31 throughout the rest of this chapter, it may be helpful to refer back to this simple set of data and analysis for physical understanding of the architectures under consideration.

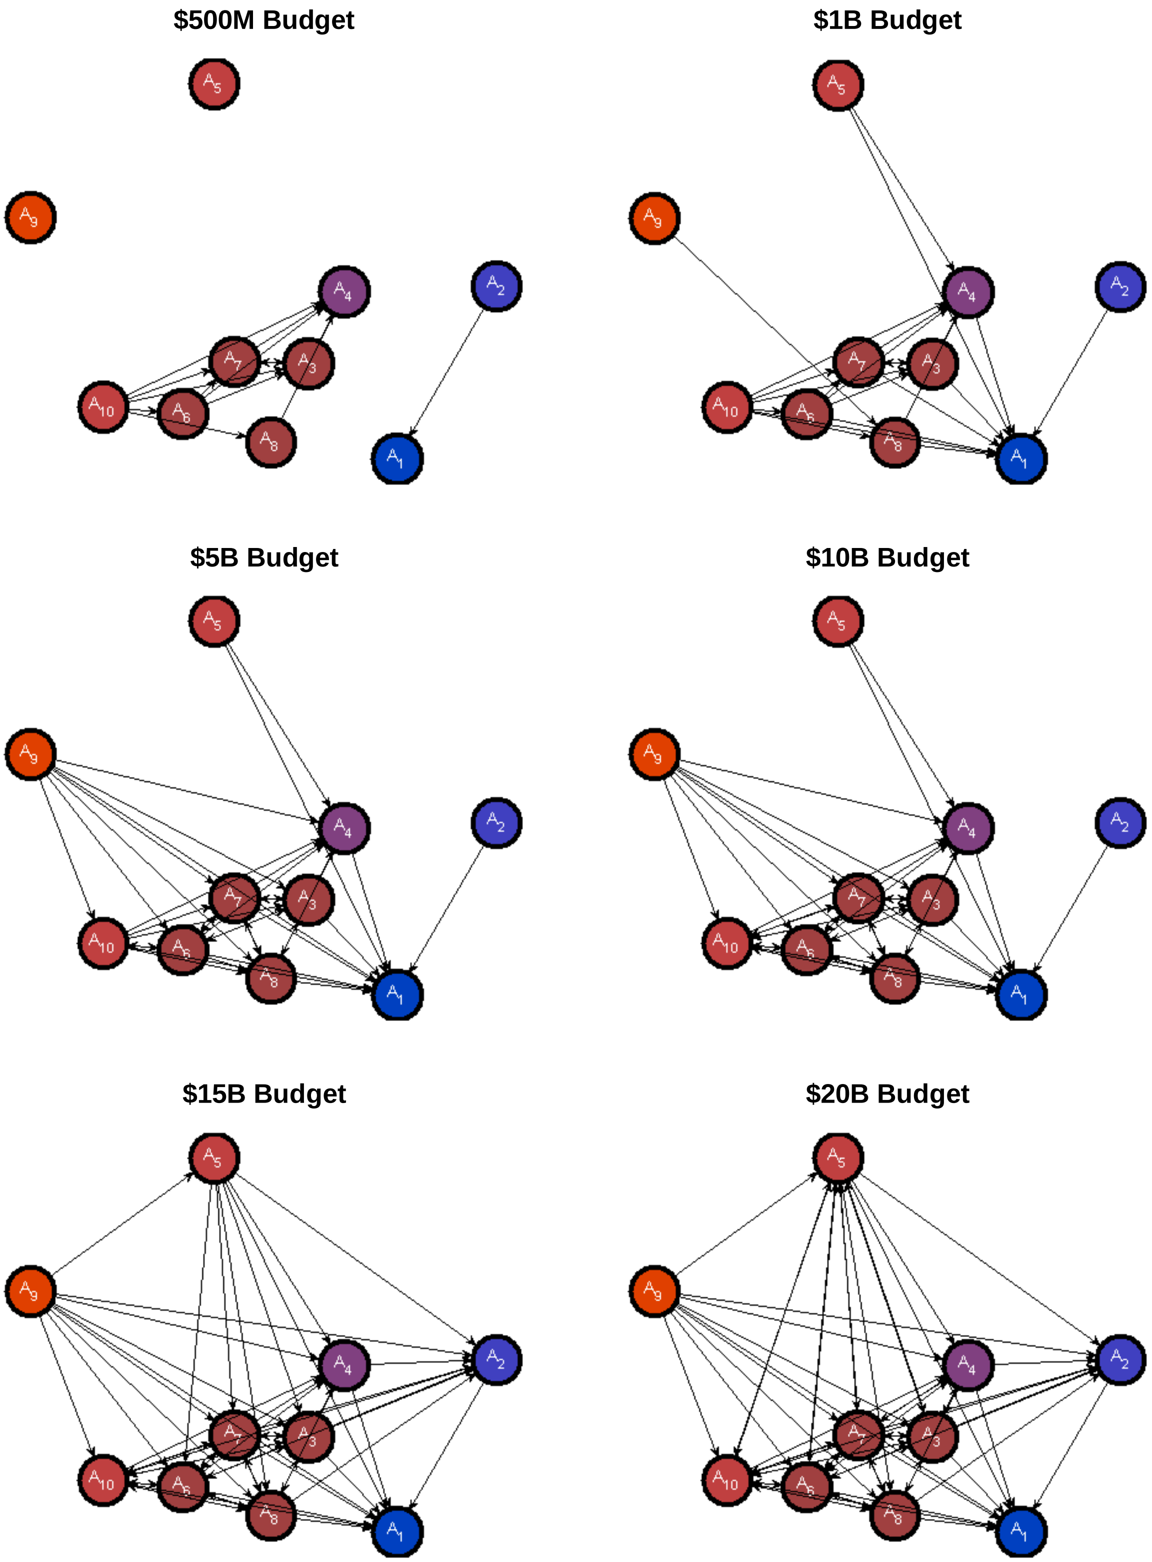

$500M Budget

$1B Budget

$5B Budget

$10B Budget

$15B Budget

$20B Budget

Figure 67. Available architecture transitions for $0.5, 1, 5, 10, 15, and 20 billion budgets.

6.1.3.2. Architecture and Configuration Transitions vs. Budget

As in Chapters 4-5, it is possible to extend the concept of the state-space visualizations from the previous analyses to consider a continuum of budgets with the goal of better understanding the sensitivity of transition options to available resources.

6.1.3.2.1. Architecture Transitions vs. Time-Independent Budget

Prior to considering the full costs of transition, it is instructive to begin with the analysis of the time-independent costs that allowed the view of available transitions in the architecture space in Section 6.1.3.1.2. Here, as in the examples of previous chapters, the per-period budget may be increased on a continuum and the number of available transitions away from a given node (or architecture) may be tracked. Since there are ten architectures in each of the plots in Figure 67, there are ten such values to be tracked, each of which is plotted as a function of budget in Figure 68.

Note that each line in Figure 68 is a monotonically increasing function of budget, but that each rises at a different overall rate. For example, note that Architecture 10 (the general Deep Space architecture) rises quickly to seven available transitions at a budget of less than $1 billion, while it takes Architecture 1 (the "Nothing" architecture) a budget of over $64 billion to reach the same number of options. In general, Figure 68 suggests that this availability of transitions for low budgets is a property of the number and type of components in an architecture: Architectures 1 and 2 have few or no components in common with other architectures and thus incur large costs to transition to any others, whereas Architectures 3, 7, 8, 9, and 10 (and especially Architecture 10) have many components in common with other architectures and incur smaller transition costs[*]. Note

---

[*] This explains the common coupling of the concepts of modularity and flexibility. In the transition cost model used to the generate the data in Figure 68, development costs were additive by component, and this modularity-representative modeling structure produced benefits when existing components needed not be re-developed for the fielding of a new system. However, it

also that since Architectures 3 and 7 require the development of all the same components (albeit that they require different production numbers of each component), their time-independent transition characteristics overlap and only Architecture 7's characteristic is seen in Figure 68.

It may further be seen in Figure 68 that the rapid rise in the number of links in Figure 67 between \$0 and \$10 billion can be easily observed as steep increases for many of the architecture transition lines. The gradual taper in the increase in number of links in Figure 67 at high budgets can also be observed as each of the lines in Figure 68 tends to plateau as it approaches the ten-transition maximum.

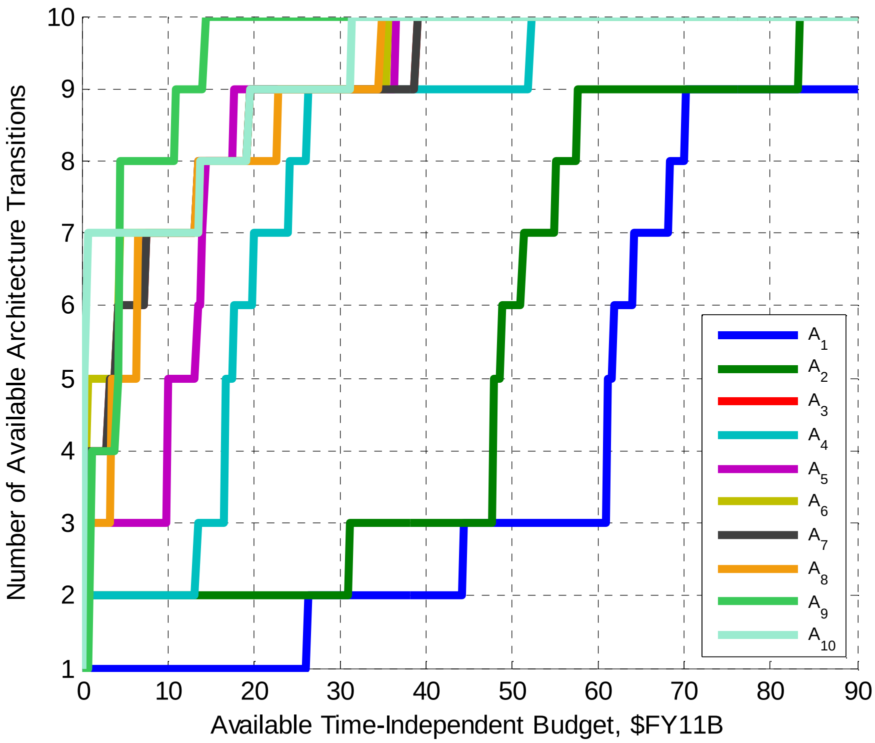

Figure 68. Available architecture transitions vs. available time-independent budget.

should be recognized that the results of Figure 68 can be produced no matter what modeling structure applies to a given problem of interest, and it is conceivable that other strategies could also produce flexibility. Thus, while modularity may in general be an important and common means to achieving flexibility, it neither guarantees flexibility in every situation, nor is it necessarily the only way to achieve flexibility.

6.1.3.2.2. Configuration Transitions vs. Full Per-Period Budget

In the context of understanding the scope of transition options for resources available in a given timeframe, it is most meaningful to conduct a continuum budget analysis for the full cost of transition (rather than only the time-independent costs, as in Section 6.1.3.2.1).

Thus, plotting the number of available transitions from each of the 3,286 nodes in Figure 61 through Figure 65 as a function of available per-period budget yields the result in Figure 69. Since there exist 3,286 configurations from which transitions can be made, there also exist 3,286 lines in Figure 69. Note that, as is typical in these transition vs. budget plots, all lines are monotonically increasing, indicating that the number of transitions available from (or, approximately speaking, the flexibility of) a given configuration cannot decrease with increasing budget.

To better facilitate analysis, each line in Figure 69 is colored by its corresponding operations architecture. This reveals, for example, that configurations with Architecture 10 in operation are distinguished by high numbers of transitions whereas configurations with Architecture 1 in operation tend to have low numbers of available transitions. This correlation with operations architecture can be attributed to the fact, for example, that the presence of a high-capability operations architecture does not only enable the operation or placement into memory of any lower-capability architecture in the following period, but makes less costly the development of subsequent architectures with common components.

One difference that Figure 69 exhibits when compared to other transition vs. budget plots in this thesis is that the lines representing each configuration no longer plateau at the same maximum Φ value. This is a consequence of the transition rules introduced in Section 6.1.2, and thus Figure 69 no longer solely conveys information about which configurations have more options than others, but it also contains information about the maximum potential a configuration has to gain options with any amount of budgetary resources. For example, Figure 69 shows that many of the configurations with Architecture 10 in operation plateau at between Φ = 400 and Φ = 600 transitions, while those with Architecture 1 in operation plateau at less than Φ = 100 transitions. Interestingly, the black dashed line in Figure 69 indicates an approximation for the current configuration of NASA's human space exploration development efforts, with a LEO-themed architecture in the second phase of development and no relevant exploration architectures in operations or memory (in shorthand notation, [2b 1 1]). This serves as a clear example of a configuration with few options even at high per-period budgets: As Figure 69 shows, this configuration plateaus at a value of just Φ = 11 available transitions by a budget of $22 billion.

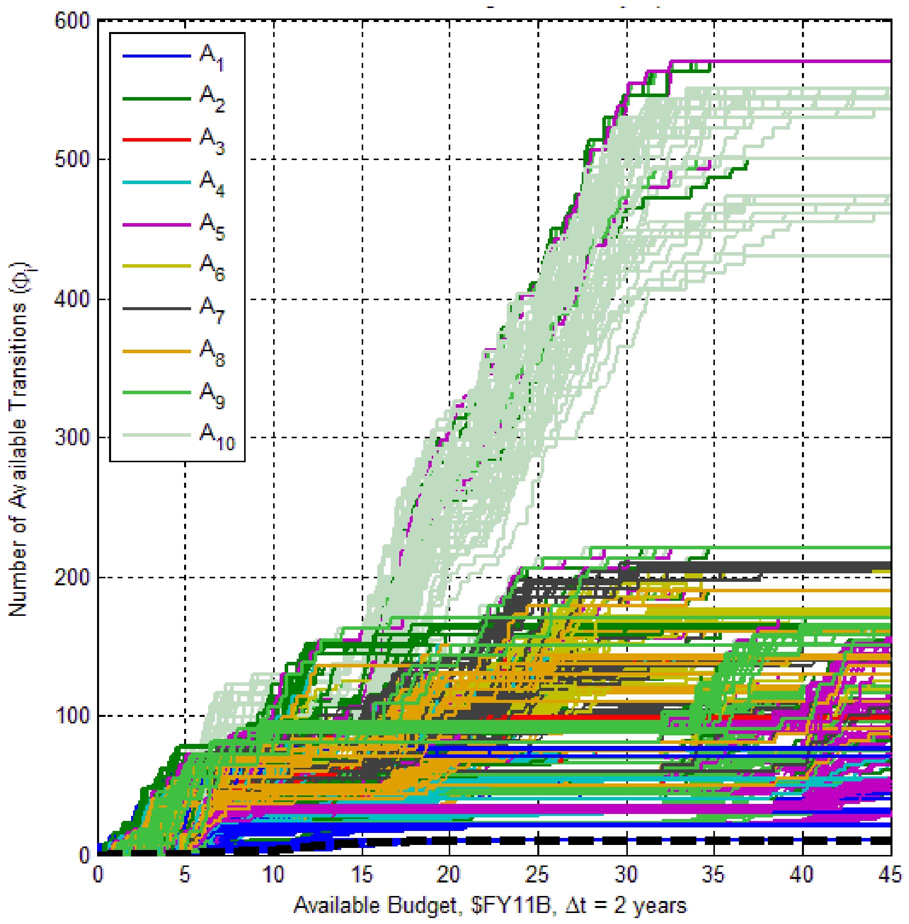

Figure 69. Number of available transitions vs. available budget over two years. *Lines are colored by operations architecture, and the black dashed line indicates the characteristic for an approximation of NASA's present configuration: [2b 1 1].*

One observation made earlier regarding the data in Figure 69 was that configurations with Architecture 10 in operation are distinguished by high numbers of transitions whereas configurations with Architecture 1 in operation tend to have low numbers of available transitions. While this statement appears substantially justified by observation of Figure 69 itself, some additional exploration is warranted.

Figure 70 and Figure 71 show cross-sections of Figure 69 taken at budget levels of $12.9 billion and $25 billion, respectively. Each histogram in the figures shows the distribution of the number of available transitions for configurations with given operations architectures. The histogram *x*-axis range internal to each of Figure 70 and Figure 71 is consistent and marked on the bottom plot, and thus the central tendencies and dispersion of Φ for different operations architectures (due to the fact that Φ is determined not only from a configuration's operations architecture, but also from its development and memory architecture) can be compared visually. Of particular note is the fact that the mean of the distribution for Architecture 10 in Figure 70 is about 28% higher than the next-highest mean, while the mean of the Architecture 10 distribution in Figure 71 is 137% higher that the next-highest mean. In other words, the budget level of interest affects the relative flexibility of one configuration (or the central tendency for an architecture) over another. In the case of Architecture 10, at a low enough budget level it would not be accurate to say that as an operational architecture it distinguishes configurations as substantially more flexible (or having substantially more options) than others.

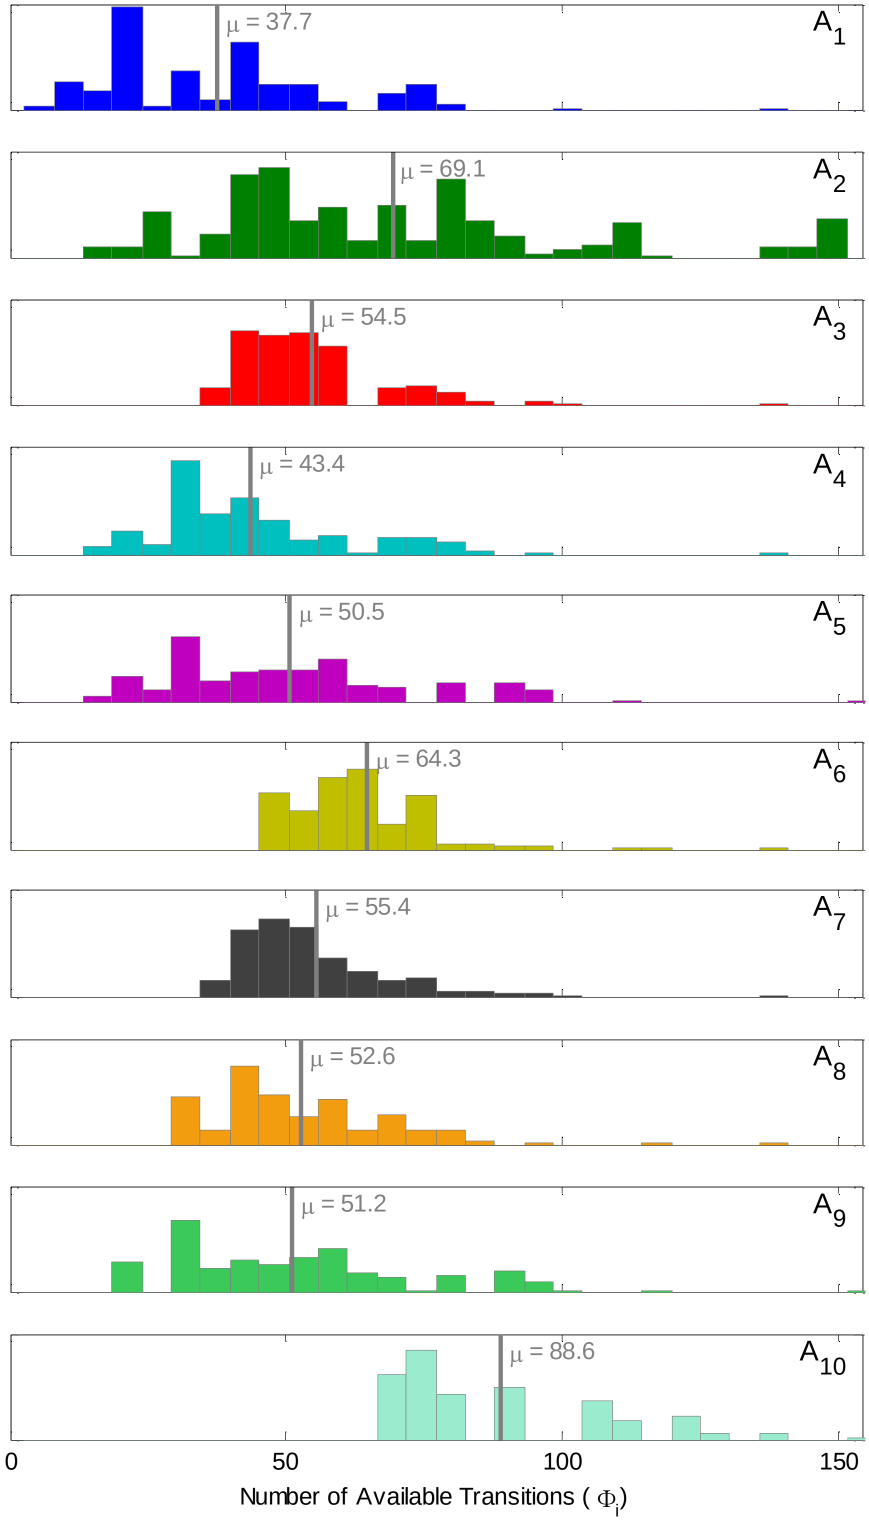

Figure 70. Distributions of number of available transitions for $12.9B per two-year period, by operations architecture. *Colors correspond to those used to represent architectures $A_1 - A_{10}$ in Figure 69.*

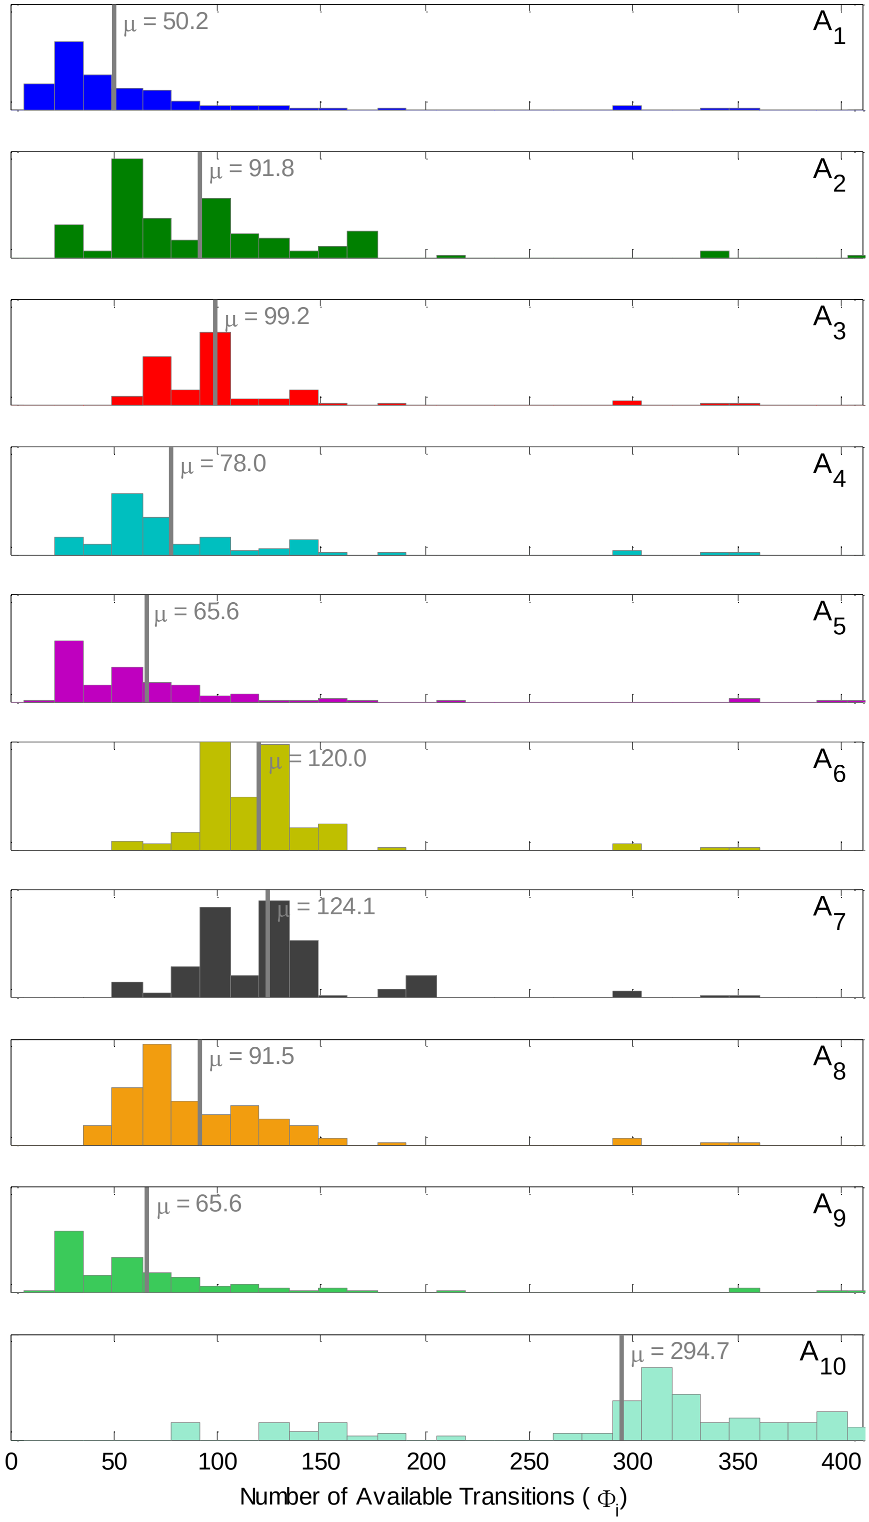

Figure 71. Distributions of number of available transitions for $25B per two-year period, by operations architecture. *Colors correspond to those used to represent architectures $A_1 – A_{10}$ in Figure 69.*

To conclude this analysis associated with the Step 1 cost transition matrix, Figure 72 shows an example in which the relative flexibility of two configurations reverses at a particular budget level (termed a “flexibility reversal” in Section 4.1.3.2). In this case, the blue line indicates a low-capability configuration of a LEO-themed architecture in the third phase of development with nothing in operation or memory, and the green line indicates a contrasting high-capability configuration of a Mars-surface-themed architecture in its first phase of development with a lunar-surface-themed architecture in operation and nothing in memory. While the high-capability configuration plateaus at a higher number of transitions, it also requires at least $4.9 billion to make its first transition because of commitments in the form of, at a minimum, termination liability and system retirement costs. This further highlights the importance of considering available resources (such as budget) when characterizing the flexibility of a space system; in this particular case, options exist over a substantial budget range with a lower-capability configuration that do not exist with a higher-capability configuration.

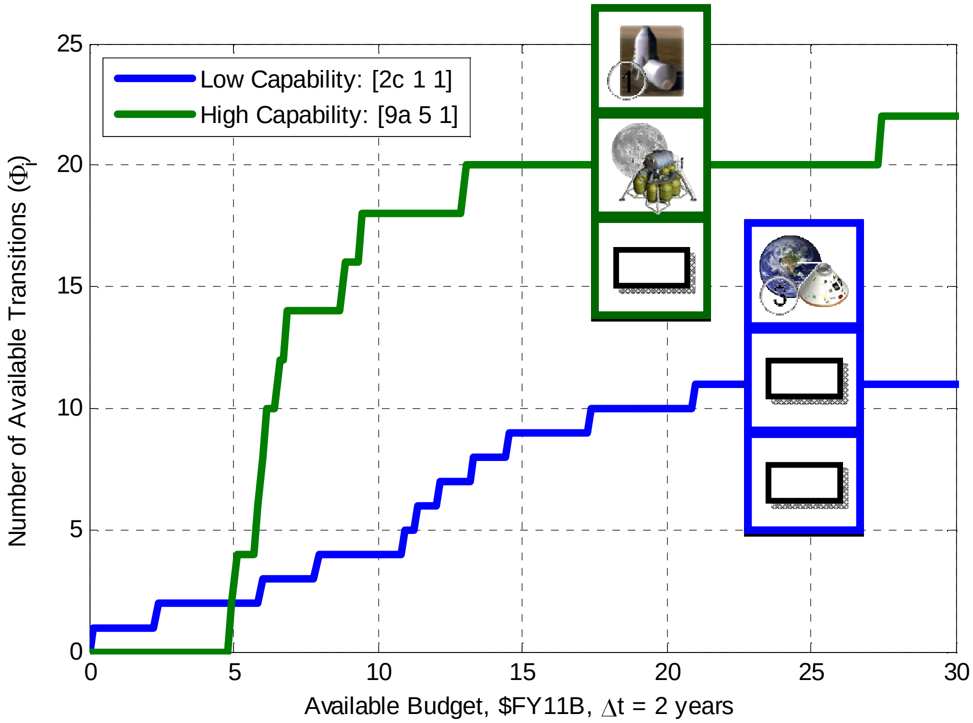

Figure 72. Example of a "Flexibility Reversal" in the NASA Application.

6.2. Step 2: Define Markovian Demand Environment Evolution

As in the examples of Chapters 4-5, Step 1 here has focused on defining the available configuration states for the system of interest, in this case comprised of development, operations, and memory architectures for the NASA human space exploration application. However, yet to be discussed is the demand environment in which the system will operate. Step 2 fills this gap with a model for the evolution of the state of the demand environment which, unlike the configuration state, evolves stochastically and is largely not under the control of the decision-maker.

In terms of the Flexible Path approach proposed by the Review of U.S. Human Spaceflight Plans Committee in 2009 (for details, see Section 1.3.3 or Ref. [34]), the mission demand environment can largely be classified in terms of mission destination. The committee's report [34] mentions that the Flexible Path approach is designed, for example, to allow decision-makers to respond to future circumstances calling for exploration of the surface of the Moon or Mars, or calling for the mounting of destination-oriented missions in response to discoveries such as life on Mars or near-Earth object threats. Additional factors influencing mission destination changes could include changes in political will that cause the reduction in scope of missions to destinations near Earth, the emergence of technological challenges from other nations that expand the scope of missions toward the Moon or beyond, and the successful (or failed) achievement of current goals in space which could have the effect of reducing or expanding the scope of mission destinations. Thus, in this human space exploration architecture application, inner solar system destination is used to characterize the mission demand state as either (1) Nothing, (2) Low Earth Orbit (LEO), (3) Geosynchronous Earth Orbit (GEO) Servicing, (4) Lunar Orbit, (5) Lunar Surface, (6) Earth-Moon L1, (7) Sun-Earth L2, (8) Venus Orbit, (9) Near-Earth Object, (10) Mars Orbit, (11) Martian Moon, or (12) Mars Surface. Note that, while some of the architectures in Table 31 are themed around some of these destinations, there does not exist a one-to-one correspondence.

Appendix C extensively details the derivation of a Markovian demand environment evolution model structured around these twelve mission destinations, based upon a survey of primarily NASA experts with substantial experience in the field of human space exploration. A key difference, however, between this model and the models in the examples of Chapters 4-5 is that it includes two conditional probability transition matrices: The first matrix, shown in Table 34, indicates probabilities of demand evolution given that current demand is fulfilled. The second matrix, shown in Table 35, indicates of probabilities of demand evolution given that current demand is not fulfilled. The separation of these matrices thus allows for modeling of configuration-dependent demand, or the reality that human space exploration mission demand does not evolve completely independently of NASA system decisions. Both matrices in Table 34 and Table 35 express transition probabilities over the two-year time step corresponding to the period length set at the initiation of the discussion in Section 6.1.

Table 34. Discrete-time Markov chain probability transition matrix for median expert inputs and Δt = 2 years, for the condition that current mission demand is fulfilled.

		To											
	Mission Demand	Noth.	LEO	GEO Serv.	Lunar Orbit	Lunar Surf.	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Mars Moon	Mars Surf.
From	Nothing	0.5180	0.2447	0.0301	0.0311	0.0928	0.0308	0.0027	0.0001	0.0372	0.0043	0.0037	0.0045
	LEO	0.0192	0.6784	0.0340	0.0577	0.1028	0.0395	0.0039	0.0002	0.0475	0.0060	0.0052	0.0057
	GEO Servicing	0.0261	0.0489	0.5192	0.0776	0.1483	0.0479	0.0157	0.0002	0.0598	0.0246	0.0190	0.0126
	Lunar Orbit	0.0101	0.0326	0.0266	0.3771	0.2868	0.0709	0.0295	0.0003	0.0664	0.0389	0.0338	0.0270
	Lunar Surface	0.0005	0.0046	0.0079	0.0080	0.8261	0.0195	0.0136	0.0002	0.0278	0.0240	0.0231	0.0447
	Earth-Moon L1	0.0095	0.0346	0.0223	0.0435	0.1491	0.5733	0.0259	0.0003	0.0522	0.0278	0.0255	0.0360
	Sun-Earth L2	0.0022	0.0439	0.0325	0.0466	0.1089	0.0448	0.4550	0.0005	0.1057	0.0637	0.0363	0.0598
	Venus Orbit	0.0018	0.0248	0.0201	0.0290	0.0957	0.0690	0.0447	0.2647	0.1950	0.0826	0.0568	0.1157
	Near-Earth Object	0.0006	0.0094	0.0076	0.0138	0.0431	0.0181	0.0141	0.0047	0.7242	0.0540	0.0453	0.0651
	Mars Orbit	0.0005	0.0106	0.0014	0.0024	0.0295	0.0207	0.0171	0.0006	0.0442	0.6123	0.0760	0.1846
	Mars Moon	0.0006	0.0100	0.0012	0.0020	0.0290	0.0187	0.0133	0.0007	0.0273	0.0439	0.6133	0.2400
	Mars Surface	0.0021	0.0011	0.0006	0.0009	0.0213	0.0068	0.0057	0.0039	0.0267	0.0025	0.0185	0.9099

Table 35. Discrete-time Markov chain probability transition matrix for median expert inputs and Δt = 2 years, for the condition that current mission demand is *not* fulfilled.

		To											
	Mission Demand	Noth.	LEO	GEO Serv.	Lunar Orbit	Lunar Surf.	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Mars Moon	Mars Surf.
From	Nothing	0.1417	0.5730	0.0495	0.0515	0.1171	0.0256	0.0029	0.0000	0.0284	0.0050	0.0029	0.0024
	LEO	0.0104	0.8055	0.0209	0.0405	0.0575	0.0234	0.0031	0.0000	0.0270	0.0063	0.0034	0.0019
	GEO Servicing	0.0151	0.1183	0.5203	0.0799	0.1572	0.0471	0.0072	0.0000	0.0436	0.0048	0.0030	0.0035
	Lunar Orbit	0.0009	0.0617	0.0267	0.5249	0.2376	0.0523	0.0197	0.0000	0.0417	0.0189	0.0093	0.0062
	Lunar Surface	0.0003	0.0221	0.0120	0.0291	0.8297	0.0220	0.0074	0.0000	0.0292	0.0180	0.0131	0.0170
	Earth-Moon L1	0.0009	0.0589	0.0269	0.0508	0.1273	0.6136	0.0104	0.0001	0.0501	0.0226	0.0098	0.0287
	Sun-Earth L2	0.0015	0.0813	0.0497	0.0687	0.1126	0.0757	0.3716	0.0001	0.1022	0.0509	0.0380	0.0475
	Venus Orbit	0.0015	0.0879	0.0429	0.0542	0.1325	0.0882	0.0441	0.1356	0.2076	0.0752	0.0521	0.0783
	Near-Earth Object	0.0004	0.0267	0.0145	0.0217	0.0608	0.0199	0.0088	0.0001	0.7579	0.0315	0.0267	0.0310
	Mars Orbit	0.0004	0.0281	0.0103	0.0219	0.0439	0.0305	0.0097	0.0002	0.0584	0.6743	0.0537	0.0688
	Mars Moon	0.0003	0.0297	0.0074	0.0094	0.0459	0.0094	0.0016	0.0005	0.0619	0.0168	0.6470	0.1701
	Mars Surface	0.0002	0.0192	0.0044	0.0093	0.0258	0.0110	0.0065	0.0037	0.0413	0.0179	0.0356	0.8252

To visualize the conditional Markov chains in Table 34 and Table 35 as is done for the Markov chains in the examples of Chapters 4-5, it is helpful to project them over more than one two-year time increment. Note that the probabilities on the diagonals of these matrices tend to quite high due to this short time step (naturally, as the time step of becomes smaller and smaller, the probability in remaining in a particular state would be expected to approach closer and closer to unity), and thus a visualization of the Markov chain on the two-year step would reveal only the obvious tendency for the system to stay in its current demand state over the coming period. Extending the time increment to an eight-year step for the purposes of visualization (by raising each matrix to the fourth power, or by using the uniformization procedure detailed in Appendix C) yields the diagrams in Figure 73 and Figure 74. In these figures, as in those depicting Markov chains in Chapters 4-5, high-probability transitions are represented as thick dark links and low-probability transitions are represented as thin light links. Also, from each demand state, a green link identifies the highest-probability transition; and if different from the green link, a red link identifies the highest probability transition given departure from a given demand state.

Thus, for example, several differences can be noticed between Figure 73 and Figure 74, which themselves represent the difference in demand evolution experts believed would exist if demand itself were fulfilled (in the case of Figure 73) versus not fulfilled (in Figure 74). Whereas the most likely transition from LEO is to a Lunar Surface demand if LEO demand is fulfilled, it is to remain in LEO if that demand is not fulfilled. Whereas the most likely transition from a Venus Orbit demand is to Mars Surface if demand is fulfilled, it is to the less ambitious Lunar Surface mission if that demand is not fulfilled; and similarly, if Mars Orbit demand is not fulfilled, the most likely demand is to continue Mars Orbit missions rather than progress to Mars Surface missions. It might also be noticed that the red link from the Lunar Surface mission (the second most likely next demand) leads to a Near-Earth Object mission rather than a Mars Surface mission in the event that the Lunar Surface demand is not being met in the current period. These examples illustrate the general characteristic of the model that the condition of demand being fulfilled favors progression of demand toward missions aimed at more ambitious destinations that are generally farther away from Earth; conversely, the condition of demand not being fulfilled tends to favor a constancy or sometimes regression of demand toward less ambitious destinations closer to Earth.

Figure 73 and Figure 74 also reveal that the mission destinations of LEO, Lunar Surface, and Mars Surface, and to a somewhat lesser degree Near-Earth Objects, form a set of long-term “sinks” for mission demand in the opinion of the expert participants. In both figures, these destinations have high probabilities of remaining in their present state and also have many high-probability incoming transitions. In contrast, mission demands like Venus Orbit, Sun-Earth L2, and Nothing tend to act almost as transient states for which demand is rare and, when it does exist, is fleeting.

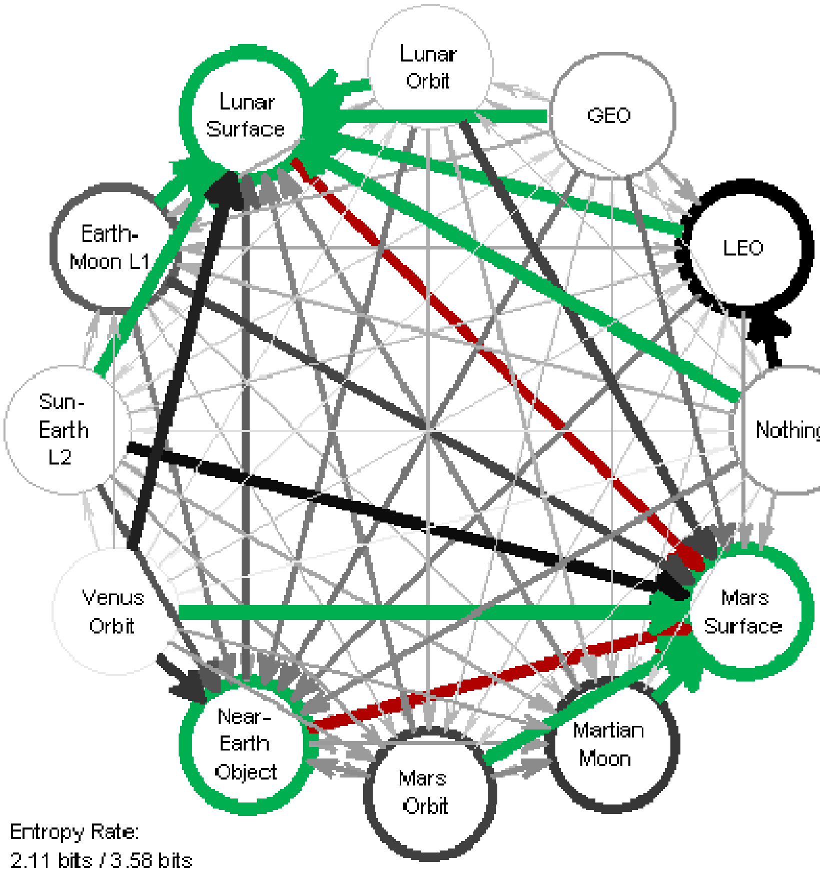

Figure 73. Visualization of the Markov chain of median expert inputs for the condition that current mission demand is fulfilled, with Δt = 8 years. *High-probability transitions are represented as thick dark links and low-probability transitions are represented as thin light links. From each state, a green link identifies the highest-probability transition. If different from the green link, a red link identifies the highest probability transition given departure from that state.*

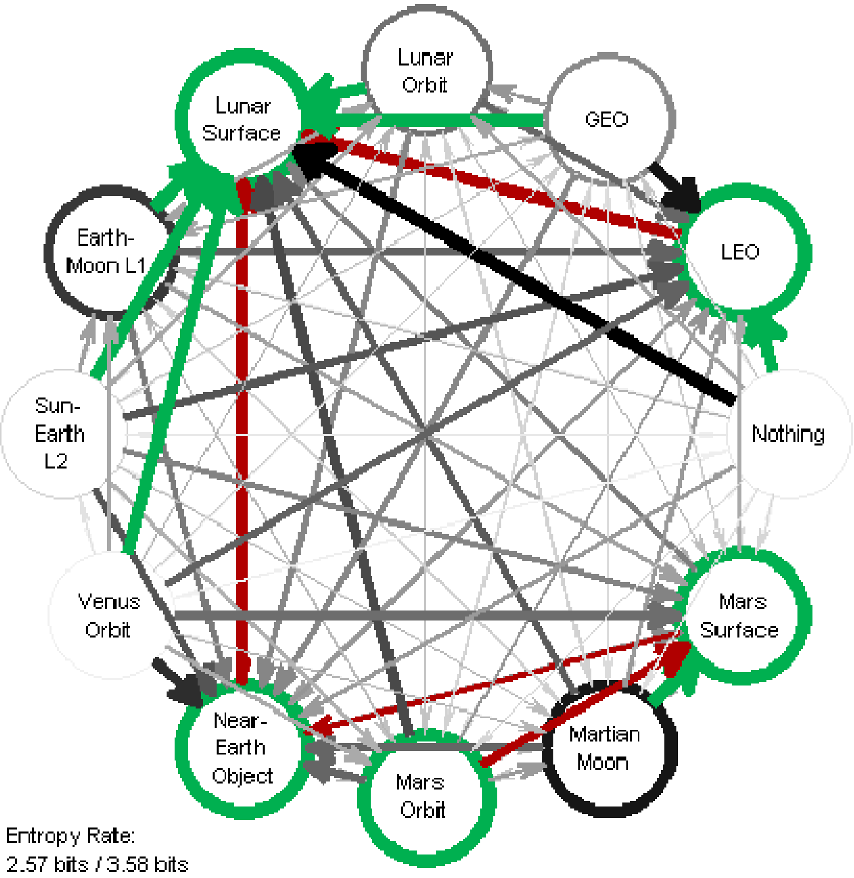

Figure 74. Visualization of the Markov chain of median expert inputs for the condition that current mission demand is *not* fulfilled, with Δt = 8 years. *High-probability transitions are represented as thick dark links and low-probability transitions are represented as thin light links. From each state, a green link identifies the highest-probability transition. If different from the green link, a red link identifies the highest probability transition given departure from that state.*

6.3. Step 3: Define State-Dependent Performance Matrix

With the set of possible engineering configurations defined in Step 1 and the set of mission demand environments defined in Step 2, the role of Step 3 is to link the configuration state to the environment state in each period with one or more quantitative performance measures. Taking the form of a matrix, each measure must inherently accumulate over time to match the formulation of Chapter 4.

6.3.1. Selecting the Performance Metric

Appendix C describes in detail the results of a survey of experts with substantial experience in the field of human space exploration. The first part of this survey requested that participants rate the relative importance of 17 candidate figures of merit for human spaceflight architecture evaluation. Four figures of merit in particular earned both the highest median score and lowest interquartile range (i.e., highest consistency) of scores among the participants: Integrated Program Lifecycle Cost, Total Spending on Production Activities, Date of First Mission to Leave LEO, and Time Between Missions.

In deciding which of these four figures of merit to use in the analysis that follows in Steps 4-5, it may be recalled that total program costs (expressed through the importance of the Integrated Program Lifecycle Cost and Total Spending on Production Activities figures of merit) will already be considered via the transition cost matrices defined in Step 1. Thus, in terms of performance, the relevant metrics to consider including are Date of First Mission to Leave LEO and Time Between Missions. Unfortunately, neither of these metrics is cumulative. For example, the Date of First Mission to Leave LEO metric tracks the occurrence of the single event in a timeline and provides no performance credit for achievements (or even the same achievement) at earlier or later times. For instance, this metric would not distinguish between a timeline involving sustained missions to the Moon in 2020 and a timeline involving a single mission to the Moon in 2020 followed by missions to LEO for the rest of the decade.

In contrast, the Time Between Missions metric is less myopic. The consistency with which it was rated with high importance by the expert survey participants is likely driven by the priority the participants placed on maintaining the skills of the human spaceflight engineering workforce and maintaining public interest through high flight rates. Furthermore, while this Time Between Missions metric itself does not accumulate over time, a surrogate for it does. If this metric is interpreted as an *average* time between missions, then for a given timeline it would be computed as the total number of missions (a cumulative metric) divided by the duration of the time horizon. Throughout the analyses in Steps 4-5, the total time will be fixed at a twenty-year horizon length, and thus a reasonable surrogate for this metric is the total number of missions flown. More specifically, the metric tracked through the performance matrix developed in Step 3 will be the Number of Missions to Demanded Destinations. To account for the fact that decision-makers may wish to place some value on missions flown to non-demanded destinations, an additional "Mission Ratio" figure of merit defined as the ratio of Number of Missions to Demanded Destinations to total missions, along with the Date of First Mission to Leave LEO metric, are considered in the genetic algorithm exploration in Step 5.

6.3.2. Populating the Performance Matrix

Thus, in Step 4 for this framework applied to the NASA human space exploration example, the objectives of interest will be Integrated Program Lifecycle Cost (to be abbreviated as "Total Cost") and Number of Missions to Demanded Destinations. The performance matrix linking the configuration state to the environment state will have dimensions 3,286 rows × 12 columns since there exist 3,286 configurations (defined in Step 1) and 12 environments (defined in Step 2). However, since the number of missions that can be flown to a demanded destination in a given period depends only upon the operations architecture available in that period, an abbreviated version of the performance matrix using as rows the 10 operations architectures will serve for display and explanation for the remainder of this step.

To populate the abbreviated 10 × 12 performance matrix, each of the ten architectures in Table 31 must be compared to mission requirements for each of the twelve mission demand environments. If an architecture has insufficient components to meet the requirements of a mission,[*] it will be assumed that the mission cannot be flown. Conversely, it will be assumed that the operator of the architecture will use the components present to either meet the present demand if possible or, if not possible, maximize the number of missions flown to the presently demanded destination. The assumed mission requirements for the mission demand environments, populated concurrently with those in Table 31, are shown in Table 36. Note that each demand environment is also associated with a particular mission rate: The LEO, GEO, and Lunar Orbit missions are associated with a demanded mission rate of four per period (two per year); the Mars missions are associated with a mission rate of one every two years; and all others except for the Nothing mission demand is associated with a rate of two missions per period (one per year).

---

[*] As will be soon described, substitutions are allowed. For example if an architecture is missing a required small chemical stage but has an extra large chemical stage, the mission can still be performed.

Table 36. Demand Environment Component Requirement Definitions for the Human Space Exploration Application.

		Demand Environment											
		Noth.	LEO	GEO Serv.	Lunar Orbit	Lunar Surf.	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Mars Moon	Mars Surf.
	Representative Mission Rate (missions per 2-year period)	0	4	4	4	2	2	2	2	2	1	1	1
Minimum Required Component Capabilities per Period	1. Crew Launch Vehicles (CLVs)	-	4	-	-	-	-	-	-	-	-	-	-
	2. Heavy Lift Launch Vehicles (HLVs)	-	-	3	4	12	3	8	8	6	4	4	10
	3. Commercial Cargo Launch Vehicles (CCLVs)	-	6	-	-	-	-	-	-	-	-	-	-
	4. Multi-Purpose Crew Vehicles (MPCVs)	-	4	2	4	4	2	2	2	2	1	1	1
	5. Commercial Cargo Logistics Modules (CCLMs)	-	6	-	-	-	-	-	-	-	-	-	-
	6. Small Chemical Stages	-	-	-	-	-	-	2	-	-	-	-	1
	7. Medium Chemical Stages	-	-	3	4	8	3	4	4	6	2	2	1
	8. Large Chemical Stages	-	-	-	-	-	-	-	4	-	2	2	6
	9. Deep-Space Habitation Modules	-	-	1	-	-	1	2	2	2	1	1	1
	10. Lunar Landers	-	-	-	-	8	-	-	-	-	-	-	-
	11. Mars Landers	-	-	-	-	-	-	-	-	-	-	-	2
	12. Multi-Mission Pressurized Rovers	-	-	1	-	8	-	2	-	2	-	2	2
	13. Unpressurized Rovers	-	-	-	-	4	-	-	-	-	-	-	2
	14. Science Rovers	-	-	-	-	4	-	-	-	-	-	-	2
	15. Surface Habitats	-	-	-	-	-	-	-	-	-	-	-	1
	16. Logistics Modules	-	-	2	-	4	2	2	-	2	-	-	-
	17. Power Generation and Storage Units	-	-	-	-	2	-	-	-	-	-	-	1
	18. ISRU Systems	-	-	-	-	2	-	-	-	-	-	-	2
	19. Surface Extravehicular Activity (EVA) Suits	-	-	-	-	20	-	-	-	-	-	-	10
	20. In-Space Extravehicular Activity (EVA) Suits	-	8	8	8	-	8	8	8	8	8	12	-
	21. Supporting Communications/Navigation Satellites	-	-	-	-	1	-	-	-	-	-	-	1
	22. RS-68-Class Engine	-	-	15	20	60	15	40	40	30	20	20	50
	23. J-2X-Class Engine	-	4	6	8	24	6	16	16	12	8	8	20
	24. RL-10B-2-Class Engine	-	-	6	8	56	6	12	28	12	14	14	44
	25. AJ-10-Class Engine	-	4	2	4	4	2	2	2	2	1	1	1

To compute the abbreviated performance matrix in Table 37, an algorithm compares each operations architecture to each set of mission demand requirements on a capability by capability basis. First, launch capability is compared, with the assumption that an architecture possessing more heavy-lift launch vehicles than are required for a given mission demand can use them to fulfill any deficit in crew launch vehicles. Second, commercial cargo logistics module requirements are considered, with the assumption that any of an architecture's logistics modules, MPCVs, and multi-mission pressurized rovers that are not explicitly required for the current mission demand can be readily outfitted to fulfill the role as a commercial cargo module if necessary. Third, chemical stage capability is compared. It is assumed that each large chemical stage not explicitly required by the current mission demand can be used to fulfill the function of a medium or small chemical stage; similarly, it is assumed that each medium chemical stage of an architecture can fulfill the function of small chemical stage. Fourth, extravehicular activity (EVA) suit requirements are checked, with the assumption that surface EVA capability in an architecture can be used to fulfill in-space needs; however, the opposite is not assumed to hold. Fifth, all components other than those listed here are checked on a one-to-one basis with the assumption that no relevant substitutions are available with other components. Given these comparisons, the algorithm identifies whether the full number of demanded missions in the column can be fully achieved with the architecture in the row; and if not, the algorithm determines the maximum integer number of missions that can be flown to the demanded destination using the available components in the architecture.

Although the resulting performance matrix in Table 37 by definition does not take cost into account, it is worth observing that some architectures, such as Architectures 6, 7, and 10 (the Sun-Earth L2, Near-Earth Object, and Deep Space architectures) perform well in a variety of mission demand environments. In contrast, architectures like

Architectures 2 and 4 (the LEO and Lunar Orbit architectures) are highly specialized and are unable to meet demands except in the environments for which they were designed.

Table 37. Performance matrix (abbreviated in rows by operational architecture, rather than configuration) quantifying the number of missions flown to the demanded destination in a given time period.

	Demand Environment											
Operations Architecture	Noth.	LEO	GEO Serv.	Lunar Orbit	Lunar Surf.	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Mars Moon	Mars Surf.
-1- Nothing	0	0	0	0	0	0	0	0	0	0	0	0
-2- LEO	0	4	0	0	0	0	0	0	0	0	0	0
-3- GEO Servicing	0	0	4	2	0	2	0	0	1	0	0	0
-4- Lunar Orbit	0	0	0	4	0	0	0	0	0	0	0	0
-5- Lunar Surface	0	4	0	4	2	0	0	0	0	0	0	0
-6- Sun-Earth L2	0	0	4	2	0	2	2	0	1	0	0	0
-7- Near-Earth Object	0	0	4	2	0	2	1	0	2	0	0	0
-8- Mars Moon	0	0	0	1	0	0	0	1	0	1	1	0
-9- Mars Surface	0	0	0	1	0	0	0	1	0	1	0	1
-10- Deep Space	0	0	4	4	0	2	2	2	2	1	1	0

6.3.3. Populating the Boolean Demand Fulfillment Matrix

An additional useful piece of information may be gathered from the data computed for Table 37. If the data in the table is converted from an integer to a Boolean (i.e., zero or one) representation, the matrix indicates whether a given operations architecture (in the row) fulfills a given demand environment (in the column). The conversion is largely trivial; every zero in Table 37 remains zero in Table 38, and all other elements become unity. The only exception is the first column, which becomes comprised entirely of ones since every architecture has the ability to fulfill the "Nothing" demand. This Boolean demand fulfillment matrix will become necessary in the definition of the modified function in Step 4 that defines the probability of transition among total states; Step 2 introduced configuration-dependent demand, and Table 38 will provide the information needed to allow selection of the proper probability transition matrix (Table 34 or Table 35, respectively a Boolean one or zero in Table 38) as a function of the current total state.

Table 38. Boolean demand fulfillment matrix (abbreviated in rows by operational architecture, rather than configuration).

		Demand Environment											
		Noth.	LEO	GEO Serv.	Lunar Orbit	Lunar Surf.	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Mars Moon	Mars Surf.
Operations Architecture	-1- Nothing	1	0	0	0	0	0	0	0	0	0	0	0
	-2- LEO	1	1	0	0	0	0	0	0	0	0	0	0
	-3- GEO Servicing	1	0	1	1	0	1	0	0	1	0	0	0
	-4- Lunar Orbit	1	0	0	1	0	0	0	0	0	0	0	0
	-5- Lunar Surface	1	1	0	1	1	0	0	0	0	0	0	0
	-6- Sun-Earth L2	1	0	1	1	0	1	1	0	1	0	0	0
	-7- Near-Earth Object	1	0	1	1	0	1	1	0	1	0	0	0
	-8- Mars Moon	1	0	0	1	0	0	0	1	0	1	1	0
	-9- Mars Surface	1	0	0	1	0	0	0	1	0	1	0	1
	-10- Deep Space	1	0	1	1	0	1	1	1	1	1	1	0

6.4. Step 4: Decision Support Analysis

Defined through Steps 1, 2, and 3 have been the set of available configurations and associated transition costs, the set of possible mission demand environments and associated transition probabilities, and the performance accumulated as a consequence of a given configuration operating in a given demand environment. With these components defined, as in the previous examples of Chapters 4-5, there now exists enough information to begin to answer the question of what is the "best" initial configuration and, furthermore, the "best" decision policy the decision-maker can choose.

Figure 75 shows a version of the assumed Figure 25 simulation timeline that has been modified to reflect the architecture triplet definition of a configuration posed in Section 6.1.1. Figure 75 also incorporates the approximation discussed in Section 6.1.3.2.2 regarding the present NASA human space exploration configuration (a LEO-themed architecture in the second phase of development and no relevant exploration architectures in operations or memory) and the approximation that the immediate demand is for LEO missions. These approximations define the initial condition of the configuration and demand environment states at $t = 0$. Also marked in Figure 75 are the two-year time increments up to the 18-year mark. Since the last period is of a two-year length, the total horizon length over which cost and performance measures accumulate will be 20 years (or ten periods). In each period, the decision-maker must decide which configuration (i.e., which architecture triplet) to select such that its implementation might begin at the start of the subsequent period. A second distinction with Figure 25, notated by the wavy arrows in Figure 75 and explained in further detail in Section 6.4.2, is that the translation of a configuration decision to an operational configuration will be modeled as probabilistic to account for the endogenous possibility of schedule slippage.

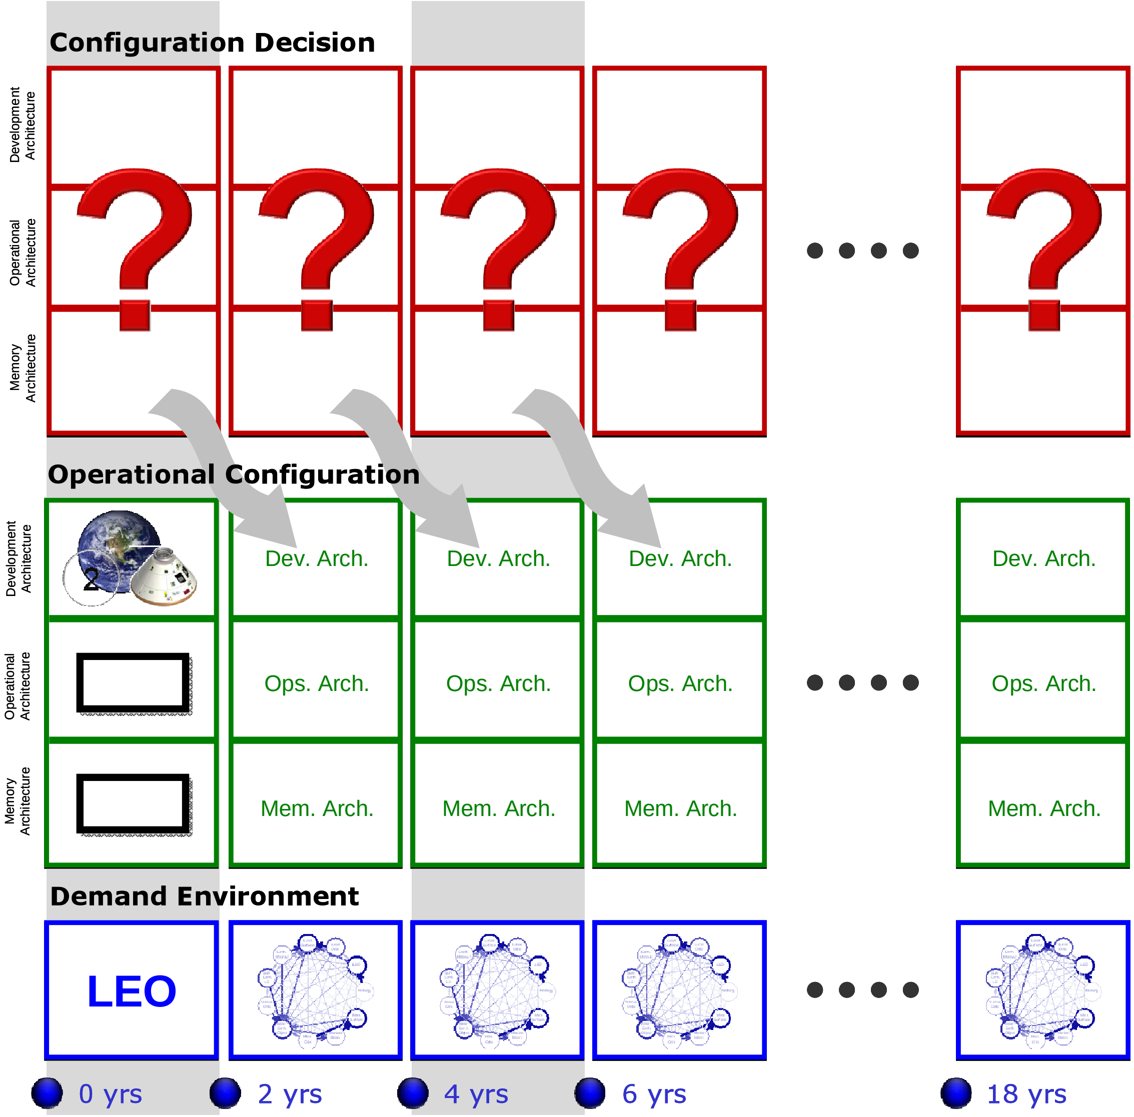

Figure 75. Visualization of Configuration Decisions and Demand Environment Evolution over Multiple Time Periods for the NASA Human Space Exploration Application.

6.4.1. Preclusion of Open-Loop Path Analysis

Unlike the previous examples in Chapters 4-5, as a matter of practicality this example foregoes the explicit identification of Pareto-optimal paths (Step 4A). In the cost matrix developed in Step 1, all rows possessed at least 10 transitions not restricted by the transition rules of Figure 59. With ten periods in the time horizon under consideration, this implies that there exist at least 10^{10} (10 billion) possible full-factorial paths, which is nearly 200,000 times more paths than were considered in the example of Chapter 5. However, this number is itself misleadingly low; some rows in the cost transition matrix have as many 600 allowable transitions, and thus an upper bound on this number is 600^{10} (about 6×10^{27}, or 6 octillion). To rigorously enumerate all these possible paths, all 3286^{10} (about 1.5×10^{35}, or 150 decillion) possible paths would need to be enumerated and then filtered according to the transition rules. Given these computational demands, only the Step 4B analysis option is utilized, employing Markov decision process techniques to preferentially identify a set of Pareto-optimal decision policies.

6.4.2. Timeline Assumptions and Expanded Probability Definitions

As in the previous examples of Chapters 4-5, the timeline depicted in Figure 75 can be modeled as occurring in the following steps (e.g., at each time step of a simulation):

1. Mission demand evolves stochastically according to the Markov chain estimate, conditioned on whether previous demand had been met.
2. The operator of the currently operational architecture attempts to use this architecture to fulfill the new mission demand, earning credit according to the performance matrix.
3. The decision-maker chooses what architectures to develop, operate, and put into memory next, paying according to the cost transition matrix. If schedule slip does not occur, this selection becomes the next-period configuration.

These three steps repeat at each time increment, and cost and performance accumulate at each increment.[*] In the descriptions of these three steps, however, two items have not been substantially discussed in prior chapters' Step 4 coverage and require additional clarification. First, as introduced in Section 6.2, since the demand environment no longer evolves independently of decisions, the probability transition matrix used to describe the evolution of the demand environment is selected as a function of whether a current configuration meets current demand. Second, a model for schedule slippage has been implemented, which describes an assumed endogenous uncertainty. This schedule slippage model is based upon the results of Dubos, Saleh, and Braun [110], who model the probabilistic schedule slippage experience of previous NASA programs as a function of Technology Readiness Level (TRL). For programs characterized by initial TRLs of 6, used in this thesis as an approximation for the initial aggregate system TRL of a human space exploration development program,[†] the regression model of Ref. [110] suggests that mean relative schedule slippage will be 28.5%. For reference, Table 39 reproduces the mean relative schedule slippage results for other TRLs. As a consequence, for a planned eight-year program starting at a TRL of 6, schedule slippage will, on average, account for approximately an additional two years of development. To approximate typical experiences of schedule slippage occurring toward the end of development as components must be aligned in schedule and integrated, this schedule slippage is modeled to occur probabilistically only for configurations in which development is in its final phase. Thus, to match the expected two-year relative schedule

---

[*] In this particular setup, the performance tracked at each period is the performance earned in the current period. The cost tracked at each period is the cost committed for the next period (equivalent to tracking in each period the necessary next-period budget that must be requested, maintained, or paid forward).

[†] This use of TRL 6 assumes adherence to U.S. Government Accountability Office recommended practices for the initiation of space system development. [111]

slippage, a 50% endogenous progression probability (i.e., a 50% probability that development will finish and permit operations) is implemented for configurations in which development is in its fourth and final phase.*

Table 39. Mean relative schedule slippage as a function of TRL, from the model of Ref. [110].

Initial System TRL	Mean Relative Schedule Slippage, percent
4	87.7
5	50.0
6	28.5
7	16.3
8	9.3

In terms of the Total State = {Configuration State, Demand State} unification of the flexibility and Markov decision process frameworks (see Section 4.4.2.2.1), both the configuration-dependent demand evolution and probabilistic schedule slippage can be integrated into the definition of the decision-dependent transition probability: Given any two total states ξ_1 *and* ξ_2 and action *a* taken from state ξ_1, the probability of reaching ξ_2 from ξ_1 in the next time increment is described by Figure 76. The flowchart in this figure illustrates how the combination of the configuration and environment in state ξ_1 determine, based on an expanded version of the the Boolean demand fulfillment matrix in Table 38, whether demand environment transition probabilities from Table 34 or from Table 35 form a basis for the total-state-to-total-state transition probabilites.

---

[*] This implementation clearly demonstrates the ability of the present framework to account for endogenous uncertainties such as schedule slippage, although there exist some impediments to modeling arbitrary schedule slippage time distributions. In particular, the distribution of project completion times is naturally geometric. In this particular application, the authors of Ref. [110] suggest a normal distribution for relative schedule slippage, while the present implementation is necessarily a geometric distribution (but with a matched mean). This limitation can be overcome, however, if enough additional states are added to allow tracking of schedule slippage history.

The flowchart also shows the logic that by default assigns these transition probabilities values of zero unless they involve a state ξ_2 with an associated configuration that matches either the decision or the "slip configuration" from the configuration corresponding to state ξ_1. This slip configuration is a configuration to which one configuration will transition in the event that schedule slippage occurs, and it is predefined for each configuration that has a development architecture in a final development phase. By definition, a configuration's slip configuration will have an identical development and memory architecture; however, since costing assumptions for the final period of development do not involve continued production of the operations architecture (see Section B.3), the operations architecture for the slip configuration is Architecture 1 (the "Nothing" architecture).

The probability associated with transitioning to a slip configuration, or the probability of schedule slip, is defined by the complement of the progression probability (50%, as discussed above). Assuming independence between the endogenous schedule slippage and exogenous demand environment evolution, the total transition probability between two total states ξ_1 *and* ξ_2, given action *a*, is computed as the product of the appropriate exogenous environment transition and endogenous schedule slip or progression probabilities. As a result, Figure 76 provides a means of integrating both exogenous and endogenous uncertainties into the optimal policy solution (and ultimately optimal initial system selection) process by capitalizing on the definition of the total state as the combination of both configuration and environment state.*

---

* This definition could be further capitalized upon, for example, if future data suggest the need to relax the independence assumption. Since the total state contains information about both the configuration and environment state, all that is necessary to relax the independence assumption is an appropriate model for the probability dependence between the exogenous (which tend to be environment-related) and endogenous (which tend to be configuration-related) uncertainties.

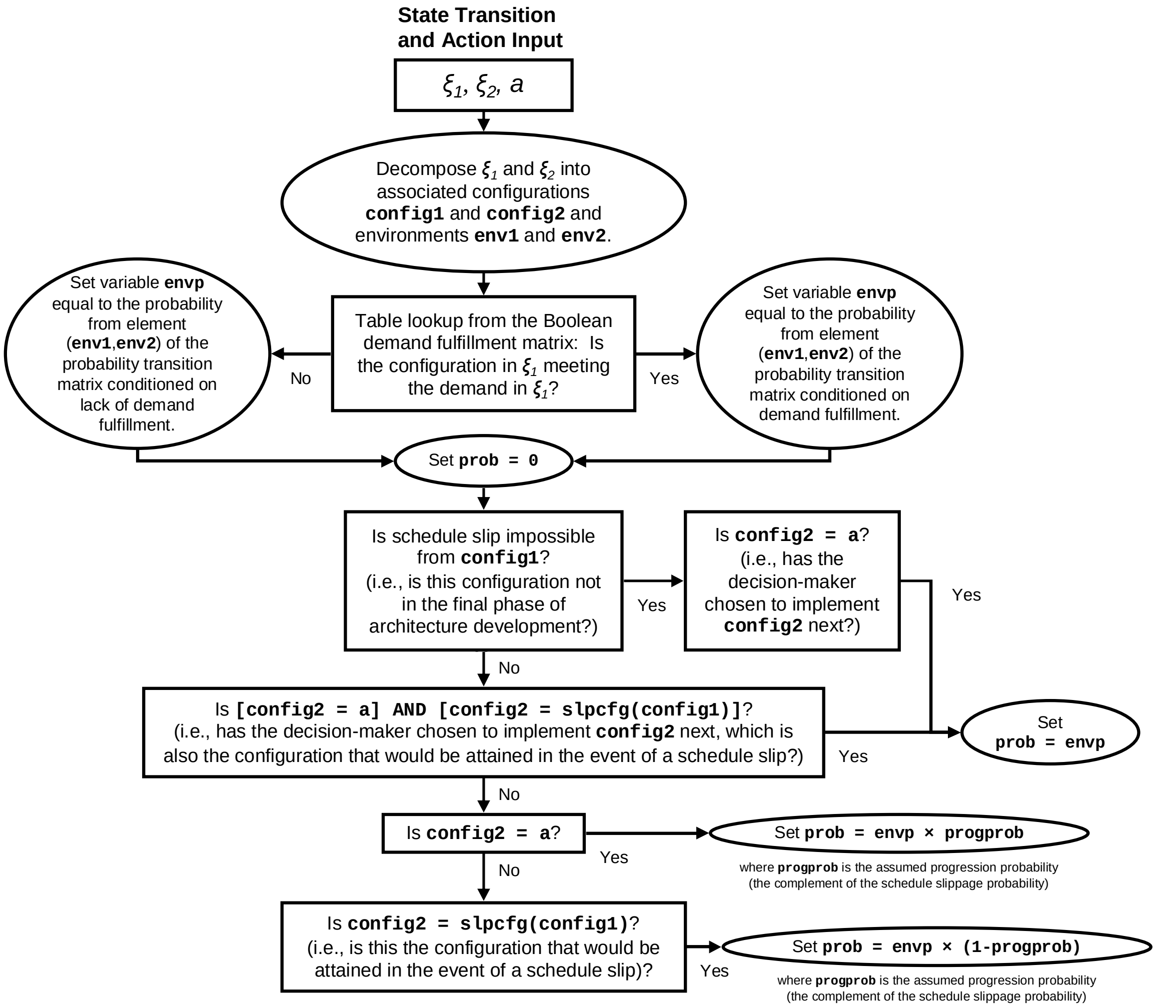

State Transition and Action Input
ξ1, ξ2, a
Decompose ξ1 and ξ2 into associated configurations config1 and config2 and environments env1 and env2.
Table lookup from the Boolean demand fulfillment matrix: Is the configuration in ξ1 meeting the demand in ξ1?
No
Yes
Set variable envp equal to the probability from element (env1,env2) of the probability transition matrix conditioned on lack of demand fulfillment.
Set variable envp equal to the probability from element (env1,env2) of the probability transition matrix conditioned on demand fulfillment.
Set prob = 0
Is schedule slip impossible from config1? (i.e., is this configuration not in the final phase of architecture development?)
Yes
Is config2 = a? (i.e., has the decision-maker chosen to implement config2 next?)
Yes
No
Is [config2 = a] AND [config2 = slpcfg(config1)]? (i.e., has the decision-maker chosen to implement config2 next, which is also the configuration that would be attained in the event of a schedule slip?)
Yes
Set prob = envp
No
Is config2 = a?
Yes
Set prob = envp × progprob
where progprob is the assumed progression probability (the complement of the schedule slippage probability)
No
Is config2 = slpcfg(config1)? (i.e., is this the configuration that would be attained in the event of a schedule slip)?
Yes
Set prob = envp × (1-progprob)
where progprob is the assumed progression probability (the complement of the schedule slippage probability)

Figure 76. Flowchart describing the effective total-state-to-total-state transition probability `prob` used in the Markov decision process solution procedure for the NASA human space exploration application. *Note that, due to the sixth transition rule in Section 6.1.2, the third to last conditional action (using the "AND" statement) is present only for probabilistic completeness. Since stagnation is not an allowed decision option, these two conditions never coincide in the present application; however, if the sixth transition rule were removed, this flowchart would still be valid.*

6.4.3. Computational Resources and Implementation

In the examples of Chapters 4-5, the total state spaces consisted of at most 120 states over 5 time periods. In contrast, the present application involves 3,286 configurations × 12 environments = 39,432 states for a total of 10 time periods. The policy matrices for which the Markov decision process dynamic programming algorithms will be searching thus consist of 394,320 elements, which is of a size nearly 660 times larger than the policy matrices in Chapter 5 (e.g., Table 28). Initial attempts to use the same MATLAB-based computer code as in Chapters 4-5 on the present NASA human space exploration example resulted in run time estimates on the order of 250,000 hours (nearly 30 years!) if executed serially. While execution on several machines in parallel was considered, use of all available MATLAB licenses in the Flight Mechanics Laboratory to which the author was granted access at NASA Johnson Space Center would reduce this time only to 19,000 hours (over 2 years), and order of magnitude improvements beyond this were required.

To solve this computational run time issue, the core MATLAB finite time horizon Markov decision process dynamic programming code was converted to Fortran and utilized the OpenMP interface to enable parallel processing among the multiple threads and processors of a single computer. Parallelization of the code is possible for computations within a given time period, since the action taken from one state at time τ has no effect on the optimal selection of the action from another state at the same time τ. Since, as discussed in Section 4.4.2.2.2, optimizations are performed for a range of weights and objective function powers to ensure satisfactory identification of the Pareto frontier, multiple instances of the code were able to be executed in parallel on each of approximately 40 eight-core, sixteen-thread, 2.93 GHz HP DL360 G6 computing nodes in the Flight Mechanics Laboratory at NASA Johnson Space Center. In the primary results that follow, weights were varied from zero to unity in increments of 0.025, and the objective powers used were 1, 2, 4, and infinity. The time required to execute a full set of these primary runs was approximately 50 hours, a significant improvement (by a factor of 5,000!) from the 250,000 hour estimate prior to the Fortran conversion and parallelization.

6.4.4. Primary Results: The Potential of an Unconstrained Per-Period Budget

As in the previous examples of Chapters 4-5, the dynamic programming solution to the present problem posed as a Markov decision process permits the identification of Pareto-optimal decision policies. However, unlike the previous examples of Chapters 4-5, in which the policies could each be displayed in tabular form on a single page, each of the optimal policies in the NASA human space exploration example is defined by a matrix of nearly 400,000 elements and would require several hundred pages to display. In lieu of identifying policies in this unwieldy form, each policy will be identified simply by the weight placed on cost and the objective function power used to obtain it as a solution (e.g., W0.1-N4 refers to the policy solution to use of a 0.1 weighting on cost, 0.9 weighting on performance, and power 4 objective function). Each such identification number has a single optimal policy associated with it, and the number itself contains some information about the character of the policy solution; for example identification numbers with high weights on cost will be associated with low-cost policies, and those with large objective function powers will have a tendency to fall away from the convex portion of the Pareto frontier. The one exception to this notation will be a notional anticipatory policy, a seemingly sensible but suboptimal policy that will illustrate the benefits of exploring the policy space.

6.4.4.1. Definition of an Anticipatory Reference Policy

Before proceeding to the full results of the MDP policy optimization, it is instructive to consider the time histories of states, costs, and system performance that may be obtained if a reasonable pre-specified policy is run through the simulation described via the steps of Section 6.4.2. As in the previous examples of Chapters 4-5, this policy will be named an anticipatory policy for the reason that it simulates the logic a decision-maker might normally follow to plan for anticipated future demands without the benefit of the techniques proposed by the current thesis. This anticipatory policy will be defined by the three basic rules below, governing conditions for transitioning development, operations, and memory architectures:

1. Development begins or continues for the architecture that most closely corresponds to most likely next demand after a typical development 8-year period. This most likely next demand is obtained from the appropriate conditional probability transition matrix projected over eight years (i.e., the two-year matrices in Table 34 and Table 35 raised to the fourth power). In most cases, the architecture that most closely corresponds to each demand environment shares the same name as the demand environment; based on the component demands, Architecture 3 is assigned as most closely corresponding to the Earth-Moon L1 demand, Architecture 10 is assigned to the Venus Orbit demand, and Architecture 8 is assigned to the Mars Moon demand. To be competitive with the optimal finite-horizon policies for which the MDP algorithm solves, no new development projects are started within four periods of the end of the simulation since these projects will not result in a fielded operations architecture with performance benefits.
2. Operations continue with the previous operations architecture unless the prior configuration involved a development architecture in its final phase, in which case the just-finished development architecture is placed into operation. To be competitive with the optimal finite-horizon policies for which the MDP algorithm solves, costs are reduced by selecting in the final period the action not to continue operations into the next period.
3. Previous memory architectures are retained unless a configuration is in the last phase of development for another architecture, in which case the current operations architecture is placed in memory.

Simulation of this anticipatory policy produces the time history results in Figure 77, shown in a format similar to those used in the state and objective time histories shown for the examples in Chapters 4-5. In the plots on the left in Figure 77, the size of each yellow dot indicates the likelihood of a configuration or demand being in a particular state (on the *y*-axis) at a given time (on the *x*-axis); here, the configuration itself is decomposed into its component architectures for clarity. The plots on the right indicate the evolution of per-period cost and performance metrics.

Starting from the first time step (approximated as the year 2011), all simulations utilize the same initial decision to continue with development of the LEO architecture since all simulations start at the same initial configuration and demand environment defined at the beginning of Section 6.4. This decision is based on the anticipation, from the Markov chain visualization in Figure 74, that lack of fulfillment of the current LEO mission demand will lead to stagnation of the demand environment and continuation of the LEO demand in the future. Development of the LEO architecture continues through its third and fourth phases until, in some simulations, it is fielded as the operations architecture in the year 2017. At this point the first missions to demanded destinations can be flown to LEO, which by 2017 continues to characterize mission demand in 56% of simulations. Due to schedule slippage, fielding of the LEO architecture is delayed in some simulations, and by 2021 the LEO architecture is operational in 70% of simulations.

Once the LEO configuration is fielded and, in many cases, begins to meet mission demand, Figure 73 suggests to the decision maker that the lunar surface mission is the most likely next demand. Thus, in many simulations the lunar surface architecture is developed throughout the early 2020s and fielded in the late 2020s. As anticipated, by 2029 the demand environment has shifted much more in favor of the lunar surface missions, with 33% of simulations exhibiting Lunar Surface demand in 2029 and only 15% of simulations exhibiting LEO mission demand. Over the total ten-year time span of the simulation, the expected total cost of this anticipatory policy is $179.9 billion for an expected 7.1 missions to demanded destinations.

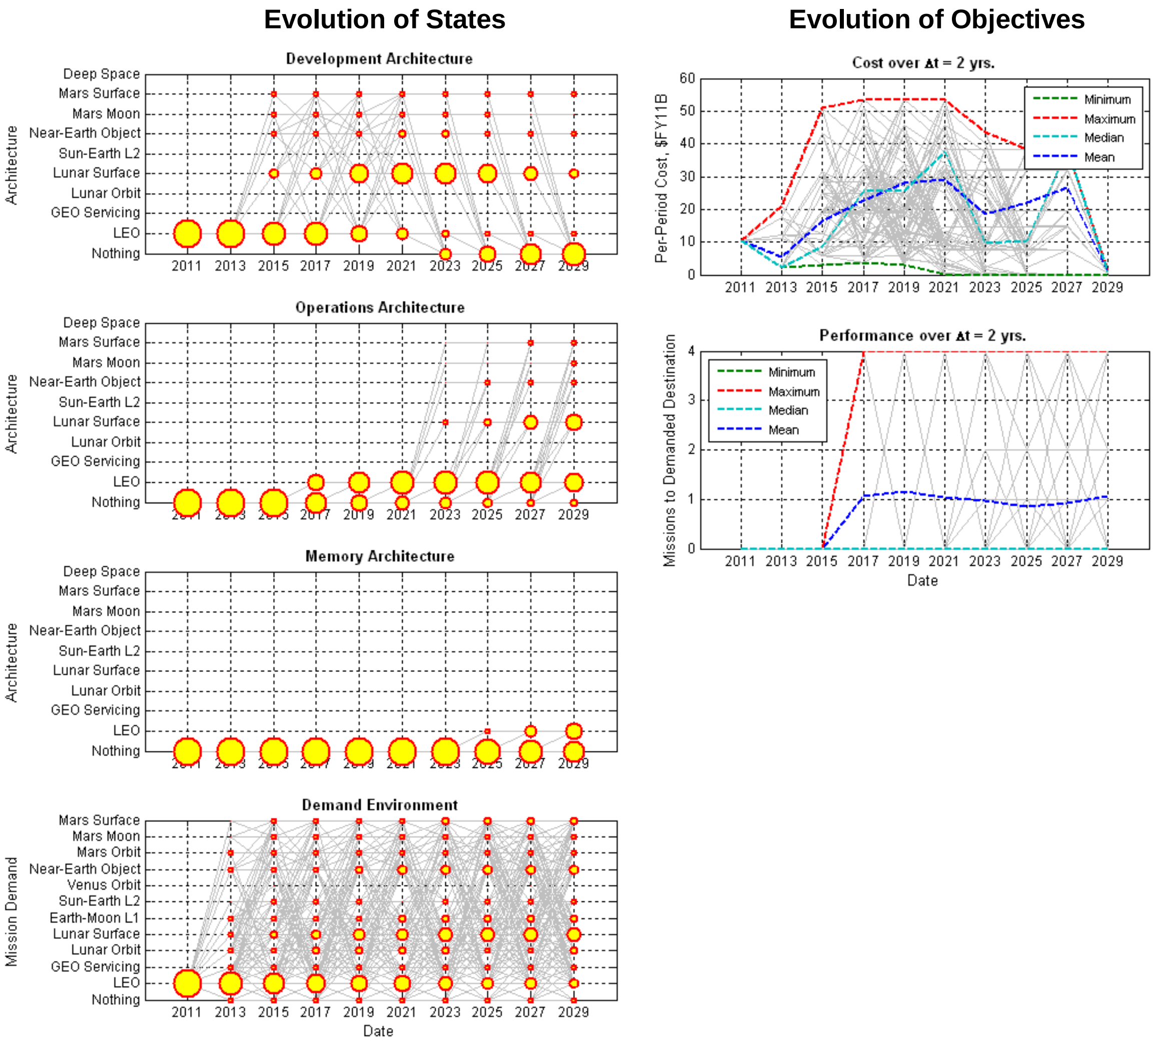

Figure 77. Evolution of states and objectives for the anticipatory reference policy. *In the plots on the left, the size of circles indicates the relative number of Monte Carlo simulation cases that exist in a given configuration or demand environment state (on the y-axes) at a given time (on the x-axes). The configuration state at each time is decomposed into its development, operations, and memory architectures. The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation.*

6.4.4.2. Pareto Frontier of Policies

Figure 78 shows the performance of the anticipatory reference policy, marked as a yellow triangle, in comparison with the performance of the full set of 63 available Pareto-optimal policies obtained from the dynamic programming optimization procedure. The

Pareto frontier in particular is of interest to decision-making because it comprises the set of policies for which performance cannot be increased without increasing cost, or for which cost cannot be decreased without sacrificing performance. Especially interesting on a Pareto frontier are regions of steep or shallow slopes, which indicate regions of compelling trades. Figure 78 shows that, for the NASA human space exploration application, the frontier is nearly linear and quite shallow above a total cost of $50 billion (regressing, with an R^2 value of 0.97, to an average slope of 0.009 missions to demanded destinations per billion dollars added) and substantially steeper below the $50 billion total cost. As a result, optimal performance at the $35 billion total cost level entails an average cost of $6.2 billion per mission to demanded destination, a value that grows substantially to $8 billion per mission at the $50 billion total cost level, $18 billion per mission at the $124 billion total cost level, and $28 billion per mission at the $226 billion total cost level.

The overlay of the anticipatory policy performance on the same plot as the Pareto frontier is of interest because doing so reveals not only that the anticipatory policy is dominated by others discovered in the MDP optimization process, but also that the anticipatory policy is just one of many options; even if it were nondominated, selection of this particular policy carries with it no options regarding cost and performance preferences.

Also marked in Figure 78 are three policies of likely interest to a decision-maker, details of which are provided next. Covered first is policy W0.125-N1, a policy that attains both substantially lower cost and higher performance than the anticipatory policy. Covered next is policy W0.050-N2, the highest-performance (and highest-cost) Pareto-optimal policy in Figure 78. The final policy examined in detail in W0.600-N∞, a policy that has an expected total cost on par with NASA’s current non-International-Space-Station human space exploration budget projected over twenty years.

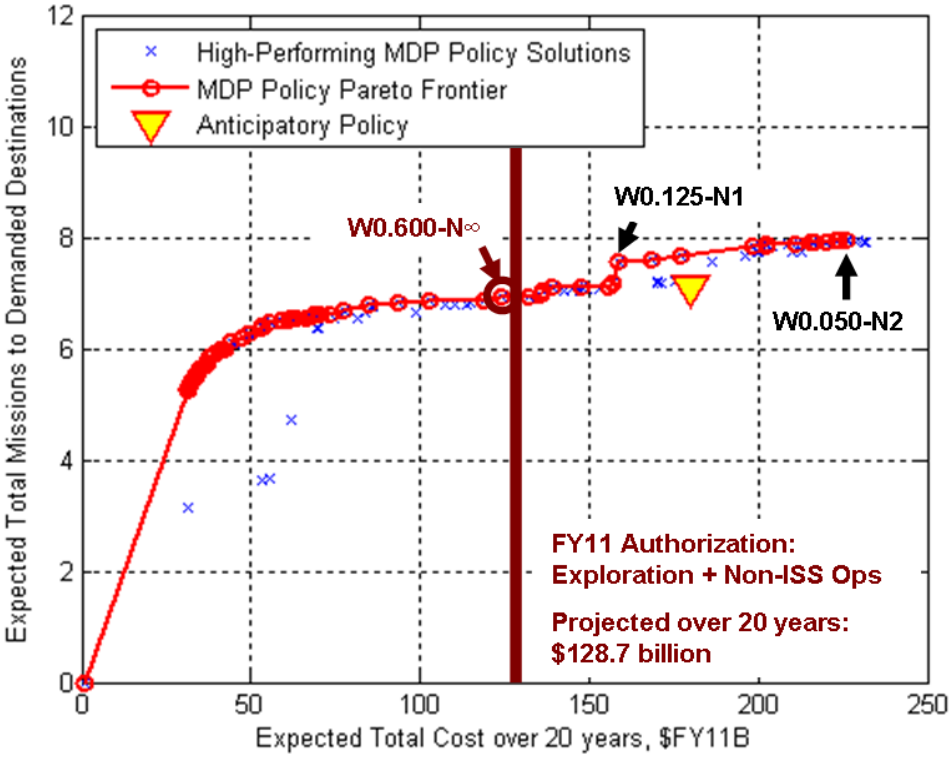

Figure 78. Trade between expected total missions to demanded destinations and expected total cost for MDP policy solutions. *Marked on the plot are three policies of varying long-term cost and performance, as well as a vertical line representing a reasonable long-term budget expectation for human space exploration activities.*

6.4.4.2.1. Policy W0.125-N1: Dominating the Anticipatory Policy

Shown in Figure 79 are the time histories of states and objectives for policy W0.125-N1, which provides 7% more expected performance than the anticipatory reference policy for 12% less expected cost. A clear example of a policy that dominates the anticipatory policy, Figure 79 provides some clue about why this is the case: Figure 79 shows no development of the Mars Surface, Mars Moon, or Near-Earth Object themed architectures, which stands in contrast to the anticipatory policy of Figure 77. Instead, the policy of Figure 79 favors greater focus on development of the core LEO and Lunar Surface themed architectures, with occasional (in less than 6% of simulations) focus on developing the Deep Space architecture starting in 2017, depending on the evolution of the demand environment early during the timeline. Among all simulations, this policy begins operations of the LEO themed architecture no earlier than 2017, the Lunar Sruface themed architecture no earlier than 2023, and the Deep Space architecture no earlier than 2025. The expected total cost of policy W0.125-N1 is $158.8 billion for an expected 7.6 missions to demanded destinations.

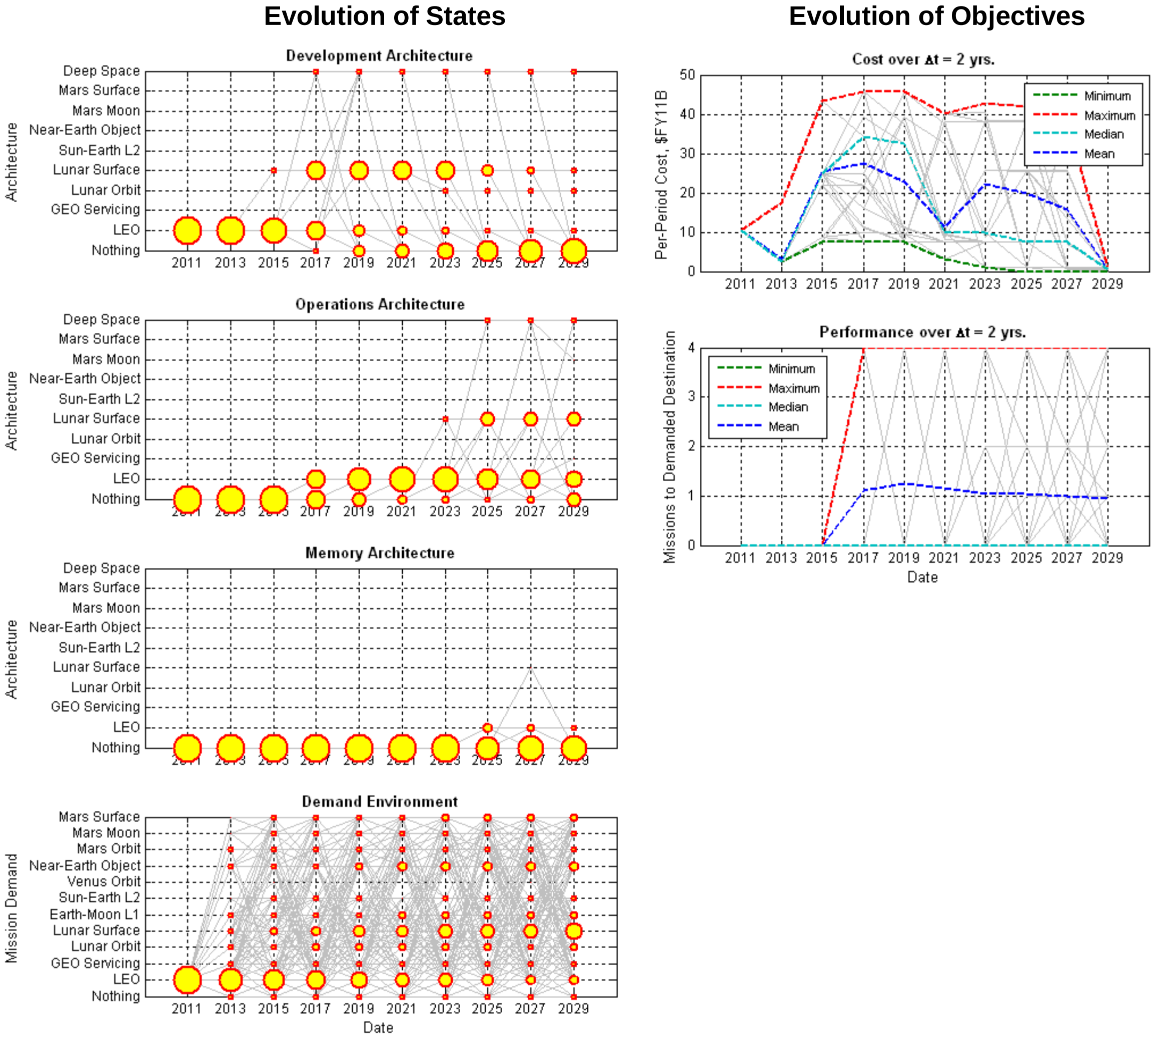

Figure 79. Evolution of states and objectives for policy W0.125-N1. *In the plots on the left, the size of circles indicates the relative number of Monte Carlo simulation cases that exist in a given configuration or demand environment state (on the y-axes) at a given time (on the x-axes). The configuration state at each time is decomposed into its development, operations, and memory architectures. The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation.*

6.4.4.2.2. Policy W0.050-N2: Maximizing Performance

Also of some interest in Figure 78 is the highest-performance policy at the upper right of the Pareto frontier. This policy, W0.050-N2,[*] has state and objective time histories shown in Figure 80. A variation on Figure 79, policy W0.050-N2 shares the predominant characteristic of continued LEO architecture development, followed by fielding of the same LEO architecture in 2017 and subsequent development and fielding of the Lunar Surface themed architecture as early as 2023. Compared to policy W0.125-N1, this policy exhibits a greater focus on developing the Deep Space architecture, starting as early as 2015 with operations starting as early as 2023. By 2029, the Deep Space architecture is operational in 8% of simulations, the Lunar Surface themed architecture is operational in 68% of simulations, the LEO themed architecture is operational in 23% of simulations, and no architecture (the “Nothing” architecture) is operational in the remaining 1% of simulations. Interestingly, in the year 2023 the Deep Space architecture is used in 2% of simulations as a starting point for development of the GEO Servicing themed architecture, and at the same point in 8% of simulations the Lunar Surface themed architecture is used as a starting point for development of the Lunar Orbit themed architecture; however, these architectures never see operation because of their development late in the simulation. The expected total cost of policy W0.050-N2 is $226.1 billion for an expected 8.0 missions to demanded destinations.

The results associated with this highest-performance policy also reveal some fundamental insights regarding the evolution of human space exploration capabilities and

---

[*] As is evident in Figure 78, there is significant clustering of candidate MDP policy solutions near the maximum-performance point of the frontier, and all have nearly identical performance. The reason for a 0.05 weighting rather than a 0.00 weighting achieving the distinction of the maximum-performance policy can be reasonably attributed to numerical sensitivity associated with using a Monte Carlo simulation to generate cost and performance results for the near-equivalent policies in this region.

architectures. First, the fact that the policy of Figure 80 maximizes performance indicates that, despite the materialization of a substantial Mars Surface mission demand in many (at least 10% of) simulations, development and fielding of the Mars Surface themed architecture is never an optimal use of time to maximize the number of missions flown to demanded destinations. This is due in part to the low mission rate of the Mars Surface mission, in part to the existence of the Deep Space architecture as an option, and in part to the transience of the Mars Surface mission demand: As detailed in Table 37, the maximum number of missions that can be achieved in a given time period in an environment of Mars Surface mission demand is one. If, once a Mars Surface mission demand materializes, the time that might intuitively be spent developing a Mars Surface architecture is instead spent developing the Deep Space architecture, up to four times as many missions could be flown per period in the reasonably likely event that future demand shifts to a different mission before development finishes.[*] This example thus illustrates that the transience (or stability) of a current mission demand is an important consideration in system decision-making.

Second, both policies W0.050-N2 and W0.125-N1 respond to strong demands for LEO and Lunar Surface missions with, predominantly, sequential development and operation of LEO and Lunar Surface themed architectures. The Lunar Surface themed architecture development is able to capitalize upon the previous development of the Multi-Purpose Crew Vehicle (MPCV) from the LEO-themed architecture to reduce development costs. This progression is intuitive but is also supported by the fact, as illustrated in Table 37, that the Lunar Surface themed architecture has the ability to operate missions in the LEO demand environment as well as the Lunar Surface demand environment.

---

[*] From Figure 74 and from the matrix of Table 35 raised to the fourth power, the 8-year (4-period) probability of remaining in the Mars Surface demand environment while demand is not being fulfilled is approximately 50%.

Third, as alluded to in the first point, this performance-optimal policy involves notable development of the Deep Space architecture (by 2029, operational in 8% of simulations) to allow non-LEO and non-Lunar-Surface mission demands to be met in cases where demand for such missions can be planned for with reasonable confidence (for example, in cases where demand early in a simulation evolves to an ambitious deep space mission and is likely to remain within the family of deep space missions). These non-LEO and non-Lunar-Surface mission demands particularly tend to occur toward the end of simulation timelines, collectively with a high probability; however, since no single deep space mission carries enough probability to justify development of a dedicated architecture, development of the Deep Space architecture provides a means of meeting mission demands for a wide variety of deep space mission expectations. To a substantial degree this idea is similar to that of the Review of U.S. Human Spaceflight Plans Committee (Augustine Committee); however, it is worth emphasis that in Figure 80, the presence of the Deep Space architecture is notable but not predominant. This optimal-performance policy calls for the Deep Space architecture's development only in special situations and, as mentioned in the previous paragraph, calls predominantly for the development of LEO and Lunar Surface themed architectures.

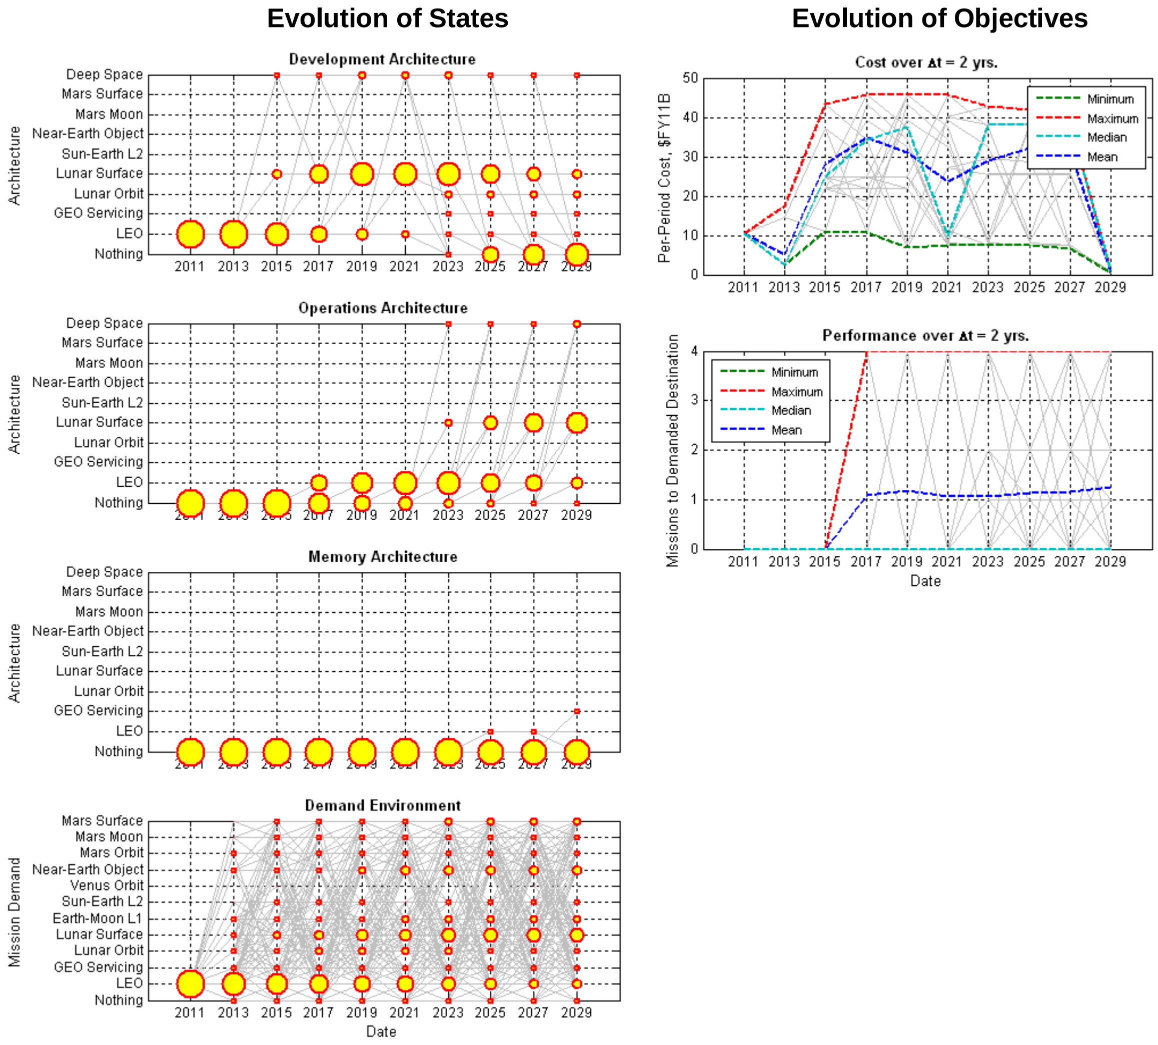

Figure 80. Evolution of states and objectives for policy W0.050-N2. *In the plots on the left, the size of circles indicates the relative number of Monte Carlo simulation cases that exist in a given configuration or demand environment state (on the y-axes) at a given time (on the x-axes). The configuration state at each time is decomposed into its development, operations, and memory architectures. The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation.*

6.4.4.2.3. Policy W0.600-N∞: Matching an Expected Long-Term Budget

A final policy of particular interest from Figure 78 involves an expected 6.9 missions to demanded destinations and lies at a total expected cost of $124.1 billion, just below the $128.7 billion budget that is obtained when the NASA FY11 authorization for exploration plus non-International-Space-Station operations [109] is projected over a 20-year time span. The history of states and objectives from this policy, shown in Figure 81, reveals that the Pareto-optimal policy for this budget involves substantially less focus on non-LEO architecture operations, although there is an increase in focus toward development and operations of the Deep Space architecture that also translates into operations of the Near-Earth Object architecture (since this is a subset of the Deep Space architecture). Interestingly, coupled with the observations in Section 6.4.4.2.2, this would suggest that the case for the Deep Space architecture and the Augustine Committee's flexible path recommendation becomes stronger at lower budgets.

Predominantly, however, it should be emphasized that the shift in development focus following completion of LEO-themed architecture development is toward no development project at all. As a result, the majority of simulations describe a scenario in which development of the LEO-themed architecture is completed and the same LEO-themed architecture is subsequently operated for the remainder of the timeline. By 2029, the Deep Space architecture is operational in 15% of simulations, the Near-Earth Object themed architecture is operational in 2% of simulations, the Lunar Surface themed architecture is operational in 10% of simulations, the LEO themed architecture is operational in 67% of simulations, and no architecture is operational in 6% of simulations. In a negligible 0.2% of simulations each, the Sun-Earth L2 and GEO Servicing architectures were also operational.

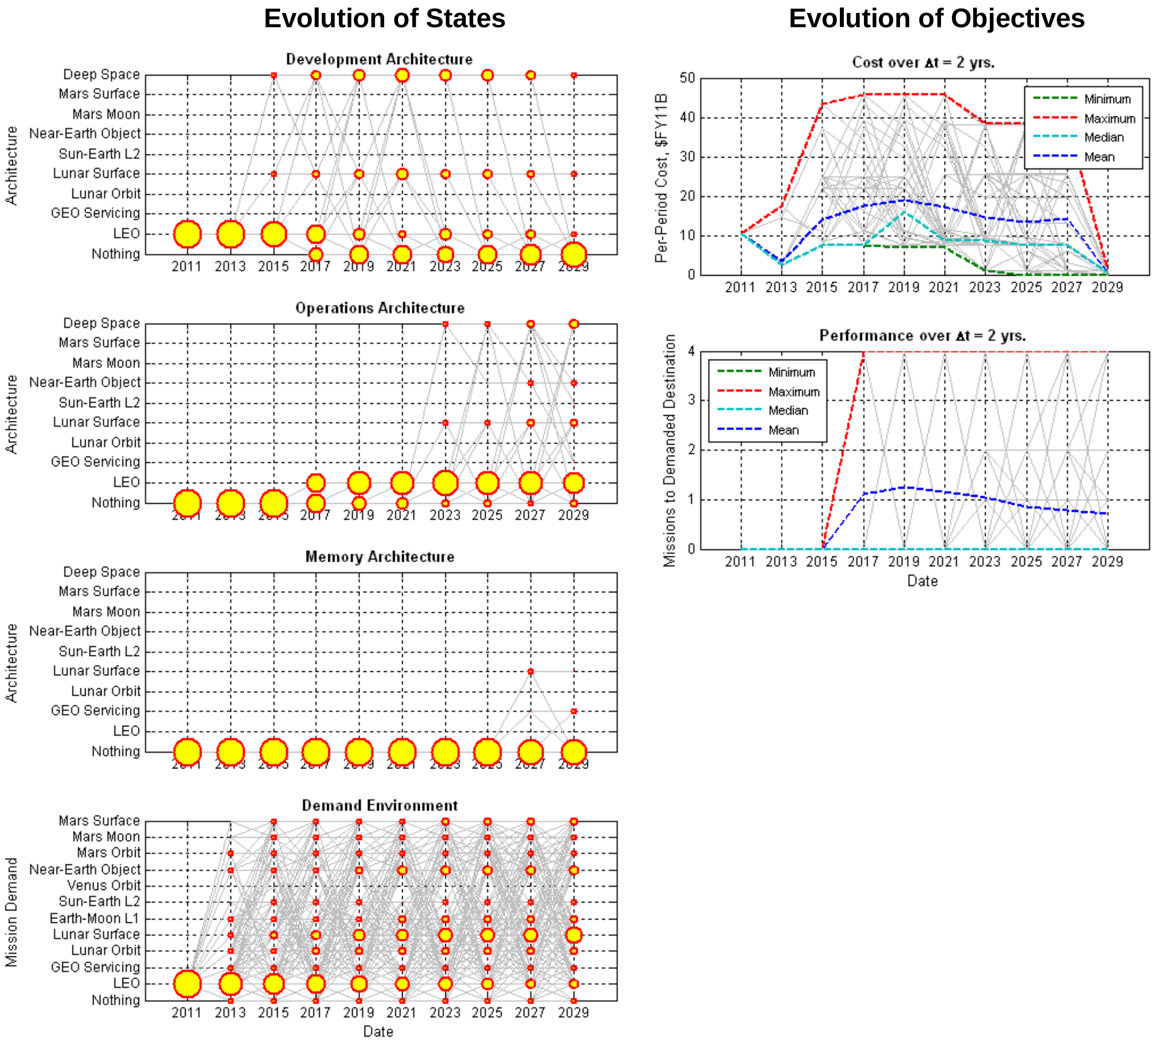

Figure 81. Evolution of states and objectives for policy W0.600-N∞. *In the plots on the left, the size of circles indicates the relative number of Monte Carlo simulation cases that exist in a given configuration or demand environment state (on the y-axes) at a given time (on the x-axes). The configuration state at each time is decomposed into its development, operations, and memory architectures. The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation.*

6.4.5. Implications of a Per-Period Budget Constraint

As introduced in Section 6.4.4.2.3, W0.600-N∞ presents a policy option with an expected long-term cost commensurate with a reasonable long-term NASA human space exploration budget expectation. However, examination of Figure 81 produces the disconcerting revelation that in many simulations (including in the mean and median cost profiles across all simulations), the $12.9 billion per-period budget assumed to be allotted to the agency for this human space exploration program is breached. As a result, even though this policy has long-term costs that match agency budget expectations, *per-period* budget constraints make this policy unreasonable.

To ascertain whether there exist policies that do not overspend on a per-period basis, a seventh transition rule is added to the existing set of six (see Section 6.1.2). This rule prevents a configuration transition from being considered if it costs more than the $12.9 billion per-period budget. As a result, 59% of the transitions previously possible in the cost transition matrix discussed in Section 6.1.2 are no longer allowed.

The Pareto-optimal policies that result for this new, constrained condition are summarized by performance in Figure 82. Most noticeable in Figure 82, in comparison with Figure 78, is the limited extent of the Pareto frontier. While the frontier in Figure 78 extends well past $200 billion expected costs and to expected missions to demanded destinations numbering near 8, the frontier in Figure 82 extends to under $70 billion and expected missions to demanded destinations numbering less than 6.6. If the maximum-performance and maximum-cost point on the new frontier is compared to policy W0.600-N∞ from Section 6.4.4.2.3, which is targeted for spending at the expected long-term budget level, Figure 82 illustrates a performance gap of 0.3 expected missions and moreover, a cost gap of about $60 billion. This cost gap is of particular interest: At an approximately constant budget of $12.9 billion per period, $129 billion will be allocated and presumably spent on human space exploration programs over the ten-period simulation. However, the existence of the per-period budget constraint results in a situation whereby no Pareto-optimal policies exist that spend the entire $129 billion budget; that is, while it is certainly possible to identify *inefficient* policies for spending these funds toward the goal of accumulating missions to demanded destinations, there exist policies that achieve the same performance at lower total costs. The cost gap of $60 billion pointed out in Figure 82 thus highlights the total amount of funds that would be used inefficiently by virtue of a constant use-or-lose $12.9 billion per-period budget.

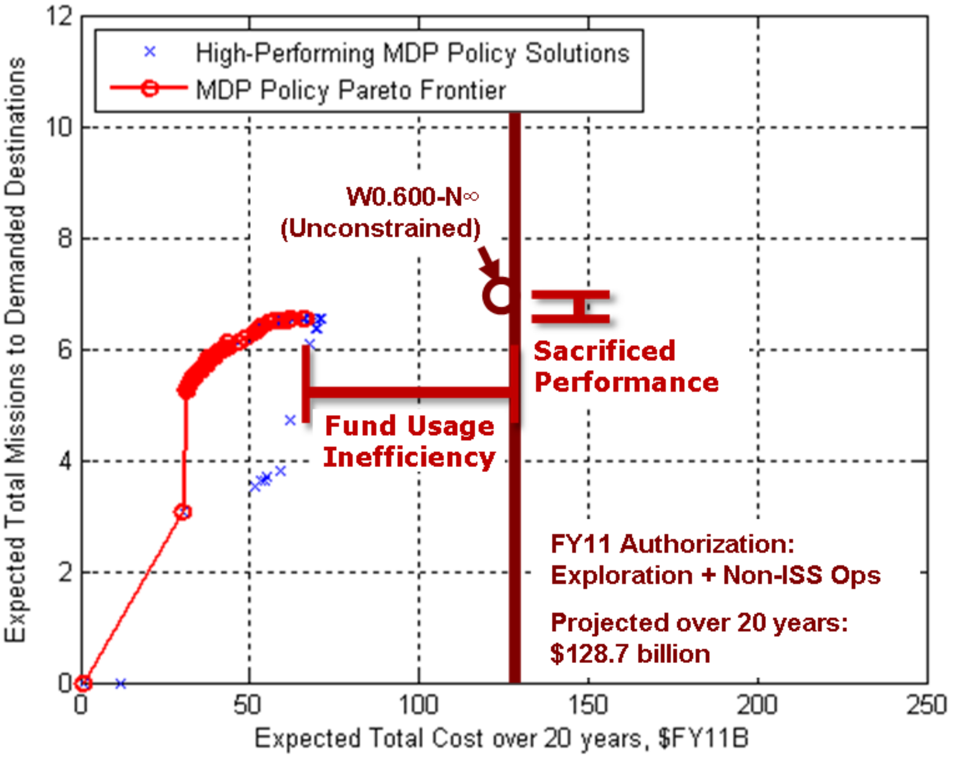

Figure 82. Trade between expected total missions to demanded destinations and expected total cost for MDP policy solutions subject to a $12.9 billion per-period budget constraint. *Marked on the plot is the long-term budget expectation and W0.600-N∞ policy from Figure 78, which serves as a reference for assessing the impacts of the per-period budget constraint.*

What do time histories of states and objectives for policies on this new Pareto frontier look like? As an example, plotted in Figure 83 are the probabilistic time histories for the maximum-performance point in Figure 82, corresponding to a $66.5 billion expected total cost and an expected 6.55 missions to demanded destinations. The plots on the left in the figure illustrate the magnitude of the restriction that the per-period cost constraint places on system development: In contrast with the time histories displayed throughout Section 6.4.4.2, which focus initially on LEO-themed architecture development but then diversify to Lunar-Surface-themed and other architecture development projects, the optimal performance available in the case of the per-period budget constraint concludes development of the LEO-themed architecture and replaces it with no development project at all. Instead, the optimal action (from a now very limited set of actions) is found to be to devote available budget resources toward continuation of

LEO-themed architecture operations into the foreseeable future. In doing so, demand for exploration beyond Earth orbit cannot be met, but the substantial probability of LEO mission demand that exists even ten periods into the future allows substantial accumulation of missions to this demanded destination. Moreover, as the plots on the left in Figure 83 illustrate, per-period spending remains in all simulations and in all time periods below the critical $12.9 billion cap.

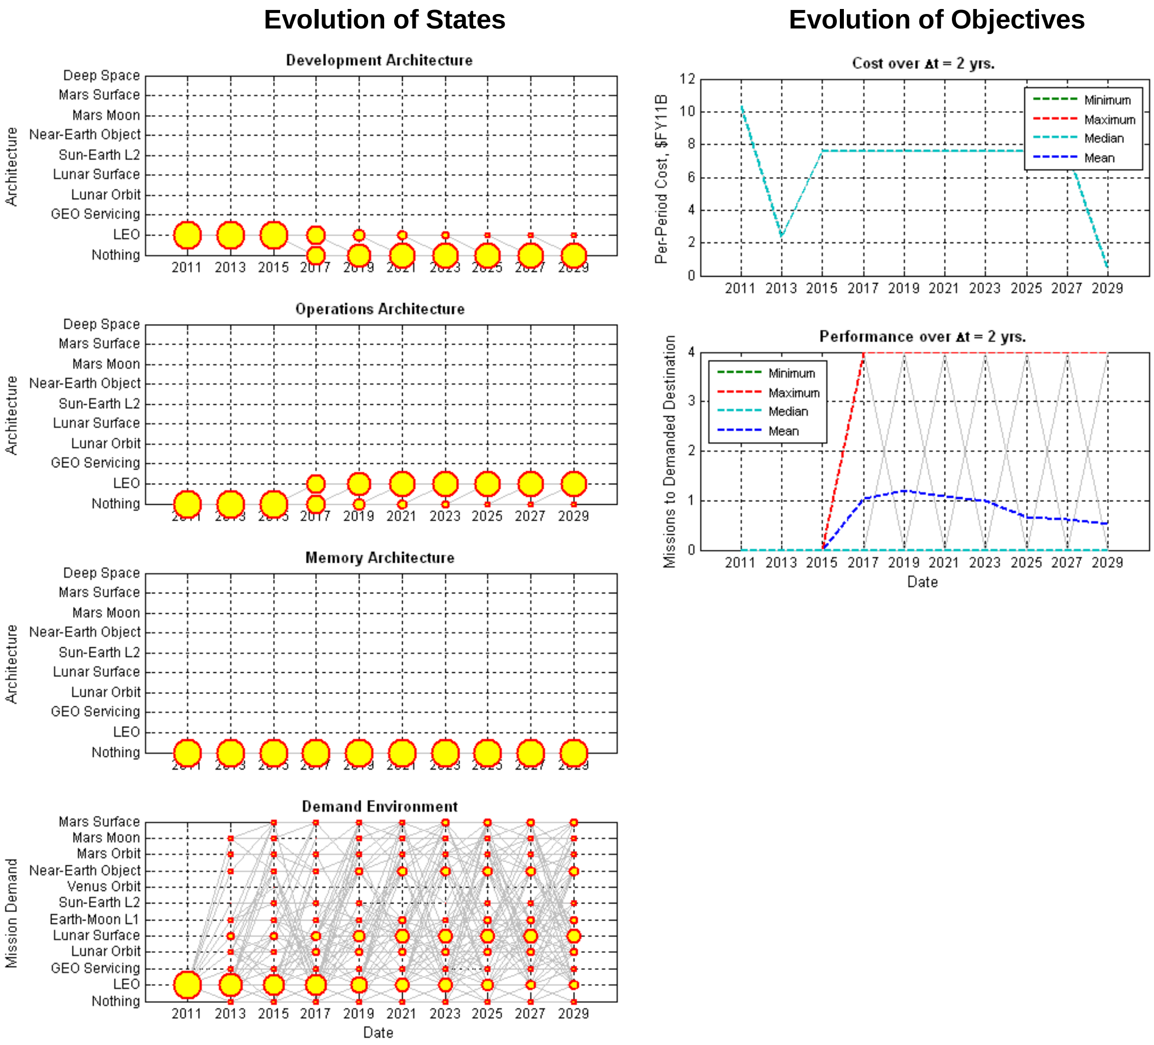

Figure 83. Evolution of states and objectives for maximum-performance policy in the presence of a per-period cost constraint. *In the plots on the left, the size of circles indicates the relative number of Monte Carlo simulation cases that exist in a given configuration or demand environment state (on the y-axes) at a given time (on the x-axes). The configuration state at each time is decomposed into its development, operations, and memory architectures. The plots on the right indicate the associated evolution of per-period cost and performance. In all plots, gray lines indicate transitions made in at least one simulation.*

6.5. Step 5: Implications for Initial System Selection

While the analysis of Step 4 has produced a large set of important and necessary data on optimal policies to follow for the *entire* system time horizon, the most relevant information to a decision-maker from this data set is likely to be the optimal decision to make at the *initial* time step.* To address this question, Step 5 builds upon the analysis results of Step 4 to provide tools and data to support this decision.

6.5.1. Implications based on the Expected-Value Pareto Frontier

As described in Sections 4.5.1 and 5.5.1, the initial decision implied by a policy is identified by locating the initial condition state in the row of the policy matrix and examining the element in the first column. In the case of the NASA human space exploration application, the initial condition corresponds to Config. 3 (a LEO-themed architecture in the second phase of development and no relevant exploration architectures in operations or memory) and the approximation that the immediate demand is for LEO missions. This converts to Total State 3,289 of the 39,432 total states that the system can take at any given time. Since the optimal policy depends upon a decision-maker's relative cost vs. performance preference along the Pareto frontier of Figure 78, the optimal initial configuration decision is a function of an appropriate coordinate along the Pareto frontier. The initial configurations thus found from the Pareto-optimal policies in Figure 78 are identified in Figure 84. In this figure, each initial configuration solution is

---

* In cases where the time step of interest is very short compared to the time required to conduct the analysis suggested by this framework, decisions over multiple future time steps may have particularly great value. In emphasizing the likely interest in the initial decision over others, it is assumed that the time required to implement this analysis (e.g., to assemble all tools, gather all cost, probability, and performance data, run the dynamic programming optimization codes, and analyze all results; likely on the order of weeks or months, depending on the availability of data, number and experience of personnel implementing the process, and number of configurations and environments considered) is shorter than the time step of interest.

decomposed into its architecture triplet components, which are displayed on the *y*-axes of the three separate plots of Figure 84 and identified by the policy's expected total cost on the *x*-axis. Since 63 Pareto-optimal policies exist in Figure 78, 63 yellow circles exist in each of the three plots of Figure 84 to identify the architectures corresponding to each policy's initial configuration. In the case of the development architecture, each yellow circle contains a number identifying the next phase of development selected for the architecture.

As Figure 84 makes evident, from the present configuration and demand environment for the human space exploration application there exists an initial configuration decision (for the next two-year time increment) that is consistent for nearly all long-term cost and performance preferences. This initial configuration involves continuation into the third phase of development of the LEO-themed architecture and operations and retention in memory of no architecture (since none exists yet to operate or retain in memory). Considering in combination (1) that the stochastic demand environment model of Figure 74 indicates substantial stability of the initial LEO mission demand and (2) that the performance matrix of Table 37 indicates the LEO-themed architecture can permit a high number of missions to be flown to LEO in response to this demand, this initial decision to continue LEO-themed architecture development rather than incur the termination liability penalty of switching to a different and likely lower-performing architecture makes sense. The only initial decision that differs exists at a small $660 million cost, which involves cancellation and payment of termination liability for the remainder of the LEO-themed architecture development in favor of developing no architecture at all.

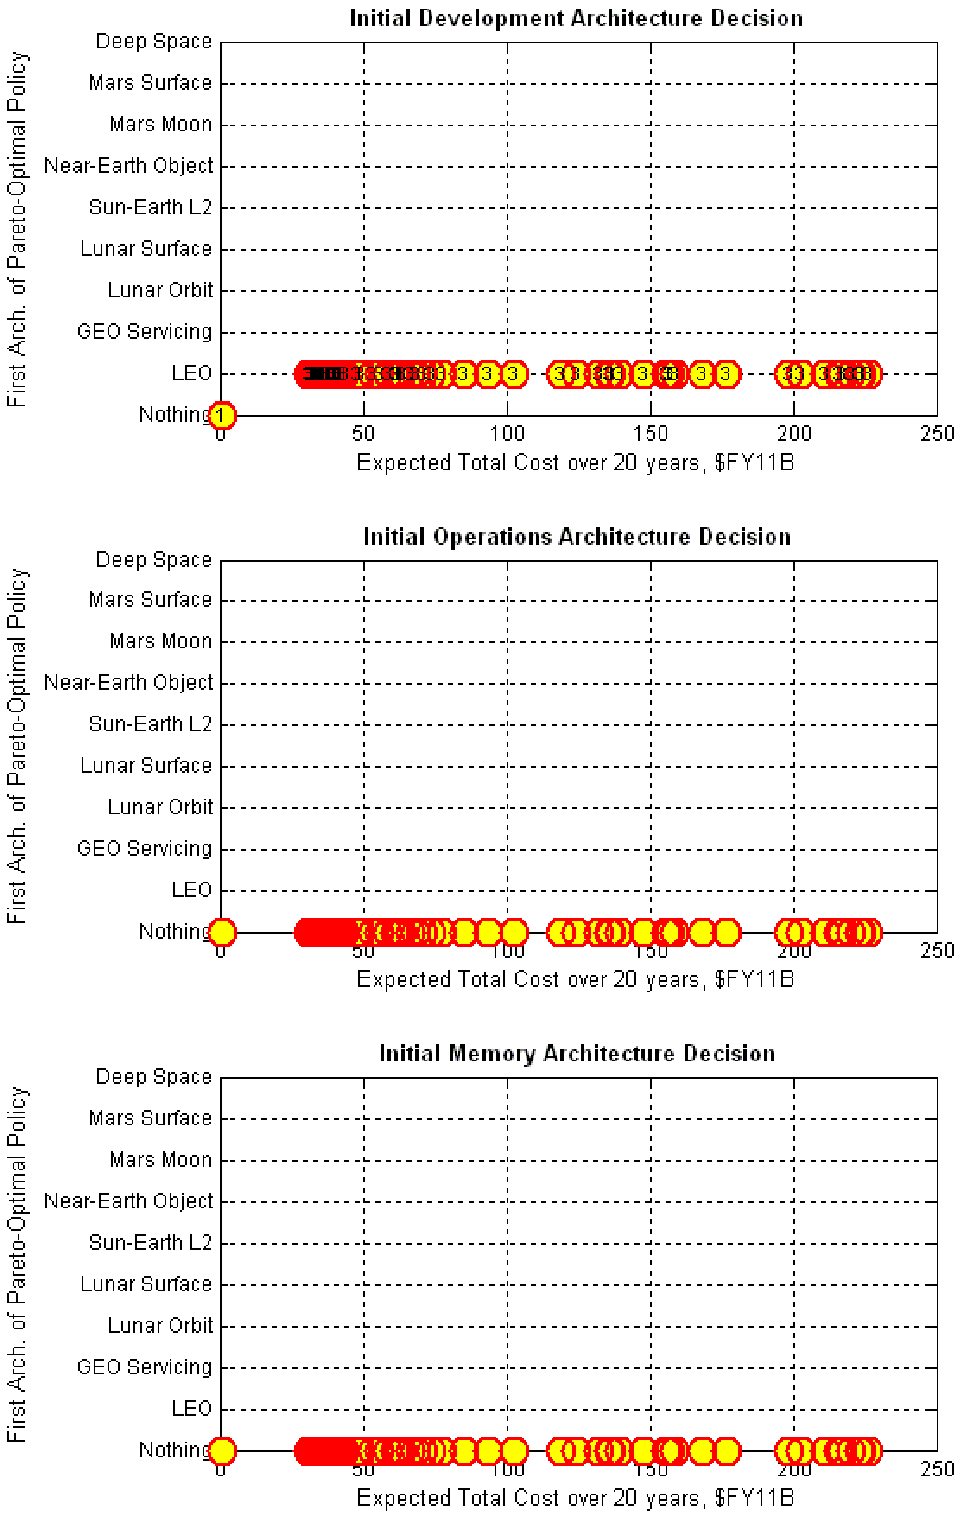

Figure 84. Initial architectures of configurations implied by Pareto-optimal policies as a function of expected policy total cost.

6.5.1.1. Cost-Constrained Policies

For completeness, Figure 85 shows initial configuration decisions implied by the per-period cost constrained Pareto-optimal policies of Figure 82. As might be expected from the example time histories shown in Figure 83, these initial decisions in almost all cases favor the continuation of LEO-themed architecture development. As in Figure 84, the one exception to this is the lowest-cost (and lowest-performance) option, which involves cancellation and payment of termination liability for the remainder of the LEO-themed architecture development in favor of developing no architecture at all. The implications of this result for the favored continued development of a LEO-themed configuration over most cost and performance preferences are thus nearly identical to those seen in the unconstrained results of Figure 84.

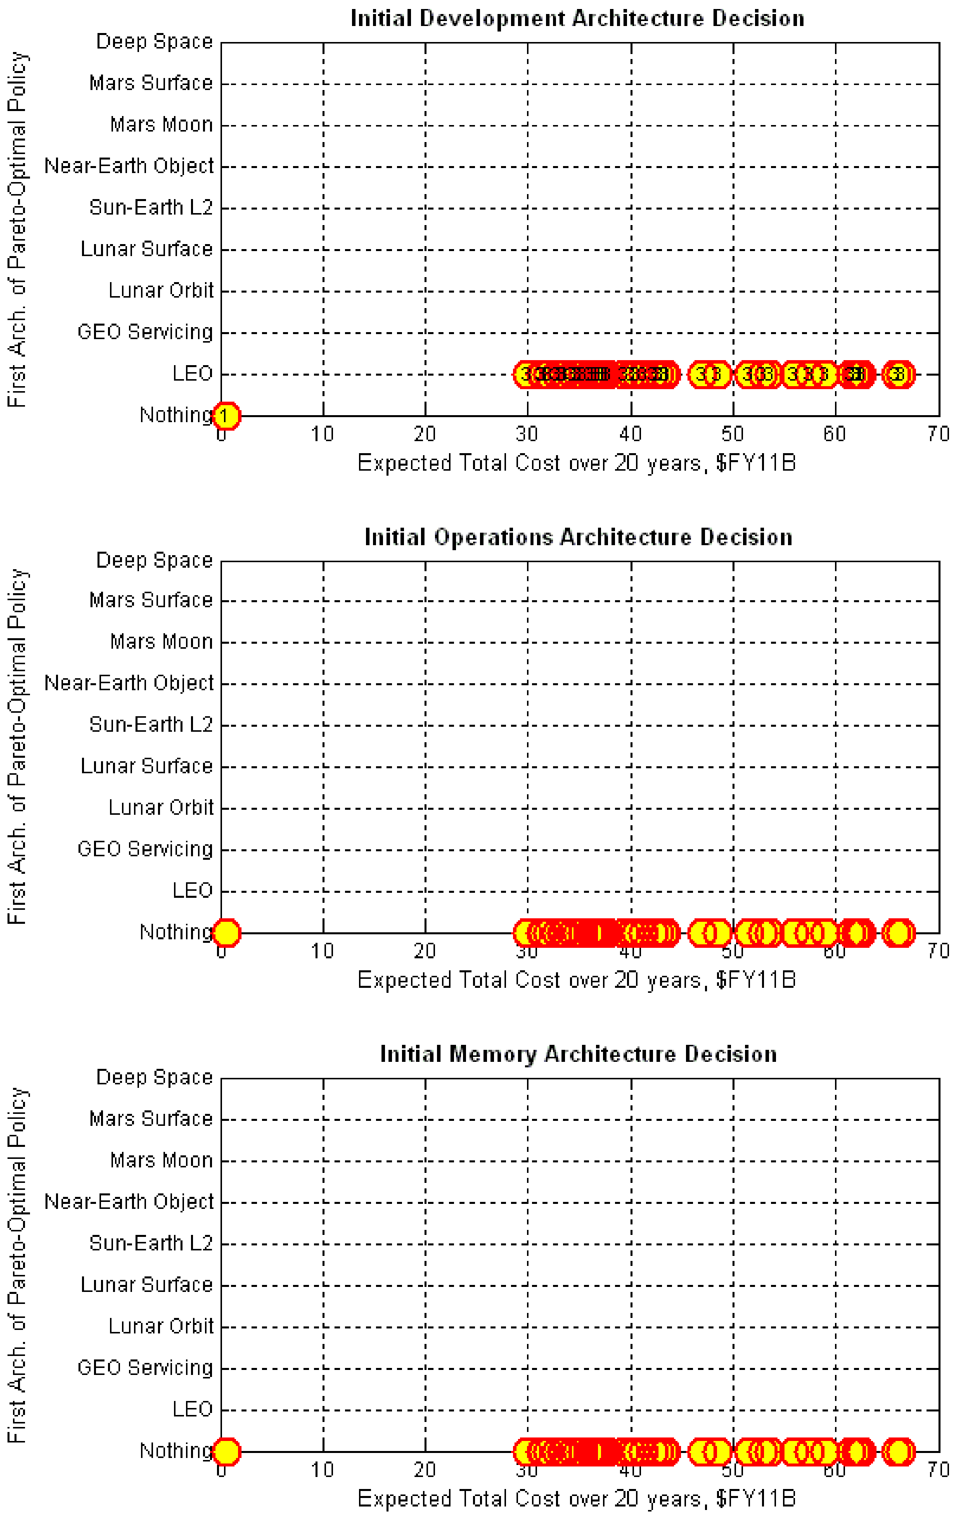

Figure 85. Initial architectures of configurations implied by Pareto-optimal policy solutions subject to a $12.9 billion per-period cost constraint as a function of expected policy total cost.

6.5.1.2. Cost-Constrained Policies using a Nondominated Expert Demand Model

Another relevant question regarding the results of Figure 84 and Figure 85 is whether the same initial decisions (predominantly to continue LEO-themed architecture development) are optimal under different assumptions for the demand environment model, which had been based upon a central tendency of expert probability estimates and may be a source of uncertainty. To address this question, the central tendencies of probability estimates from a particular subset of the original expert population are used to produce a new model and a corresponding new set of optimal policies. As discussed extensively in Appendix C, the experts chosen for inclusion in this subset are those who qualified as non-dominated within the total set of survey participants based on their number of years of experience in the four relevant experience metrics of interest.

When the model based on this uniquely experienced set of experts is substituted and carried through the analysis process of Steps 3-4 of this thesis' framework subject to a $12.9 per-period cost constraint, the Pareto frontier of Figure 86 results. Note that the frontier is similar in shape to Figure 82 but with substantial vertical stretching due to, as noted in Section C.2.3.2.2, the fact that this set of experts on average assigns a substantially higher probability of continuing demand for missions to LEO in the event that current mission demand is fulfilled (85.7% vs. 67.8% in Table 34), resulting in a longer maintenance for LEO mission demand and a higher number of missions accumulated to demanded destinations when the LEO-themed architecture enters into operation. As Figure 87 shows, however, the modified demand model has minimal impact on the optimal initial configuration decisions in comparison to Figure 84 and Figure 85: With the exception of the lowest-cost option to cancel all future development, all other policies are in agreement to initially continue development of the LEO-themed architecture.

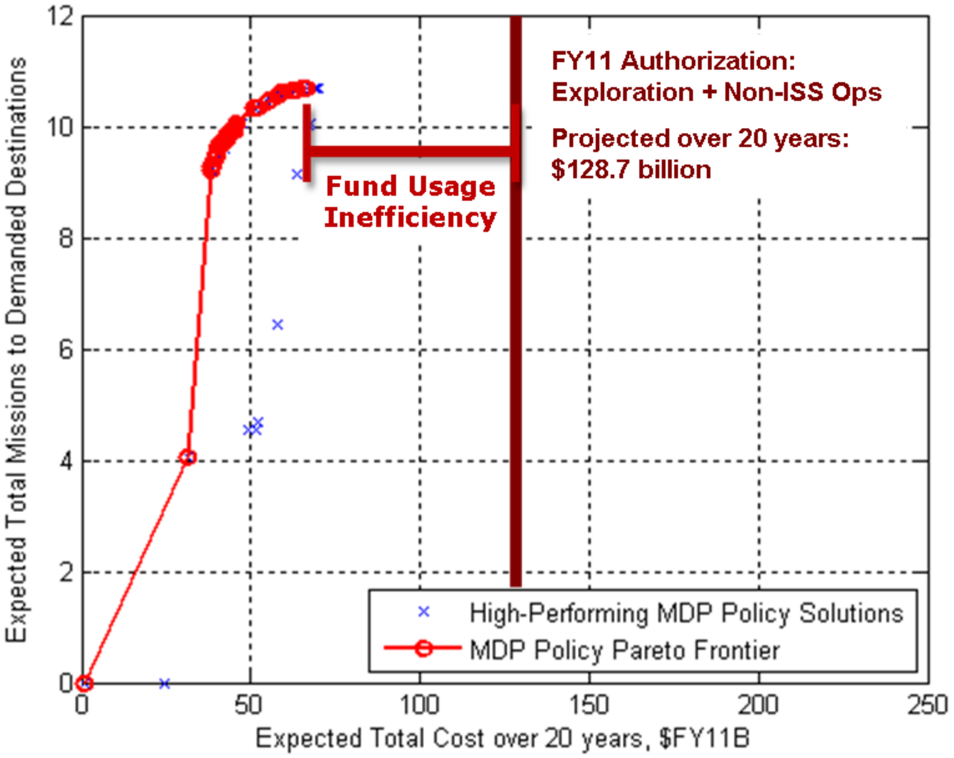

Figure 86. Trade between expected total missions to demanded destinations and expected total cost for MDP policy solutions subject to a $12.9 billion per-period budget constraint and subject to a demand model based on the central tendency of the non-dominated expert probability estimates. *Marked on the plot is an appropriate long-term budget expectation and the gap that exists between it and the highest Pareto-optimal method of spending.*

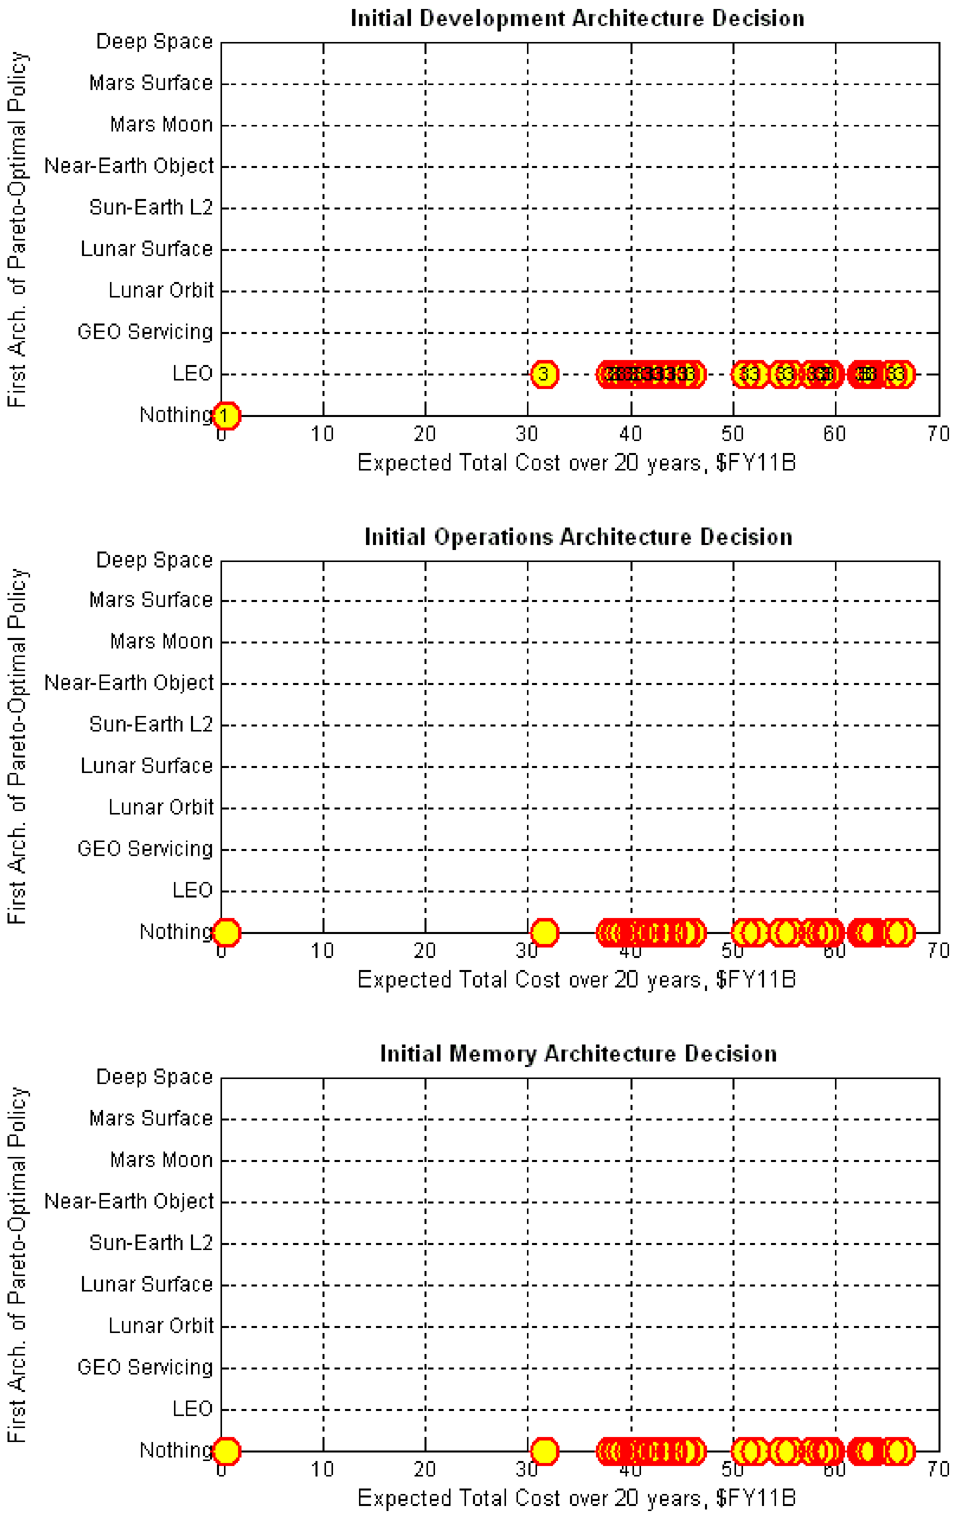

Figure 87. Initial architectures of configurations implied by Pareto-optimal policy solutions subject to a $12.9 billion per-period cost constraint and subject to a demand model based non-dominated expert probability estimates as a function of expected policy total cost.

6.5.2. Accounting for Non-Expected-Value Objectives

As discussed in previous chapters, a final relevant consideration for initial system selection is the fact that expected-value objective functions for the cumulative cost and performance metrics may not fully capture a decision-maker's objectives. While use of these cumulative expected-value objectives enables the use of MDP dynamic programming techniques to efficiently explore the astronomically large policy trade-space, consideration must in general be accorded to other objectives as well. Applying to the NASA human space exploration example the genetic algorithm developed and used in Chapters 4-5 yields the more extensive set of multi-objective optimal policy solutions presented here.

In addition to appending the 90th percentile (near-worst-case) total cost and 10th percentile (near-worst-case) performance metrics as in Chapters 4-5, this section adds two metrics (in both their mean and near-worst-case dispersion senses) implied as a result of the the figure of merit portion of the survey sent to human space exploration experts discussed in Appendix C and Section 6.3. The first metric is the date of the first mission to leave low-Earth orbit, intended for minimization. The second metric is the ratio of the number of missions flown to demanded destinations to the total number of missions flown over the simulation timeline, in short designated as "Mission Ratio". Both metrics employ an assumption, consistent with production costing assumptions detailed in Appendix B, that in states and times when a current operations architecture is unable to fulfill current mission demand (i.e., the Boolean zero elements of Table 38), the architecture can be and is flown on the missions and corresponding mission rates to which it is themed*. The mission ratio metric captures the efficiency with which missions are targeted toward demanded destinations and is intended for maximization (with a maximum possible value of unity). It also serves to capture the fact that the

---

* In the case of the Deep Space architecture, it is flown on the Mars Moon missions and rates.

original survey's figure of merit suggestion neglected the demanded destination distinction in the figure of merit definition, and thus in combination with the Number of Missions to Demanded Destinations metric allows for a measure of the total expected number of missions flown, including those to destinations that may not have been demanded or expected.

Applying the genetic algorithm described in Chapters 4-5 using each policy identified in Figure 78 as an initial member of the genetic algorithm population and searching for multi-objective optima in terms of each of the eight objectives described above (i.e., the original two cumulative objectives of Figure 78, the two new objectives described in the above paragraph, and their corresponding 90th or 10th percentile near-worst-case dispersions) yields the results of Figure 88. Each subplot in Figure 88 shows a cross-section of the the performance of the policy solutions, each of which is displayed as a data point colored by its implied initial configuration decision, in terms of two objectives. While in many cases the policy solutions show little variation in performance according to the percentile metrics (indicating that the near-worst-case results may be difficult to influence), the means tend to show substantial variation. For example, the leftmost subplot that is second from the bottom in Figure 88 shows that the expected date of the first mission beyond low-Earth orbit can be made as early as 2025 with sufficient expenditure of funds,[*] corresponding to policies with initial decisions to switch immediately to development of the Lunar Surface themed architecture (the light blue points in the subplots). As the subplot two above this subplot illustrates, these same policies also result in the highest mission ratios of about 0.6.

As the colors of the data points in Figure 88 emphasize, only five initial configuration decisions are identified among the Pareto-optimal policies in the genetic

---

[*] In simulations for some policies, no missions beyond LEO are ever flown. In such simulations, the date for the first beyond-LEO mission is recorded as 2031, which is one time step beyond the simulation time horizon. This explains the plateau for this subplot at low cost levels.

algorithm search. The blue and red points represent the Config. 1 and 4 options seen in the optimal expected-value analysis of Section 6.5.1 to either shift to no development or continue into the next phase of LEO themed architecture development. The green points indicate the option to restart development of the LEO themed architecture, the light blue points indicate the option to shift immediately to development of the Lunar Surface themed architecture, and the purple point indicates the option to shift immediately to development of the Mars Moon themed architecture. While, as mentioned earlier, the Lunar Surface themed architecture options provide clear benefits in terms of speeding the process of leaving low-Earth orbit and increasing the ratio of demanded to total missions flown, the second row in Figure 88 shows that continuation of LEO-themed architecture development (Config. 4) provides high numbers of missions to demanded destinations for both low mean and 90th percentile costs.

The usefulness of the multivariate plot of Figure 88 becomes even more evident if constraints are imposed by the decision-maker. For example, suppose that a decision-maker wishes to be 90% certain that the \$128.7 billion long-term human space exploration budget projection used earlier in this analysis will not be breached by the policy he or she adopts. Imposing this constraint eliminates many high-cost (and also high-performance) options that formerly fell into the high 90th percentile cost regions of the multivariate plot that are now gray in Figure 89. Among these eliminated options are the policies with Lunar Surface themed architectures as an initial development decision. Consequently, with the clear advantage in terms of number of missions to demanded destinations and mission ratio for over most of the cost range of interest (and little remaining variation available in terms of the first beyond-LEO mission date), policies with continuation of LEO-themed architecture development (Config. 4) are largely supported by these results, in basic agreement with the findings of Section 6.5.1.

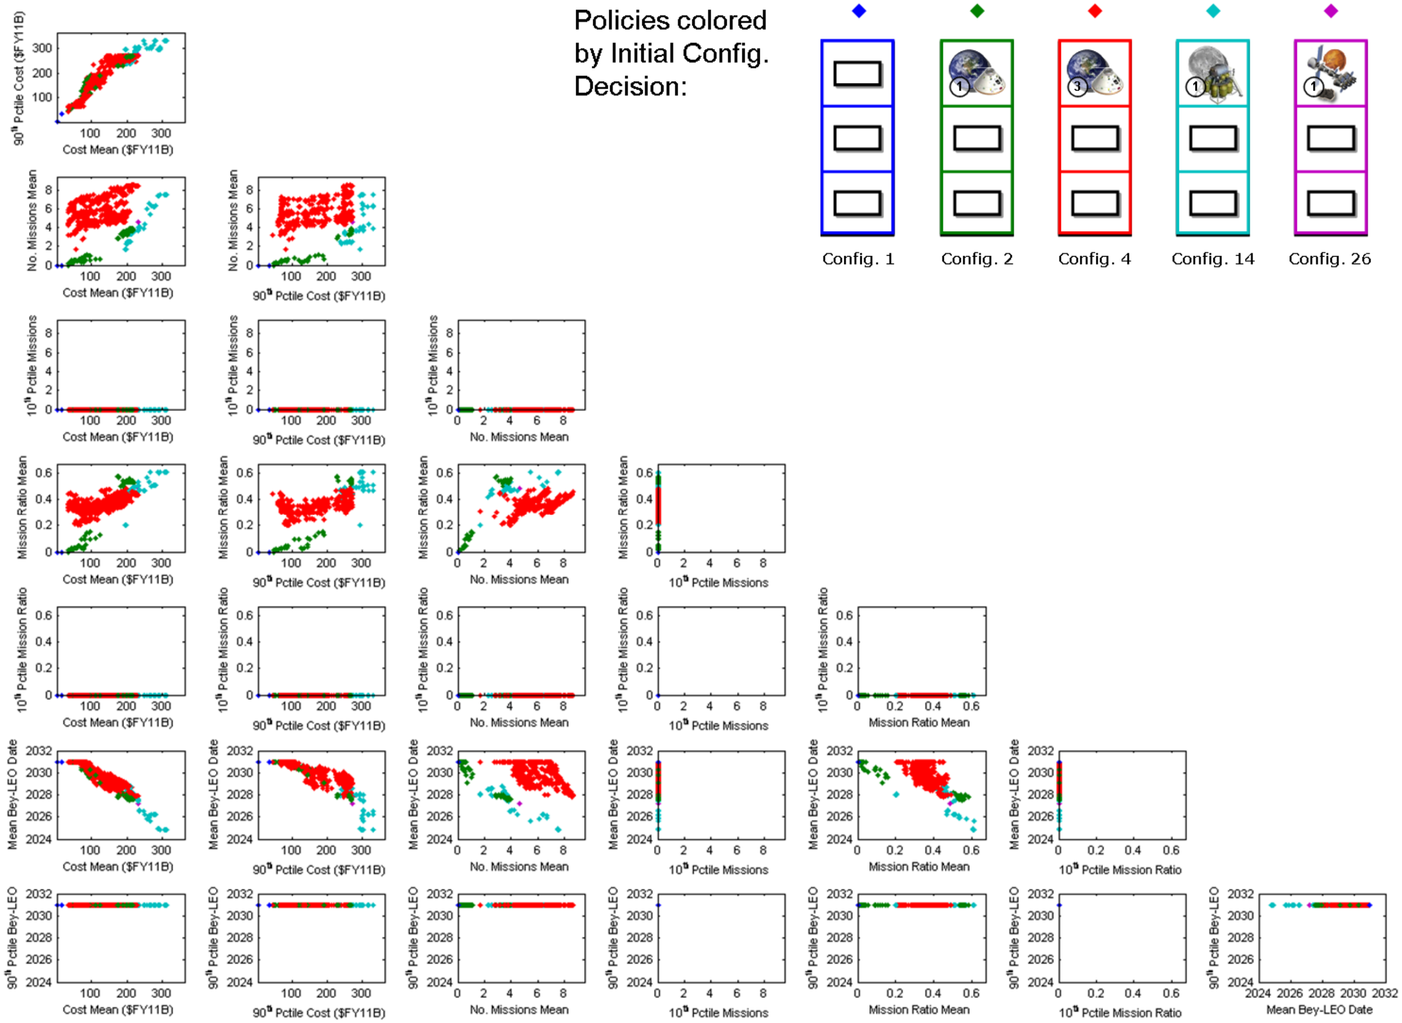

Figure 88. Multivariate plot of multi-objective genetic algorithm policy results. *Each data point indicates the performance of one policy result in terms of the eight percentile and expected-value metrics. Data points are colored by their policy's initial configuration decision.*

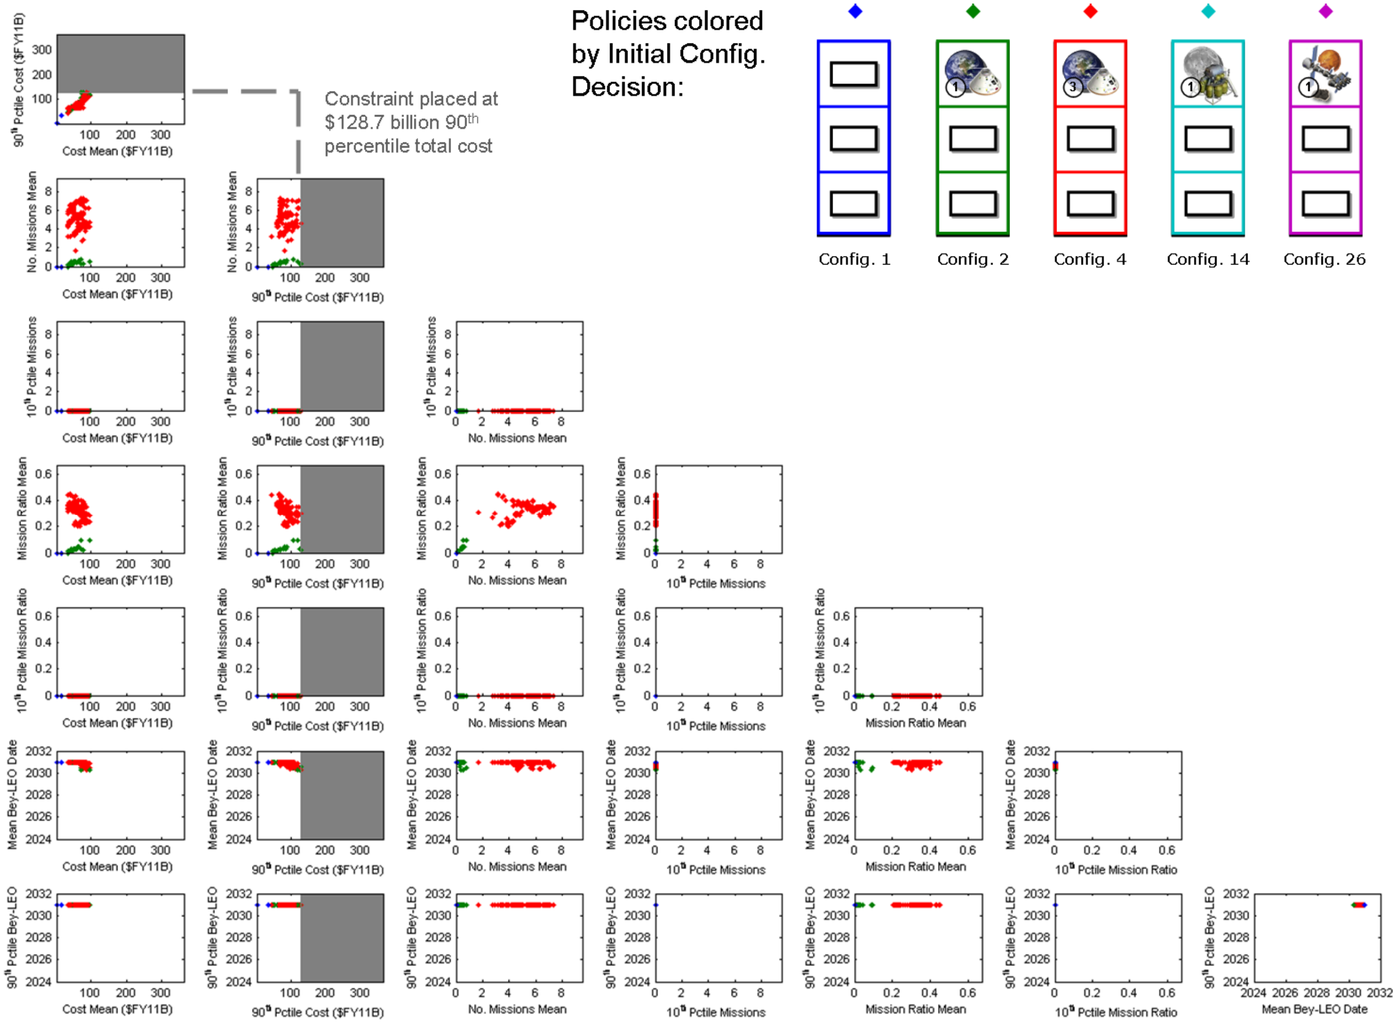

Figure 89. Multivariate plot of multi-objective genetic algorithm policy results with 90th percentile cost constraint imposed. *Each data point indicates the performance of one policy result in terms of the eight percentile and expected-value metrics. Data points are colored by their policy's initial configuration decision. Gray areas indicate regions of the space eliminated due to the cost constraint.*

6.6. Summary

This chapter has covered in substantial detail how the core theoretical framework proposed in Chapter 4 of this thesis can be extended and applied to address long-term program planning for the course of NASA's human space exploration efforts. The application has illustrated the ability of this framework to accommodate large configuration spaces of hundreds or thousands of candidate engineering configurations. To accommodate this, an automated cost model accounting for development, production, mission and ground operations, program management and systems engineering, and program termination and retirement costs was developed to facilitate population of a large (10.8 million element) cost transition matrix. The application has also illustrated the ability of the framework to model multi-period development, which introduced the need to use configuration state definitions accounting for development and operations architecture decisions as well as memory. Furthermore, the ability of the framework to model dependence between the effects of previous system configuration decisions and the demand environment was demonstrated, and an extensive survey distributed to individuals with human space exploration and systems engineering experience demonstrated how such a configuration-dependent Markovian demand model can be aggregated from multiple expert probability estimates. The ability of the framework to accommodate endogenous uncertainties, here in the form of schedule slippage, was demonstrated, as was the ability of the framework in its fifth step to account for non-cumulative objective functions.

In terms of practical implications and insights for human space exploration, implementation of the framework in this chapter has provided several: Step 1 illustrated that existence of the Deep Space architecture as a configuration's operations architecture is associated with a very high number of available transitions (i.e., options or, roughly, flexibility) with respect to other architectures at high per-period budget levels but not at low budget levels. Since the Deep Space architecture is to a large degree representative of the Review of U.S. Human Spaceflight Plans Committee's "Flexible Path" approach, the practical implication of this observation is that the flexibility of this (or any other) approach or configuration in comparison to its alternatives may be a strong function of available budget resources. In this context, Step 1 also illustrated that NASA's *current* human space exploration configuration is starkly inflexible in the context of the candidate architectures and configurations in the state space of interest, with relatively few transition options even at high budgets over the coming two-year period.

Implementation of Steps 2 and 3 of the framework involved eliciting expert opinions regarding mission demand environment evolution and figure of merit importance. The resulting Markovian demand environment model shows a general progression in demand toward the Martian surface, on the condition that mission demand is fulfilled, with secondary demands (or "sinks") at the Lunar Surface and Low-Earth Orbit and tertiary demand at Near-Earth Objects. Progression toward the Martian Surface demand is less likely under the condition that mission demand is not fulfilled, and the model exhibits the general characteristic that the condition of demand being fulfilled favors progression toward missions aimed at more ambitious destinations that are farther away from Earth; conversely, the condition of demand not being fulfilled tends to favor constancy or sometimes regression of demand toward less ambitious destinations closer to Earth. In terms of the figures of merit, consistently high-scoring metrics from the survey results lead to use of Integrated Program Lifecycle Cost and Number of Missions to Demanded Destinations as two objectives for exploration and optimization.

Step 4 identified Pareto-optimal policies over a range of long-term cost and performance preferences. Considering initially the case of no per-period budget constraints, it was shown first that an anticipatory reference policy was dominated by others that could perform at higher numbers of missions to demanded destinations at lower long-term costs. Among the Pareto-optimal policies it was shown that, due in part to the transience of the Mars Surface demand, development of the Mars Surface themed architecture was, interestingly, never optimal. Instead, in scenarios where demand evolved to ambitious destinations, many high-performing policies favored development of the Deep Space architecture. The main exceptions appeared to be development of the LEO and Lunar Surface themed architectures, which dominated development and operations plans for most policies due principally to the predominant progression of the demand environment, the fact that LEO development is partially complete as an initial condition, and the fact that the Lunar Surface themed architecture has the ability to operate missions in the LEO demand environment as well as the Lunar Surface demand environment. In the three Pareto-optimal policies examined in detail, no missions away from LEO started earlier than 2023.

The second case considered in Step 4 involved the implementation of a $12.9 billion per-period budget constraint representative of a doubling (due to a two-year period length for ths present application) of the NASA FY11 authorization for exploration plus non-International-Space-Station operations. This constraint severely limited solution options, and the highest-performing Pareto-optimal solution involved continuation of development of the LEO-themed architecture until completion and subsequent transition to LEO-themed architecture operation with cessation of any new development. Compared to the per-period budget-unconstrained Pareto-optimal solution at the long-term budget level, the highest-performance constrained solution exhibits a 0.3 expected mission performance gap and, moreover, a cost gap of about $60 billion. This cost gap indicates the total amount of funds that, by the performance measure used in this work, would be used inefficiently by virtue of a constant use-or-lose $12.9 billion per-period budget.

The first segment of Step 5 examined the policy solutions of Step 4 in terms of their implied initial decisions and found agreement over the vast range of cost and performance preferences that the optimal initial decision is to continue development of the LEO-themed architecture into its third development phase. Only the lowest-cost option involved cancellation of this architecture’s development (and replacement with no development at all). These conclusions were found to hold even under an alternative demand environment model. In considering the implications of non-expected-value and non-cumulative objectives, the second segment of Step 5 confirmed these conclusions under the constraint that a decision-maker wishes to adopt a policy that meets a $128.7 billion 20-year program cost with 90% probability.

In this way, Steps 1-5 provide a set of information to the decision-maker not only about the best immediate decision (in this case, to continue development of the LEO themed architecture), but also a cost- and performance-tailored policy and a corresponding outlook for the future. As new information becomes available or as questions arise, the approach used here also provides the analyst and decision-maker with the ready ability to modify inputs and test the robustness of his or her results to changing numerical assumptions.

CHAPTER 7

CONCLUSION AND AVENUES FOR FUTURE WORK

7.1. Summary

At the outset in Chapters 1-3 of this thesis, a review of the state of the art and practice in aerospace engineering revealed that, at present, there exists no comprehensive quantitative, stochastic, multi-objective, and multi-period framework for integrating flexibility into space system design decisions. Moreover, it was recognized that a substantial need for such a framework exists: Flexibility is well-recognized as important to space system success, to the extent that DARPA and NASA have in recent years proposed flexible spacecraft and flexible paths, respectively, as future program directions with substantial budgetary and resource implications. Because this property of flexibility is by definition linked to the ability of a decision-maker to make choices in response to [typically uncertain] changing environments or requirements over multiple periods, a framework that considers the integration of flexibility into decision-making must be both stochastic and multi-period in nature. Because most engineering applications involve trades among multiple objectives, such a framework must be multi-objective in order to completely consider the breadth of decision-maker interests. Finally, to permit the use of objective performance metrics as opposed to unitless subjective ratings, such a framework must also be quantitative.

The framework that this thesis introduces in Chapter 4 consists of five practical steps intended for implementation by engineering systems analysts, the first three of which focus on defining and characterizing a set of state spaces representing system options and environment demands. The fourth step employs multi-period decision analysis techniques, including Markov decision processes from the field of operations research, to find Pareto-optimal paths and policies a decision-maker may follow in a stochastically changing demand environment. With a set of full Pareto-optimal multi-period decision paths policies thus identified, the final step examines the implications of these paths and policies for the selection of an initial system. The end product is a quantitative, stochastic, multi-objective, and multi-period framework for integrating flexibility into space system design decisions. This thesis, moreover, illustrates that not only is the two-period state-centric notion of flexibility prevalent in the literature compatible with a comprehensive decision support framework, but that it is naturally adapted for use with Markov decision process solution techniques from the operations research community.

Three examples have been used to illustrate the application of this framework to space systems decision-making. The first and simplest example in Chapter 4 presented a scenario in which decisions were to be made regarding numbers of communications and reconnaissance satellites to be fielded to meet future national needs. This example was used as a means for exploring the present thesis' framework in great depth: The chapter began with a foundational two-period state-centric concept of flexibility from the economics literature and showed how, through the proper interpretation of this concept for space systems and linkage to the environments in which these systems may be required to operate, it can be unified with powerful dynamic programming techniques already in existence to solve Markov decision process problems. Along the way, several additional insightful analyses were developed, particularly in Step 1, in which the number of available transitions from a given configuration state at a given budget $\Phi_i(b)$ was developed as a surrogate metric for flexibility. In particular, it was illustrated that "flexibility reversals" are possible due to interactions between existing capabilities, existing commitments, and available resources. In these situations, more transitions are available from Configuration i than Configuration j at a budget level b_1, but fewer transitions are available from Configuration i than Configuration j at a higher budget level b_2 (i.e., that $\Phi_i(b_1)>\Phi_j(b_1)$ but $\Phi_i(b_2)<\Phi_j(b_2)$).

Later in Chapter 4, Step 4 made the important distinction between paths and policies; while paths are a more traditional method of planning and consist of the simpler task of laying out a set of actions to execute in future years, they preclude a decision-maker from considering the full “playbook” of if-then possibilities when making his or her decisions. Step 5 illustrated how the complicated policy (and, to some extent, path) results of Step 4 can be distilled into information that a decision-maker can use to make an initial system selection. Step 5 addressed how the expected-value optima of Step 4 can be used as reasonable initial guesses for more local design space searches in the case that decision-makers have non-expected-value or non-cumulative objectives in mind. Finally, Step 5 also addressed the intriguing point that flexibility has a particular niche in environments of neither very high nor very low uncertainty, but rather in environments in which the present gives just *some* information about future demand. Emphasized was the inherent link between flexibility and policy, which specifies the conditions under which a system’s flexibility is exercised.

The example of Chapter 5 demonstrated how the theoretical framework posed in Chapter 4 can be applied to a problem motivated by recent DARPA fractionated spacecraft development efforts. This chapter defined a scenario in which a hypothetical Department of Defense decision-maker was faced with a decision about what combination of payloads to launch upon potentially multiple distributed, free-flying satellites. Step 1 of this analysis illustrated how the number of available transitions metric Φ clearly captured the relatively high flexibility of a three-payload fully-fractionated configuration over a three-payload monolith over most budget levels. Step 4 of the analysis revealed examples in which, subject to a notional demand environment evolution model, an optimal path involved a one-period delay prior to fielding of a three-payload monolith and an optimal policy identified an efficient compromise between maximum performance and minimum cost by only developing the three-payload monolith if an appropriate level of demand for particular payloads materialized early during the program timeline. These examples illustrated the ability of this thesis' approach to identify non-intuitive high-performing, low-cost paths and policies that might otherwise be overlooked. Step 5 of this DARPA-motivated application objectively narrowed initial system selection decisions to just four candidate configurations, and the imposition of budget and performance constraints strongly suggested selection of the three-payload monolith as the initial configuration. This result highlighted the important conceptual point that finding a minimum-cost, maximum-performance solution in a changing demand environment may not be equivalent to finding a solution with maximum flexibility.

Chapter 6 presented the extension of the basic theoretical framework in Chapter 4 toward addressing long-term program planning for NASA's human space exploration efforts. New elements addressed included incorporating a large state space of thousands of configurations, multi-period development and associated operations and memory architecture decisions, configuration-dependent demand modeling, elicitation of expert-opinion Markov chain probabilities, incorporation of endogenous schedule-slippage uncertainties, and exploration of non-cumulative as well as non-expected-value objectives. Results of this chapter provided several practical implications and insights for human space exploration. For example, implementation of Step 1 within this chapter illustrated that the relative flexibility of a configuration utilizing a Deep Space architecture in operations can be a strong function of available budget resources. Thus, the availability of these resources must be considered prior to classifying a configuration or approach as more flexible than, less flexible than, or equally flexible as its alternatives. Step 1 also illustrated that NASA's *current* human space exploration configuration is starkly inflexible in the context of the candidate architectures and configurations in the state space of interest, with relatively few transition options even at high budgets over the coming two-year period.

Using a Markovian demand environment model derived from the central tendency of expert probability inputs describing human space exploration mission demand evolution, Step 4 within Chapter 6 identified Pareto-optimal policies over a range of long-term cost and performance preferences. Interestingly, among the Pareto-optimal policies it was shown that due in part to the transience of the Mars Surface demand, development of the Mars Surface themed architecture was never optimal. Instead, in scenarios where demand evolved to ambitious destinations, many high-performing policies favored development of the Deep Space architecture. However, in most scenarios, demand remained at less ambitious missions and prompted development of the LEO and Lunar Surface themed architectures. Also considered within Step 4 of Chapter 6 was implementation of a \$12.9 billion per-period budget constraint. This constraint was found to severely limit solution options, and the highest-performing Pareto-optimal solution given the constraint involved continuation of development of the LEO-themed architecture until completion and subsequent transition to LEO-themed architecture operation with cessation of new architecture development. Compared to the per-period budget-unconstrained Pareto-optimal solution at the long-term budget level, the highest-performance constrained solution exhibited a cost gap of about \$60 billion which, by the performance measures used in this work, would be used inefficiently by virtue of a constant use-or-lose \$12.9 billion per-period budget. However, regardless of whether a constrained or unconstrained per-period budget assumption was used, and even with the inclusion of additional non-cumulative and non-expected-value metrics, the initial system decision analysis of Step 5 supported for virtually all long-term cost and performance levels continuation of present LEO-themed architecture development as an immediate next step for human space exploration.

Overall, the applications of this thesis' framework demonstrated throughout the preceding pages have not only fulfilled the framework's intent of informing initial system selection, but also have provided (1) cost- and performance-tailored policies and a corresponding outlooks for future costs and utilization, (2) insights regarding future options and flexibility, (3) useful models for examining demand evolution, and (4) the ability to re-execute the quantitative analysis and examine decision, cost, or performance sensitivity (or robustness) as assumptions change or new information becomes available.

7.2. Contributions

Summarized, the main contribution of this thesis is a quantitative, stochastic, multi-objective, and multi-period framework for integrating flexibility into space system design decisions. While Chapter 2, and particularly Table 1, note that some of the individual elements of this framework have been suggested at various times and by various analysts and engineers in the aerospace industry over the past decade, no works to date have unified them in a way to enable the comprehensive analysis, trade-space exploration, and decision-making capability demonstrated within this thesis.

More specifically, this main contribution is enabled by several component contributions, including (1) formulation of the two-period state-centric notion of flexibility as a formal configuration-state-based concept for space system analysis and design, (2) formulation of a state-centric stochastic multi-period model capable of describing evolution of the demand environment in which an engineering system operates, and (3) incorporation of system modification policy into initial system selection by using the above formulation to pose integration of flexibility in design as a solvable sequential decision-making problem.

A fourth component contribution is the implementation and demonstration of the utility of solving for the Pareto-optimal sequential decisions enabled by flexibility, including optimal “open loop” sequential system configuration paths and “closed loop” system configuration policies. Enabling tools utilized from the operations research community are the formulation and probabilistic dynamic programming solution techniques for Markov decision processes. In addition, Appendix A contributes a new heuristic technique for identifying concave portions of Pareto frontiers in dynamic programming problems. Fifth, these Pareto-optimal configuration paths and policies are used systematically to recommend initial system configuration decisions.

The final component contribution of this thesis is the application and illustration of this framework to relevant space system design problems using the examples of (1) communications and reconnaissance satellite system selection, (2) multiple- or distributed-payload satellite selection, and (3) NASA human space exploration architecture selection. A requirement for execution of these examples is the development of transition cost and stochastic demand environment evolution models, contributed for the NASA example in Appendices B and C.

7.3. Avenues for Future Work

As noted in Section 7.2, the main contribution of this thesis is a framework for integrating flexibility into space system design decisions. Despite its positive qualities and advances over previous work, however, it would be naïve to claim this is *the* framework for integrating flexibility into space system design decisions. By necessity, this framework *approximates* the true systems, environments, and selection process that a decision-maker must consider. With this in mind, the contributions of this thesis should be viewed in two contexts: First, in the form presented in this thesis, this framework is a powerful tool for *informing* design decisions. The ultimate choice remains with the decision-maker, who may consider and trade the full set of factors, effects, and constraints (technical, programmatic, political, or otherwise) for a design problem; however, as an approximation to the full problem, this framework may still (1) reveal high-performance and/or low-cost policy solutions that would otherwise be nonintuitive, (2) support or challenge the performance, cost, and approximate optimality of hypothesized policies, and/or (3) allow investigation into why certain paths and policies perform well or poorly. Moreover, beyond its computational capability to examine the optimality of paths and policies, the framework provides a set of concepts useful in framing decision-making thought on flexibility.

Second, this thesis' contributions provide a step forward, upon which future investigators may build to improve modeling detail and realism, in the consideration of flexibility in the design of space and other engineering systems. To this end, the following discussion identifies several interesting questions and avenues for future work that engineers and anlysts may choose to investigate in the future.

7.3.1. Multi-Period Expansion of the Φ Transition Metric

Introduced as an element in the analysis of the cost transition matrices generated in Step 1 of this thesis' framework, the metric $\Phi_i(b)$ expresses the number of configuration transitions available from Configuration i given budget b. This metric has a physical meaning and shares conceptual similarities with the idea of flexibility. However, as noted in Section 4.1.4, it is limited in that it accounts only for transitions one period into the future. An expansion of this metric is certainly possible and of interest for future investigation. Such a metric (e.g., Φ') could be defined recursively in terms of the simple two-period versions Φ as in Eq. (31). Here, $b(\tau)$ represents a schedule of budget levels b_1, b_2, b_3, …, b_T in T total periods, and Φ' expresses the total number of transitions accessible in the T-period tree originating from Configuration i in the first period. This metric is clearly just a start to tracking multi-period transition availability, and variant metrics might also be proposed, for example, to distinguish between options that involve expansion or downscaling of configurations.

$$\Phi_i'\left(b(\tau)\right) = \Phi_i\left(b_1\right) + \sum_{j \text{ s.t. } c_{ij} < b_1} \left(\Phi_j\left(b_2\right) + \sum_{k \text{ s.t. } c_{jk} < b_2} \left(\Phi_k\left(b_3\right) + \ldots \right) \right) \tag{31}$$

7.3.2. Decision-Makers with Authority Limitations

In many scenarios, particularly for complex and high-value systems under consideration by modern republics, there may exist a practical difficulty in identifying a single decision-maker or decision-making body. Instead, the decision-making process might be more accurately approximated as one in which a particular decision-making body has the ability to decide to pass a decision input to one or more other decision-making bodies. This input may then be used by the receiving decision-making bodies to produce inputs or recommendations for additional bodies, perhaps in iteration with the first decision-making body, until a final decision is reached. This chain of decision inputs, in which no one decision-maker has complete control over the final configuration decision but each has influence, might well be termed as a negotiation process. If elements of the negotiation process (e.g., the preferences of members of the other decision-making bodies) are uncertain, then to any individual decision-maker the transformation from one's own configuration recommendation to the final decision might appear probabilistic. In this case, the framework provided by this thesis already provides a means to model this scenario, provided that the decision-maker can provide an estimate of the probability that a particular final decision will be made given that his recommendation is for a given configuration (similar in implementation to the use of endogenous schedule slippage probabilities for the human space exploration example in Chapter 6). Given the complexities that exist in the flow of recommendations among decision-making bodies in governments and other organizations, the rigorous estimation of these probabilities may prove a rich and fruitful avenue for future research. Ultimately, correctly modeling this effect could permit decision-makers to maximize their influence on the choices of decision-making bodies utilizing their recommendations.

7.3.3. Robustness to Changing Decision-Makers

Another intriguing area for future work originates from a paradox of the flexibility sequential decision-making problem. As described in this thesis, the selection of the best initial configuration must consider decision options and demand evolution through future periods. This is central to the study of flexibility, since the existence of options and choices over multiple time periods distinguishes the flexibility problem from related problems of robustness and optimization. However, over a long enough time horizon, tomorrow's decision-maker will be different from – and have different preferences from – today's decision-maker. Thus, while today's decision-maker may be able to solve for the Pareto-optimal decision policy over a long time horizon through Steps 4 and 5 of this framework, he or she may not be around to implement this policy in the future. It may be, therefore, that this decision-maker can only count on being able to influence *today's* decision.

In such a scenario, the decision-maker would be interested in making a strategic choice of system or architecture configuration initially such that performance remains high and cost remains low regardless of the preferences of future decision-makers. This consideration may serve to reduce the likelihood of costly program cancellations and major redirections at the appointment of new decision-makers. Future work examining this area of future work may be enabled by the formulation of the flexibility problem presented here and may begin by examining not only Pareto-optimal paths and policies from Steps 4 and 5 of the framework, but also near-optimal sequences and policies.

Somewhat related to this area of future work is additional development of strategies to update and maintain probability, cost, and performance model information as time passes. The thrust of such development would be to allow a decision-maker to use this thesis' framework to make decisions at multiple future time periods without the need to repeat the entire analysis process from scratch.

7.3.4. Additional Theory and Algorithm Investigations

Some additional areas for further theory and algorithm development have arisen during discussions throughout this thesis and merit note here. First, with the exception of the brief coverage in Section 4.1.3.2.2, the uncertainties considered in this thesis have been aleatoric rather than epistemic. That is, the uncertainties in demand environment evolution and schedule slippage are not considered to exist due to cost or performance modeling limitations but rather due to inherent uncertainties in how events in the world will unfold in the future. This thesis has generally considered, for example, that the cost and performance matrices are associated with negligible uncertainty. Future development may consider methods for assessing the impacts of parameter uncertainties or the impacts of investments intended to change model parameters (for example, technology investments to reduce launch vehicle production costs) on optimal system decision results.

Second, future algorithm development is invited in two areas. Further algorithm development toward the goal of seeking concave Pareto frontiers for multi-objective dynamic programming problems (see Appendix A) would improve the quality of Pareto frontiers for applications with concave frontiers. Furthermore, if such a method could guarantee the identification of all Pareto-optimal policies it may eliminate the need for Step 4A of this framework, since the paths sought by Step 4A are special cases of the policies sought by Step 4B. In addition, algorithm development toward the goal of perturbing the Pareto-optimal cumulative expected-value objective policies in order to discover efficient policies in terms of non-cumulative, non-expected-value objectives is another area in which this thesis has only scratched the surface.

7.3.5. Additional Application-Specific Questions

Finally, some practical questions have arisen in the execution or discussions of the example applications in Chapters 5 and 6 that are beyond the scope of the present investigation but worth consideration in future studies. In the distributed payload versus monolithic satellite trade study of Chapter 5, for example, an interesting and open question remains of what combination of cost, performance, and probability inputs are required to favor monolithic versus fractionated spacecraft as optimal initial configuration solutions.

In the NASA human space exploration example of Chapter 6, a number of interesting variations on the basic study performed in this thesis could be performed, including (1) gradual tightening or relaxation of the $12.9 billion per-period cost constraint to study properties of the resulting Pareto-optimal solutions, (2) use of "personalized" rather than central-tendency expert probability estimates to examine the Pareto-optimal system implications of each expert's views of future mission demand evolution, and (3) implementation of discounting at various rates to simulate preferences for current over future cost and performance. Ideally, studies like this would be performed with decision-maker interaction to provide an understanding of how (or whether) changing assumptions changes the optimal initial system decision.

Additionally, the configuration and demand environment definitions of the NASA application may be modified as different problem scopes become of interest, candidate systems change, or other updated information becomes available. For example, in the time since work on the NASA application for this thesis was initiated, developments in the commercial space sector have driven NASA toward use of commercial systems only for International Space Station resupply (rather than a combination of commercially- and traditional government-developed systems). In this context of largely decoupled LEO and beyond-LEO human spaceflight programs, a reasonable modification to the scope of the human exploration application may be consideration of only beyond-LEO activities, which would involve not only removal of the LEO-themed configuration and LEO demand environment from Steps 1 and 2 of the application, but also subtraction from the available budget the funds that NASA plans to devote to commercial flights to the

International Space Station. Recent developments have also seen a change in NASA’s heavy-lift launch vehicle of choice from the Ares V to the Space Launch System, and additional destinations of interest within the Earth-Moon system have arisen, both of which could be accounted for with minor modifications to configuration costing assumptions and demand environment definitions.

Future advanced development of the NASA example might include consideration of (1) any costs associated with the retention of architectures in memory, (2) development time benefits associated with the existence of already-developed components, (3) additional commonality benefits below the level considered in this thesis (e.g., below the level of treating stages, crew vehicles, landers, rovers, etc. as basic components), and (4) demand evolution that is not only dependent on the current operations architecture’s interaction with the current demand environment, but also dependent on the current development and/or memory architecture’s interaction with the current demand environment. Finally, a potentially useful and complementary approach to this thesis’ use of the substantial computing power described in Section 6.4.3 would be to re-execute the NASA analysis using an 8-year time step, sacrificing the modeling of multi-period development but gaining the ability to analyze many more (on the order of several hundred, rather than the ten in Table 31) architectures in order to seek potentially non-intuitive architectural solutions.

7.4. Closing Remarks

This thesis began with the 40-year-old story of Skylab and the Apollo program’s largely accidental flexibility. With any luck, the pages of this thesis have conveyed that there now exist the tools necessary to analyze flexibility and, where appropriate, select space system designs, tailored to a decision-maker’s cost and performance preferences, that have the flexibility to suitably respond to uncertain and changing future environments and requirements. Ideally, the concepts and techniques provided by this thesis will make flexibility stories like Skylab – where flexibility existed but was accidental, or worse, where flexibility might have been desirable but was unattainable – relics for the history books.

Following these final remarks is a quote that Charles Darwin never said, but perhaps that he might have said[*] given the ideas in *On the Origin of Species*. A reflection on the ability of a species to survive, the excerpt conveys that a species' ability to adapt to change is of paramount importance toward its survival. This is an interesting and important closing thought to bear in mind. However, an equally important distinction exists between Darwin's natural world and the engineering world: In the world of natural selection, no species – and certainly no individual – has control over its genetic predisposition to adapt to new or changing climates, floods, famines, droughts, diseases, or predators. However, in the world of engineering systems, *humans control the "genes" (or design variables) of the system*. Engineering decision-makers have always had control not only of physical properties of engineering systems, but also of the inherent flexibility these systems have to adapt to the changing environments in which they find themselves. With this fundamental degree of control, it is the responsibility of space system engineers, analysts, and decision-makers now and in the future to continue to develop and utilize decision-making tools that will allow the engineering of the best possible "genetics" into tomorrow's space systems.

---

[*] Versions of this particular saying are, in fact, so widely quoted and misattributed to Darwin that even the California Academy of Sciences had it etched on the floor of its San Francisco museum.

According to Darwin’s *Origin of Species*, it is not the most intellectual of the species that survives; it is not the strongest that survives; but the species that survives is the one that is able best to adapt and adjust to the changing environment in which it finds itself.

Leon C. Megginson, Ph.D., Capt. USAAF (Fmr.), 1963
Southwestern Social Science Quarterly, Vol. 44, No. 1

APPENDIX A

A HEURISTIC METHOD FOR IDENTIFYING CONCAVE PARETO FRONTIERS IN MULTI-OBJECTIVE DYNAMIC PROGRAMMING PROBLEMS

A.1. Introduction

The past few decades have seen a significant rise in the use of Pareto frontiers in aiding aerospace system decision-making. Defined as the set of non-dominated design solutions, or the set of solutions for which one design objective cannot be improved without the sacrifice of another (e.g., see Refs. [87] and [91]), the concept of the Pareto frontier has for many become a cornerstone of aerospace systems analysis, both in theory and in practice. The fundamental advantage of the Pareto frontier is that it allows an analyst to objectively identify inferior design points without the need for information from a decision-maker on the relative priority of one design objective over another.

Frequently, obtaining a Pareto frontier requires only a representative scan of a problem's design space, or the space spanned by the range of a problem's input variables or options. During this design space exploration, the performance of each candidate design among the multiple metrics of interest is tracked, and the Pareto frontier can be identified by eliminating (filtering [112]) dominated designs from consideration. However, this procedure can be computationally intractable for applications in which the design space is very large, leading to the need for methods that are able to preferentially seek out Pareto frontiers (e.g., see Refs. [113]-[115]). Conceptually the simplest method for accomplishing this relies on an alternative interpretation of the Pareto frontier as the set of optimal designs over all possible decision-maker preferences. In short, this translates into weighting and aggregating all Ω objectives of interest into a single objective function J (often a simple additive weighting, such as in Refs. [116]-[122]), finding the optimum design(s) with respect to this function, changing the weights w_i, and repeating this process over the entire (or a representative) set of possible weights.

While this alternative interpretation is largely correct, a body of literature with its origins as early as the 1970s recognizes that a simple additive weighting aggregate objective function will fail to capture concave portions of a Pareto frontier [123]-[130]. This is illustrated in Figure 90 through Figure 92 and Eq. (A1). Figure 90 shows an example of an objective space defined by incommensurate objectives Γ_1 and Γ_2. Both Γ_1 and Γ_2 are normalized on a scale from zero to unity, such that larger values of both objectives are preferred. The Pareto frontier is clearly concave with respect to the origin. If the iterative procedure described in the previous paragraph is employed using a simple additive weighting objective function (i.e., Eq. (A1) where $n = 1$), only convex portions of the frontier are identified, as shown at the left in Figure 91. If the same procedure is applied but with an objective function of increasing order (e.g., $n = 2$ and $n = 4$), the concave Pareto frontier is captured more fully [123]. If the Tchebycheff norm is used, denoted in this work as $n = \infty$, all Pareto-optimal points may be captured [127]-[128], limited in resolution only by the discrete weightings considered.

$$J \equiv \begin{cases} -\sum_{i=1}^{\Omega} w_i \left(1-\Gamma_i\right)^n & , n < \infty \\ -\max\left(w_1\left(1-\Gamma_1\right), w_2\left(1-\Gamma_2\right), \ldots, w_\Omega\left(1-\Gamma_\Omega\right)\right) & , n = \infty \end{cases} \tag{A1}$$

The dependence on n in an aggregate objective function's ability to capture the Pareto frontier can be explained graphically via Figure 92. Contours in Figure 92 represent values of the aggregate objective function J as n is increased. Notice how the maximum curvature of each contour increases as n is increased: At $n = 1$ there exists no curvature, and by $n = \infty$, there exists a point of infinite curvature on each contour. In effect, as n is increased, each contour penetrates more deeply toward the concave portion of the Pareto frontier. As weights w_i on each objective are varied (Figure 92 illustrates only the case $w_1 = w_2 = ½$), the relative location of the point of maximum curvature changes, and different points on the concave frontier maximize J and are recorded as Pareto-optimal.

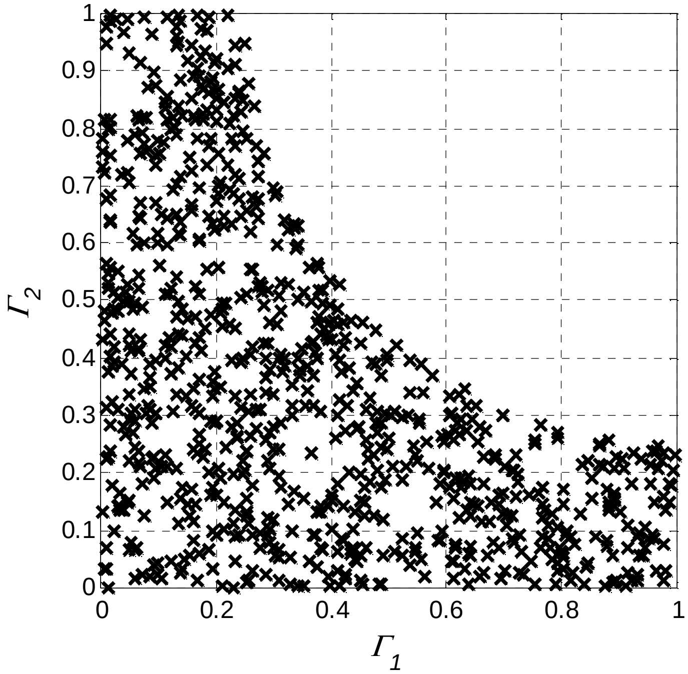

Figure 90. Example of an objective space with a concave frontier of nondominated points.

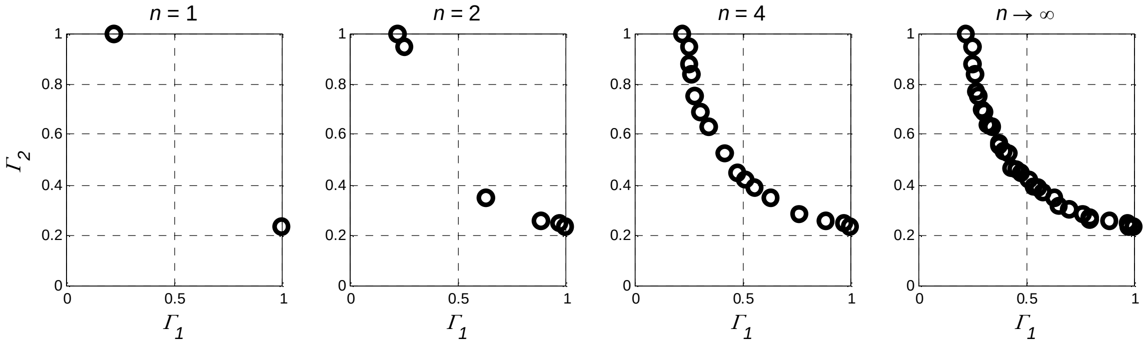

Figure 91. Nondominated solutions identified via the aggregate objective function of Eq. (A1) for varying *n*.

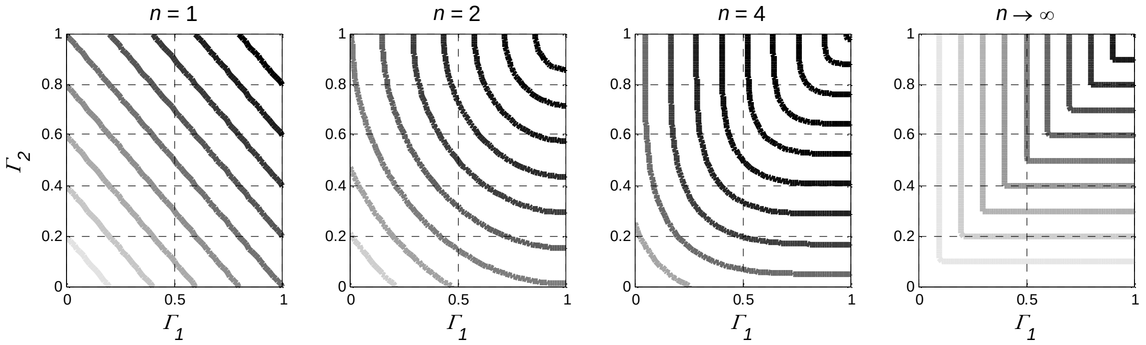

Figure 92. Contours of the aggregate objective *J* (see Eq. (A1)) as *n* increases. Γ_1 and Γ_2 are weighted equally.

A.2. Application to Dynamic Programming

Many aerospace applications make use of dynamic programming as an efficient optimization procedure for multi-stage decision problems. First introduced in the 1950s by Bellman [93],[131], dynamic programming takes advantage of the recursive structure of many multi-stage objective functions in order to decompose the optimization problem into a series of more tractable single-stage optimization problems. Commonly, the recursion in dynamic programming problems takes the form of stage-to-stage addition. That is, the objective Γ is comprised of summed single-stage γ_τ objective functions as in Eq. (A2). Alternatively (but nearly equivalently), Γ may be comprised of an objective function derivative $d\Gamma/dt$ integrated over time as in Eq. (A3). In both cases, Bellman's principle of optimality forms the basis for efficient optimization. This fundamental principle states that "an optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision." [93]

$$\Gamma_i = \sum_{\tau=1}^{T} \gamma_{i,\tau} \tag{A2}$$

$$\Gamma_i = \int_{t=0}^{T} \dot{\Gamma}_i \, dt \tag{A3}$$

A subtle difficulty exists when applying dynamic programming to multi-objective problems. A computationally appealing method for solving these problems is to aggregate the multiple objectives into a simple additive objective function in each stage, apply single-objective dynamic programming algorithms as usual, and then scan over the possible aggregating weights to identify the Pareto frontier. The aggregate objective function for this method is shown in Eq. (A4) (or in Eq. (A5) for the continuous-time case). Note that η indicates the aggregate objective function at each stage (i.e., the per-stage version of J), and it is assumed the sum of all weights w_i is unity.

$$\eta = -\sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T} - \gamma_i\right) \tag{A4}$$

$$\eta = -\sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T} - \dot{\Gamma}_i\right) \tag{A5}$$

A property of this simple additive weighting form of η is that it sums (or integrates) to J as defined in Eq. (A1) for $n = 1$, as shown in Eqs. (A6)-(A7). However, as a consequence, the technique suffers from the limitation that it cannot detect concave portions of Pareto frontiers.

$$\sum_{\tau=1}^{T} \eta = -\sum_{\tau=1}^{T} \sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T} - \gamma_i\right) = -\sum_{i=1}^{\Omega} w_i \sum_{\tau=1}^{T} \left(\tfrac{1}{T} - \gamma_i\right) = -\sum_{i=1}^{\Omega} w_i \left(1 - \Gamma_i\right) = J \tag{A6}$$

$$\int_{t=0}^{T} \eta \cdot dt = -\int_{t=0}^{T} \sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T} - \dot{\Gamma}_i\right) \cdot dt = -\sum_{i=1}^{\Omega} w_i \int_{t=0}^{T} \left(\tfrac{1}{T} - \dot{\Gamma}_i\right) \cdot dt = -\sum_{i=1}^{\Omega} w_i \left(1 - \Gamma_i\right) = J \tag{A7}$$

An appealing solution to this problem is the application of a higher-order aggregate objective function, per the observations in Section A.1. Unfortunately, this solution has a problem since, in general, summing a per-stage or integrating a time-derivative objective function η of this form (see Eqs. (A8)-(A9)) does result in the total objective function J (see Eqs. (A10)-(A11)). Thus, use of this per-stage objective function η in a standard single-objective dynamic programming algorithm will not properly represent the objective to be maximized or minimized.

$$\eta = -\sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T} - \gamma_i\right)^n \tag{A8}$$

$$\eta = -\sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T} - \dot{\Gamma}_i\right)^n \tag{A9}$$

$$\sum_{\tau=1}^{T} \eta = -\sum_{\tau=1}^{T} \sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T} - \gamma_i\right)^n = -\sum_{i=1}^{\Omega} w_i \sum_{\tau=1}^{T} \left(\tfrac{1}{T} - \gamma_i\right)^n \neq J \tag{A10}$$

$$\int_{t=0}^{T} \eta \cdot dt = -\int_{t=0}^{T} \sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T} - \dot{\Gamma}_i\right)^n \cdot dt = -\sum_{i=1}^{\Omega} w_i \int_{t=0}^{T} \left(\tfrac{1}{T} - \dot{\Gamma}_i\right)^n \cdot dt \neq J \tag{A11}$$

However, one motivating observation can be made. In cases where the normalized objective γ_i (or its counterpart, $d\Gamma_i/dt$) is small compared to $1/T$ and where the normalized objective Γ_i is small compared to unity, the binomial approximation can be applied to Eqs. (A10) and (A11) to show (in Eqs. (A12)-(A13)) that the sum of the individual per-stage aggregate objective functions of order n nearly equals the total aggregate J, multiplied by a correction factor. In other words, using a nonlinear power-n *per-stage* aggregate objective function will properly sum to the power-n *cumulative* objective function and thus permit detection of a concave Pareto frontier *in the region of the objective space where designs perform poorly* (e.g., near the coordinate $\Gamma_1 = \Gamma_2 = 0$ in Figure 90).

At first glance, this observation appears to have limited utility, since poorly performing designs are generally of little interest. However, consider a simple two-dimensional concave Pareto frontier consisting of three points: (0,1), (ε, ε), and (1,0), where $\varepsilon << 1$. The point that produces the frontier's concavity, namely (ε, ε), is indeed poorly performing and thus might be accurately be identified using η of the form in Eqs. (A8)-(A9). While this example is extreme, it highlights the fact that when searching for concave Pareto frontiers, poor (but nondominated) designs are still of interest. In this example, for instance, finding that (ε, ε) is indeed on the frontier would provide the decision-maker critical information about the available trades. In this case, the decision-maker would almost certainly choose the single-objective maximum (1,0) or (0,1), rather than spend additional time and resources attempting to identify compromise solutions.

$$\begin{aligned} \sum_{\tau=1}^{T} \eta &= -\sum_{i=1}^{\Omega} w_i \sum_{\tau=1}^{T} \left(\tfrac{1}{T} - \gamma_i\right)^n \approx -\sum_{i=1}^{\Omega} w_i \sum_{\tau=1}^{T} \left(\tfrac{1}{T}\right)^n \left(1 - nT\gamma_i\right) \\ &= -\sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T}\right)^{n-1} \left(1 - n\sum_{\tau=1}^{T} \gamma_i\right) \approx -\sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T}\right)^{n-1} \left(1 - \Gamma_i\right)^n = J \cdot \left(\tfrac{1}{T}\right)^{n-1} \end{aligned} \tag{A12}$$

$$
\begin{aligned}
\int_{t=0}^{T} \eta \cdot dt &= -\sum_{i=1}^{\Omega} w_i \int_{t=0}^{T} \left(\tfrac{1}{T} - \dot{\Gamma}_i\right)^n \cdot dt \approx -\sum_{i=1}^{\Omega} w_i \int_{t=0}^{T} \left(\tfrac{1}{T}\right)^n \left(1 - nT\dot{\Gamma}_i\right) \cdot dt \\
&= -\sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T}\right)^{n-1} \left(1 - n \int_{t=0}^{T} \dot{\Gamma}_i \cdot dt\right) \approx -\sum_{i=1}^{\Omega} w_i \left(\tfrac{1}{T}\right)^{n-1} \left(1 - \Gamma_i\right)^n = J \cdot \left(\tfrac{1}{T}\right)^{n-1}
\end{aligned} \tag{A13}
$$

A.3. A Heuristic Method for Identifying Pareto-Optimal Solutions

In many cases, concave Pareto frontiers will not be as sharply defined as the (ε, ε) example earlier, and thus it would be incorrect to claim that applying per-stage aggregate objective functions of the form in Eqs. (A8)-(A9) will always result in the intended quantity J being maximized. However, this appendix's proposed method is motivated by the hypothesis that applying such an objective function can provide a greater likelihood of finding concave Pareto frontiers in such multi-objective dynamic programming problems.

The proposed method is summarized in Figure 93. In the first step, the objectives for the problem of interest must be identified and normalized such that the cumulative totals Γ_i are each no less than zero and no greater than unity. Implicitly, each of these objectives is additive such that Eqs. (A2)-(A3) hold. Secondly and thirdly, a set of powers n is selected and a set of weights $\{w_1, w_2, \ldots, w_\Omega\}$ is selected for testing. The set of weights is used to scan for Pareto-optimal points across a representative set of possible decision-maker preferences, and the set of powers is selected to increase the likelihood that concave portions of the Pareto frontier will be identified.

For each combination of the power n and set of weights, an aggregate per-stage objective function η is used, as specified in Eqs. (A8)-(A9). This approach of converting the multi-objective problem into a single-objective problem permits the use of traditional single-objective dynamic programming algorithms. Once such an algorithm is applied, the design variables leading to the optimum solution for this set of n and $\{w_1, w_2, \ldots, w_\Omega\}$

are recorded, and the process repeats for a new set of weights and/or a new power n. These recorded designs are the candidate Pareto-optimal solutions.

In the penultimate step, each of the candidate solutions is evaluated in terms of each of the Ω objectives of interest. As expressed in Eqs. (A10)-(A11), in general the sum of power-n aggregate per-stage objectives η (the quantity optimized) is not equal to the power-n aggregate of the cumulative objectives Γ_i (the quantity that would ideally find points on concave portions of the Pareto frontier, for $n > 1$). As a result, some of these candidate solutions are likely to be dominated solutions. Thus, the final step of this method is the discarding of dominated points to find the final estimate of the problem's Pareto frontier.

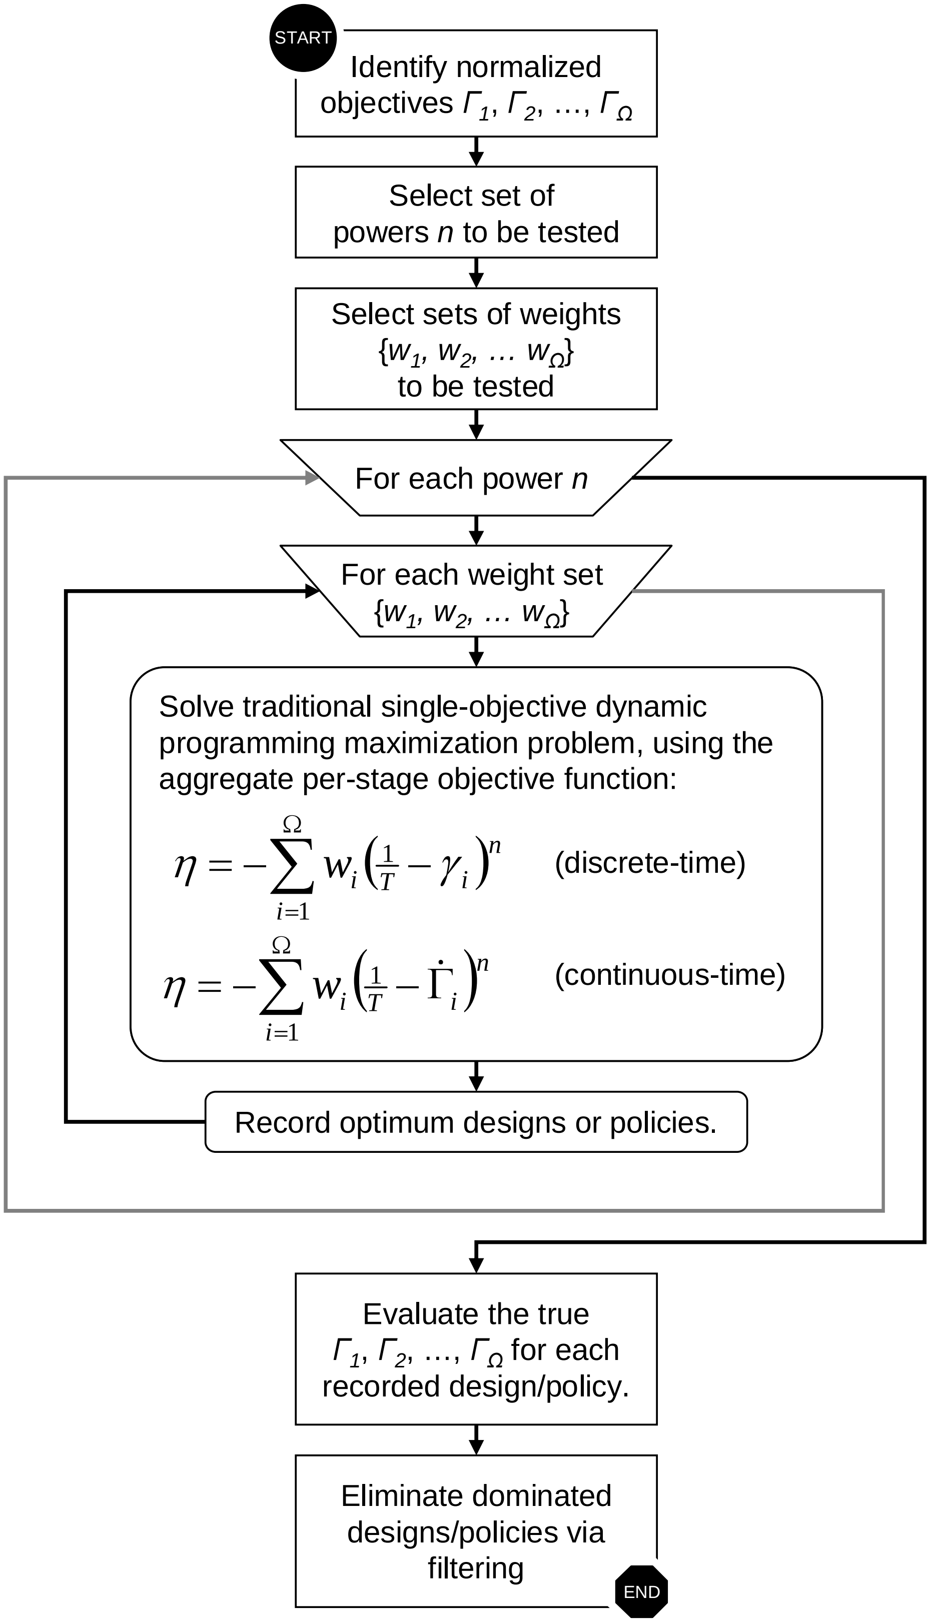
START
Identify normalized objectives $\Gamma_1, \Gamma_2, \ldots, \Gamma_\Omega$
Select set of powers n to be tested
Select sets of weights $\{w_1, w_2, \ldots w_\Omega\}$ to be tested
For each power n
For each weight set $\{w_1, w_2, \ldots w_\Omega\}$
Solve traditional single-objective dynamic programming maximization problem, using the aggregate per-stage objective function:
$\eta = -\sum_{i=1}^{\Omega} w_i \left(\frac{1}{T} - \gamma_i\right)^n$ (discrete-time)
$\eta = -\sum_{i=1}^{\Omega} w_i \left(\frac{1}{T} - \dot{\Gamma}_i\right)^n$ (continuous-time)
Record optimum designs or policies.
Evaluate the true $\Gamma_1, \Gamma_2, \ldots, \Gamma_\Omega$ for each recorded design/policy.
Eliminate dominated designs/policies via filtering
END

Figure 93. Flowchart for this appendix's heuristic multi-objective optimization algorithm.

A.4. Example Application

To demonstrate this method on a relevant and illustrative multi-objective dynamic programming problem, the following scenario is selected: A stealth aircraft loaded with enough ordnance to neutralize five hostile targets is to be flown across unfriendly territory, starting at a friendly airfield at coordinates (0, 100) and ending at a second friendly airfield at coordinates (500, 100) miles. The 500-mile stretch between the two fields is divided into five zones of equal length and each with breadth 200 miles. In sequence, one target in each of the five zones is to be neutralized. Each target has a particular strategic value (for example, measured in number of weapons or vehicles rendered inoperative), and it is desirable to maximize the total value of all sites neutralized during the mission. However, it is also desirable for the aircraft to minimize the total distance it travels during the mission (for example, to minimize its time at risk). Thus, this is a five-stage problem with two incommensurate objectives. Coordinates of candidate targets are listed in Table 40, and corresponding target values are listed in Table 41.

Table 40. Target Coordinates for Example Application.

Target No.	Target Coordinates (miles)				
	Zone 1	Zone 2	Zone 3	Zone 4	Zone 5
1	(55.2, 114.0)	(179.7, 193.6)	(260.5, 177.5)	(390.0, 186.5)	(457.9, 181.4)
2	(87.1, 113.3)	(191.8, 151.9)	(299.8, 161.1)	(312.9, 128.1)	(479.6, 164.2)
3	(12.3, 81.9)	(120.8, 125.4)	(232.9, 155.2)	(308, 124.1)	(486.6, 146.8)
4	(61.8, 78.2)	(137.9, 106.1)	(276.3, 123.5)	(349.2, 62.3)	(416.3, 46.1)
5	(26.8, 68.3)	(171.0, 53.9)	(268.0, 90.6)	(398.8, 56.6)	(462.4, 27.2)
6	(68.8, 42.3)	(124.3, 19.6)	(214.8, 76.5)	(388.5, 36.7)	(464.0, 7.4)
7	(75.1, 7.0)	(166.4, 13.8)	(221.9, 48.8)	(300.8, 19.9)	(403.3, 5.7)
8	(13.7, 4.7)	(148.4, 8.3)	(282.6, 44.9)	(354.6, 9.8)	(443.9, 4.5)

Table 41. Target Values for Example Application.

Target No.	Target Value				
	Zone 1	Zone 2	Zone 3	Zone 4	Zone 5
1	1	80	1	95	1
2	1	42	4	25	6
3	7	18	5	20	12
4	10	1	13	1	69
5	18	16	30	1	83
6	38	33	48	6	98
7	65	35	83	10	99
8	67	38	88	12	100

A.4.1. Full-Factorial Pareto Frontier

In this case, the problem is small enough to permit use of a computer to enumerate and evaluate all 8^5 = 32,768 possible routes (in general, such enumeration may not be practical, but this small example is selected to allow comparisons between the heuristically-generated and true Pareto frontiers). Distances between sites in sequential zones (as well as between the start and end sites and their neighboring zones) are precomputed and stored in an 8 × 8 × 6 array. In total, the full factorial evaluation of all possible routes requires 393,216 array lookups. The results of this evaluation are visualized in Figure 94, with the Pareto frontier outlined in dark gray. The frontier consists of 58 points and extends from a total distance of 524.5 miles and total target value of 189 (the distance optimal solution, shown as the black circle) to a total distance of 1006.8 miles and total target value of 430 (the value optimal solution, shown as the light gray triangle). Note that the frontier has two major concave segments, in the 600-750 mile range as well as the 800-1000 mile range. Three smaller concave segments exist within the 530-600 mile range. Also marked on the chart is an example compromise solution which attains a total value score of 302 with a 596-mile traverse.

Figure 95 graphically illustrates the locations and values of the targets listed in Table 40 and Table 41 as well as the physical solutions implied by the three solutions marked in Figure 94. Darker sites indicate sites of higher value. Notice that the distance optimal solution (black) takes the most direct route across the map but neglects high-value targets at the map's edges. On the other hand, the target value optimal solution (light gray) visits the highest-value targets in each zone but must fly in a costly and risky zig-zag pattern. The example compromise solution (dark gray) is similar to the distance optimal solution but deviates to visit high-value targets toward the bottom of the map.

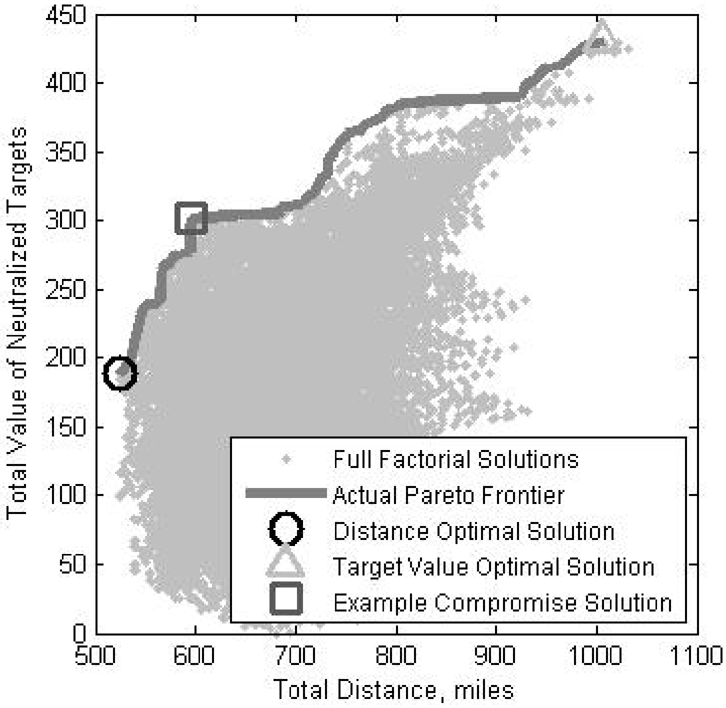

Figure 94. Performance trades and the full frontier for the example application.

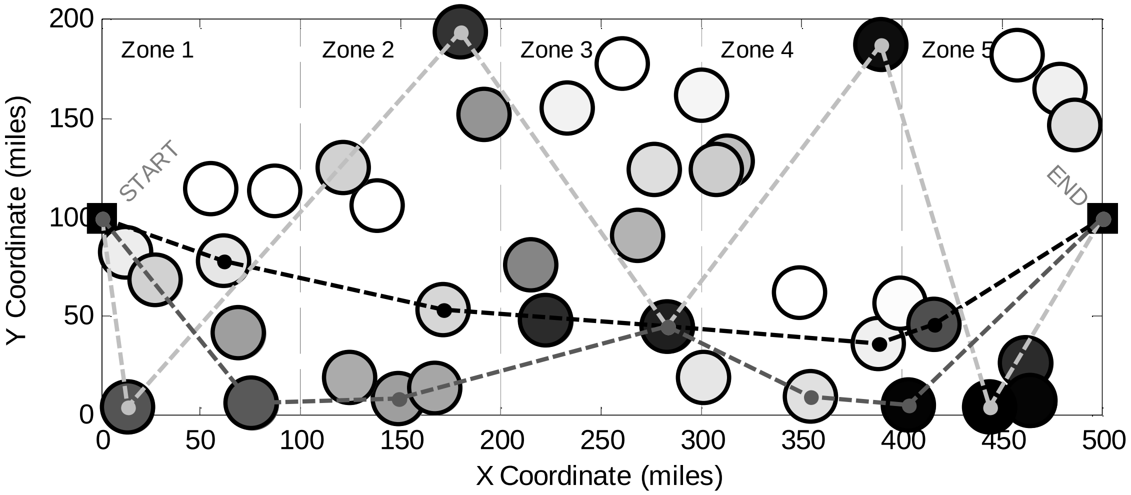

Figure 95. Graphical representation of the targets listed in Table 40 and Table 41. Darker sites indicate sites of higher value, and three sample paths are shown that correspond in color to the distance optimal solution (black), target value optimal solution (light gray), and an example compromise solution (dark gray) in Figure 94.

A.4.2. Heuristically-Generated Pareto Frontier

Approaching this example in the manner outlined in Section A.3 illustrates several advantages of this appendix's heuristic method. To begin the process, in this application the precomputed distance array is negated and subsequently offset and scaled such that the smallest element (previously the greatest distance) is zero and the sum of the maximum distances in each zone transition is unity. The target value matrix is offset and scaled such that the smallest elements (those with values of 1 in Table 41) are zero and the sum of the maximum target values in each zone is unity.

Following the remainder of the process outlined in Section A.3, the solid black line in Figure 96 shows the result for the selection $n = \{1, 2, 150, \infty\}$ and weights $w_1 = \{0, 10^{-8}, 10^{-6}, 10^{-4}, 10^{-2}, 0.04, 0.08, 0.12, \ldots, 0.88, 0.92, 0.96, 0.99, 0.9999, 1 - 10^{-6}, 1 - 10^{-8}, 1.00\}$, with $w_2 = 1 - w_1$. Note that the resulting frontier (in black) closely approximates the true frontier (in dark gray). In particular, the existence of both major concave segments is captured, as are the three smaller concave segments identified earlier. Also visible in the plot are candidate solutions (black triangles) from the heuristic algorithm that were not Pareto optimal when evaluated with respect to the true, rather than power-*n*, objective functions.

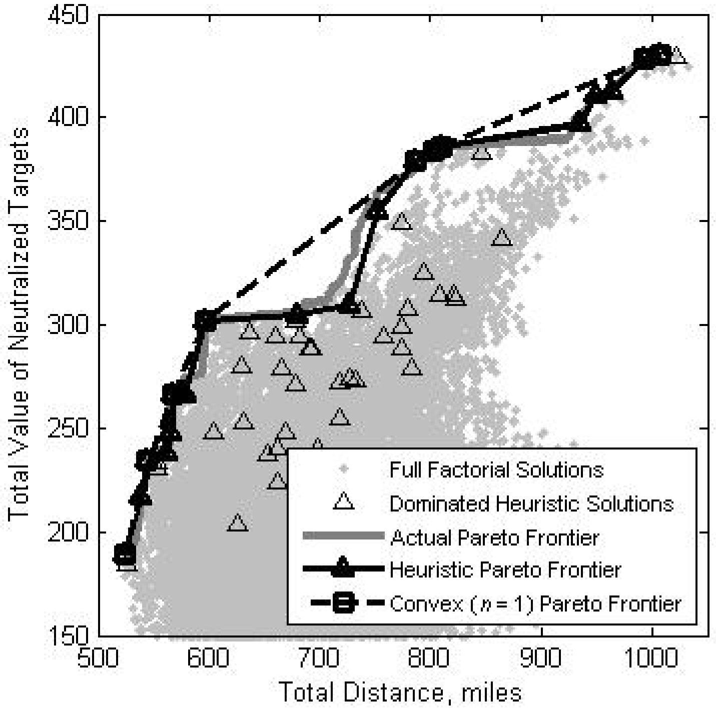

Figure 96. Comparison of Pareto frontiers generated using heuristic and full factorial methods.

Table 42 provides a useful comparison of the accuracy and efficiency of the three methods discussed in this appendix in the context of this example application. Accuracy is tracked here by two statistics. The first, which is the value of the coefficient of determination (R^2), indicates the degree to which the interpolated approximate frontiers explain the variations exhibited in the interpolated true frontier. While the simple additive weighting method achieves an R^2 value of 0.9389, the heuristic method performs significantly better with an R^2 value of 0.9899. The second measure numerically integrates the absolute value of the difference between the target value metric for each of the interpolated approximate frontiers and the target value metric for the true frontier over the range of the true frontier. Normalized such that the area between the simple additive weighting frontier and the true frontier is unity, this metric illustrates that the geometric area discrepancy is improved by more than a factor of three when the heuristic method is used in place of the simple additive weighting method.

In terms of efficiency, Table 42 highlights that the heuristic method requires less than 23% as many function calls (here measured in terms of the number of table or array lookup operations required) as the full factorial analysis for this example. Furthermore, 32% of the points identified from the heuristic method are nondominated, in contrast with the 0.18% ratio for the full factorial analysis. In this sense, the heuristic method is efficient in preferentially seeking points on the Pareto frontier. Furthermore, it might be reasonably hypothesized based on the advantages that dynamic programming provides that these indicators would more highly favor the heuristic method as the size of the problem (number of zones and number of sites per zone) increases. In terms of the simple additive weighting technique, it is notable that this method requires only about 19% as many function calls as the heuristic method and has 100% success in identifying Pareto-optimal points (in the sense that all points it identifies are Pareto-optimal). However, this metric does not reflect the number or importance of Pareto-optimal points on concave segments that the simple additive weighting method omits.

Table 42. Comparison of Pareto Frontier Search Methods in the Example Application.

	Pareto Frontier Search Method		
Metric	Full Factorial Analysis	Simple Additive Weighting (Dynamic Programming)	Heuristic Approach (Dynamic Programming)
Accuracy			
Coefficient of Determination (R^2)	1.00[a]	0.9389	0.9899
Integrated Area Discrepancy (normalized)	0.00[a]	1.00	0.3136
Efficiency			
Number of Function Calls (array lookups)	393,216	17,164	89,936
Ratio of Non-dominated to Total Points Evaluated	0.0018	1.00	0.3167

[a] By definition

A.5. Summary

In summary, this appendix has presented a heuristic method for identifying concave Pareto frontiers in multi-objective dynamic programming problems that employ additive recursion. Using a power-*n* (instead of simple additive weighting) per-stage aggregate objective function, the method possesses an important advantage of being easily integratable with existing single-objective dynamic programming algorithms; that is, only definition of the per-stage objective function need be modified, eliminating the need for a unique multi-objective dynamic programming algorithm. Because simple additive weighting is a special case of this method (the case $n = 1$), this heuristic method's results are at least as (and generally more) capable of identifying concave segments of Pareto frontiers, and in theory the technique becomes better able to identify points on concave frontiers as overall concavity of the frontier increases.

The example aircraft route selection application shown in this appendix illustrates how the heuristic method can substantially increase the accuracy of the detected Pareto frontier over simple additive weighting. Furthermore, the $R^2 = 0.9899$ coefficient of determination for the detected Pareto frontier is obtained with 4.4 times fewer function calls than required for the full factorial analysis.

It may be reasonably hypothesized that the computational advantage of this approach over full factorial analysis substantially increases as the number of fully enumerated paths increases. In the aircraft route selection application demonstrated here, the 32,768 paths could be enumerated, evaluated, and compared by a standard desktop computer within about one-half of a second. However, had the number of zones and available sites per zone each doubled, the number of paths would have increased to 1.10 trillion (a factor of 33 million greater!). In the full factorial approach, all these paths must be evaluated, potentially at a large expense of time and computational resources, because there exists no *a priori* knowledge about which paths are likely to be Pareto-optimal. In contrast, the number of function evaluations required by this appendix's heuristic method is controlled largely by the user's selections of sets of trial weights and powers, and in all cases the algorithm preferentially searches for Pareto-optimal solutions. As a result, although the ideal combination of power and weight sets cannot be known in advance, for large problems the computational expense of sweeping through a wide range of possible sets may easily be more efficient than the full factorial analysis. For example, using the extended w_1 weight set {1, 2, 3, …, 148, 149, 150, ∞} in this appendix's example application produces a nearly exact match to the true frontier, improving the fit of the detected frontier to an R^2 value of 0.9969.

This method is termed heuristic in the sense that no formal proof assures that the true Pareto frontier will be converged upon, even if the power and weight sets are increased infinitely in range and resolution. While the method is motivated by fundamental properties of the per-stage objective function η (noted in Sections A.1 and A.2), the fact that convergence tends to occur has only been observed empirically. Furthermore, the astute reader may notice two additional details which contribute to this heuristic characterization:

First, the definition of J in Eq. (A1) for $n = \infty$ is not the true limit of the $n < \infty$ expression as $n \to \infty$; for this to be true, the summation within the $n < \infty$ expression would need to be raised to the power $1/n$. However, this modification would nullify the theoretical accuracy for highly concave frontiers noted at the end of Section A.2, since the derivation of this property required the exchange of the per-stage and per-objective summations. Thus, strictly speaking, the utility of selecting $n = \infty$ as weighting is itself heuristic in nature.

Second, the large finite powers n (e.g., 150) used in the weighting sets for the example application present numerical difficulties since they are applied to aggregate objective functions η which are normalized to fall between zero and unity. When the resulting very small numbers, which can differ by many orders of magnitude, are added during the operation of the dynamic programming algorithm, some fall below computer numerical precision limits. When this occurs, it is possible for several next-stage alternatives to tie as optimum, in which case logic must exist to select among these equivalent alternatives. Rather than default to logic that selects the tied alternative which happens to appear first in the array, this appendix's implementation selects at random among the tied alternatives. Because different weight sets frequently result in identical optimal solutions, this randomization has the effect of diversifying the set of candidate Pareto-optimal solutions that are detected. Thus, the logic that handles selection among tied next-stage alternatives within the selected dynamic programming algorithm is also a heuristic element.

These components, which characterize this appendix's method as heuristic, are all worthy of future investigation and improvement. In the interim, it is hoped that the method presented here will contribute to theory and practice in multi-objective dynamic programming applications within the aerospace and broader system design and optimization communities.

APPENDIX B

TRANSITION COST MODEL FOR HUMAN SPACE EXPLORATION CONFIGURATIONS

Section 6.1.2 of this thesis describes how a custom cost model is used to estimate transition costs within Step 1 of the NASA human space exploration application. Described in its final form in this appendix, this transition cost model for human space exploration configurations was developed from publicly-available information and cost models over a period of approximately six months, the last two of which were spent obtaining feedback from systems engineers and cost analysts at NASA Johnson Space Center.

The cost model described here has the ultimate purpose of converting an input of two configurations (a “from” configuration and a “to” configuration) to a one-period transition cost. As described in Section 6.1.2, each configuration is a {Development, Operations, Memory} architecture triplet; and as described in Section 6.1.1, each architecture is defined by a set of components. Each period in the NASA application is assigned a duration of two years, and thus the cost model translates a decision to move from one {Development, Operations, Memory} architecture set to another into a two-year cost. Repeated use of this model over all possible pairwise combinations of configurations permits the population of the cost transition matrix for Step 1 of this thesis’ framework.

As shown in Section 6.1.2, Figure 58 summarizes the tools utilized and types of cost estimated by this model. The model has components coded in MATLAB (approximately 830 lines; outlined in orange in Figure 58) and in Microsoft Excel (approximately 260 lines of which are in Visual Basic; outlined in green in Figure 58). Once executed with a given set of inputs, the model can populate a full cost transition matrix within 25 minutes on a standard desktop computer.

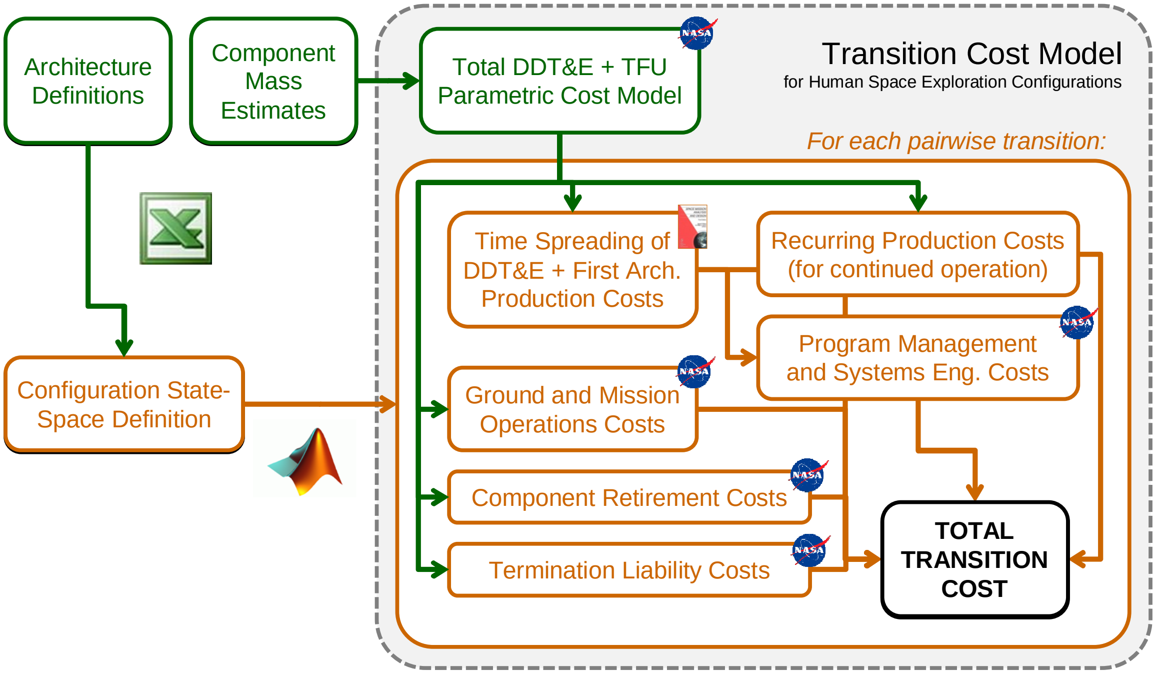

Architecture Definitions
Component Mass Estimates
Total DDT&E + TFU Parametric Cost Model
Transition Cost Model
for Human Space Exploration Configurations
For each pairwise transition:
Time Spreading of DDT&E + First Arch. Production Costs
Recurring Production Costs (for continued operation)
Program Management and Systems Eng. Costs
Configuration State-Space Definition
Ground and Mission Operations Costs
Component Retirement Costs
Termination Liability Costs
TOTAL TRANSITION COST

Figure 58. Transition Cost Model for Human Space Exploration Configurations.

In terms of the cost components estimated by this model (i.e., development, production, operations, program management and systems engineering, retirement, and termination liability), this appendix is organized to cover each in detail. Each of the following sections addresses one of these components in terms of (1) the core models upon which the parametric estimates are based and (2) any additional logic built in to the model to enforce consistency in assumptions. At the conclusion of the appendix, a validation is presented showing satisfactory results against a set of 121 cost transition estimates independently generated by NASA cost analysts in 2010.

B.1. Development and First-Period Production Costs

The most complex segment of the transition cost model is the portion that involves the estimation of development and first-period production costs. Although conceptually separable, these development and first-period production costs are covered within the same section in this appendix because both are assumed to be distributed, based on historical data [98],[132], over the entirety of a four-period (eight-year) development. Thus, estimating these costs involves (1) estimating the total development and production costs and (2) estimating the distribution of these costs among each of the four development periods. In addition, logic is included to account for the fact that existing components need not be re-developed.

B.1.1. Total Development and First-Period Production Cost Estimation

Total development and first-period production costs are based upon core estimates for design, development, test, and evaluation (DDT&E) and theoretical first unit (TFU) costs for each of the 25 architectural components listed in Table 31. With a few exceptions, these core DDT&E and TFU estimates are based directly upon the publicly-available NASA JSC Spacecraft/Vehicle Level Cost Model (SVLCM) [133], which outputs total DDT&E and TFU estimates as a function of system Earth weight $w_{vehicle}$ for several different classes of space vehicles and hardware. Of particular interest for this application are the liquid rocket engine, manned spacecraft, unmanned planetary spacecraft, and launch vehicle stage classes. The equations used to produce these estimates (in $FY11M) are given in Eqs. (B1) and (B2), and the a and b coefficients for these equations as a function of vehicle or hardware class are provided in Table 43. Note that the TFU cost equation includes a q_{TFU} term accounting for the production of multiple units; with the exception of two solid rocket boosters on the side of the assumed heavy-lift launch vehicle, two satellites per communication/navigation satellite pair, and two mobile power units included with the power generation and storage units, the per-TFU quantity $q_{TFU} = 1$. In all cases, LC is taken to equal unity (i.e., no substantial learning effects, in part due to the findings of the validation discussed in Section B.6).

$$c_{DDT\&E} = a_{DDT\&E} w_{vehicle}^{b_{DDT\&E}} \tag{B1}$$

$$c_{TFU} = a_{TFU} w_{vehicle}^{b_{TFU}} q_{TFU}^{\left(\frac{\ln LC}{\ln 2}+1\right)} \tag{B2}$$

Table 43. Coefficients for Eqs. (B1) and (B2) as a function of Vehicle/Hardware Class.

Vehicle/Hardware Class	$a_{DDT\&E}$	$b_{DDT\&E}$	a_{TFU}	b_{TFU}
Liquid Rocket Engine	24.039	0.550	0.121	0.662
Manned Spacecraft	15.089	0.550	0.435	0.662
Unmanned Planetary Spacecraft	10.152	0.550	0.674	0.662
Launch Vehicle Stage	5.951	0.550	0.129	0.662

The vehicle masses used as inputs to Eqs. (B1) and (B2) for each individual component of an architecture are provided in Table 44. Note that in many cases, an architecture component is itself comprised of multiple vehicles (e.g., multiple stages, multiple manned spacecraft modules), the costs of which are combined to produce a single DDT&E cost estimate and TFU cost estimate for the component.

Vehicle mass inputs are based on several sources. The crew launch vehicle component is modeled after the LV 13.1 option (approximately the Ares I) from the ESAS report [29], and the heavy lift launch vehicle is modeled after the LV 27.3 with EDS option (approximately the Ares V) from the ESAS report [29]. Masses for a representative deep-space habitation module, multi-purpose crew vehicle, lunar lander, multi-mission pressurized rover, and chemical stages are based on inputs from the NASA Human Exploration Framework Team (HEFT) and Human Spaceflight Architecture Team (HAT) [134]-[136]. The Mars lander is based on previous NASA design reference architecture planning [137]-[138], and the unpressurized rover, surface habitat, and ISRU systems are each based on mass assumptions within the ESAS report [29]. The logistics module is based upon the ESAS report [29] with a gross-to-dry-mass correction based on the Italian Space Agency's Multi-Purpose Logistics Modules [139]. The power generation and storage unit includes a component based upon the ESAS report's surface outpost primary power source [29] plus two mobile power units based upon Ref. [140]. The Mars Science Laboratory rover [141] is used as a representative science rover, the Mars Reconnaissance Orbiter (MRO) [99] is used as a representative communications and navigation satellite[*], and space suits are approximated as Apollo lunar space suits [143]. Engine masses are obtained from Refs. [144]-[147].

As noted by the two architectural components with no mass estimates in Table 44, the DDT&E and TFU estimates for the the Commercial Cargo Launch Vehicle and Commercial Cargo Logistics Module are not obtained from SVLCM but rather from representative NASA investment in commercial cargo vehicle development [148] (commercial launch vehicle development costs are assumed to be borne by the industry) and for SpaceX Falcon launch prices and Dragon per-flight contract rates [149]-[150].

While the NASA JSC SVLCM provides total DDT&E and TFU costs, it does not provide a breakdown of how those costs are spent by year. To accomplish this for the multi-period developments considered in the NASA application, an accepted historical model for the time spreading of program costs provided by the standard *Space Mission Analysis and Design* (SMAD) reference [98],[132] for conceptual design is discretized. As Figure 97 shows, the resulting distribution of costs among a four-period development is unimodal, with about 19% of costs incurred in the first period, 41% in the second period, 32% in the third period, and 8% in the fourth period.

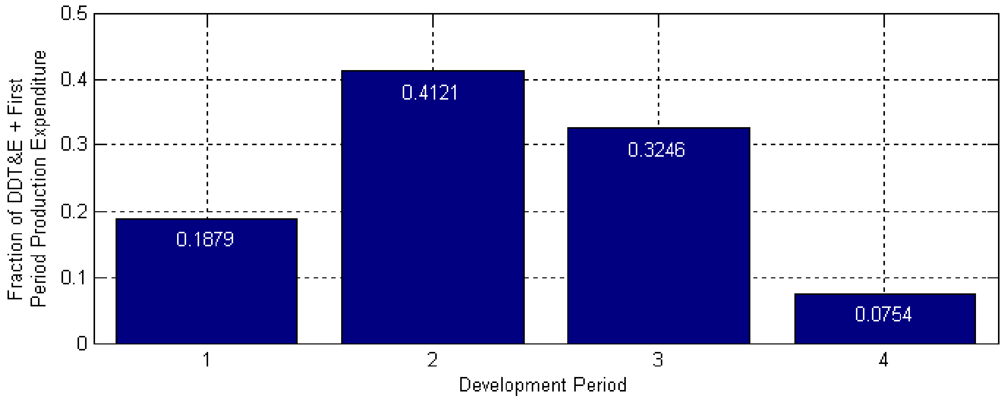

Figure 97. Normalized and four-period discretized development and first-period production cost distribution. [98],[132]

---

[*] Added to the communications and navigation satellite cost are two $100 million launch costs, representative for MRO's launch aboard an Atlas V 401 [142].

Table 44. Architectural Component Weight Assumptions and DDT&E and TFU Cost Outputs. [29],[99],[134]-[150]

Architecture Components		Liquid Rocket Engine Dry Weight (lb)	Manned Spacecraft Dry Weight (lb)		Unmanned Planetary Spacecraft Dry Weight (lb)		Launch Vehicle Stage Dry Weight (lb)			Total DDT&E ($FY11M)	Total TFU ($FY11M)
			SC #1	SC #2	SC #1	SC #2	Stage #1	Stage #2	Stage #3		
1.	Crew Launch Vehicle (CLV)						188049	39572		6746.4	544.5
2.	Heavy Lift Launch Vehicle (HLV)						221234	120617	31745	10670.0	1316.9
3.	Commercial Cargo Launch Vehicle (CCLV)									0.0	56.6
4.	Multi-Purpose Crew Vehicle (MPCV)		30159							4388.6	401.3
5.	Commercial Cargo Logistics Module (CCLM)									846.0	76.8
6.	Small Chemical Stage						34746			1871.2	131.3
7.	Medium Chemical Stage						39682			2013.0	143.3
8.	Large Chemical Stage						55119			2411.7	178.2
9.	Deep-Space Habitation Module		46680							5580.4	535.9
10.	Lunar Lander		4024	5810			11091			4222.0	302.3
11.	Mars Lander		9714				94578	18464		6920.6	530.6
12.	Multi-Mission Pressurized Rover		10095							2403.8	194.5
13.	Unpressurized Rover		551							485.7	28.4
14.	Science Rover				1709					608.8	93.0
15.	Surface Habitat		33069							4616.7	426.5
16.	Logistics Module		6842							1940.9	150.3
17.	Power Generation and Storage Unit				25353	3404				3573.5	848.5
18.	ISRU Systems				7937					1417.0	257.2
19.	Surface Extravehicular Activity (EVA) Suits		147							235.0	11.8
20.	In-Space Extravehicular Activity (EVA) Suits		147							235.0	11.8
21.	Supporting Communications/Navigation Satellites				2273					712.3	424.0
22.	RS-68-Class Engine	14876								4740.1	70.0
23.	J-2X-Class Engine	5450								2728.6	36.0
24.	RL-10B-2-Class Engine	610								818.2	8.5
25.	AJ-10-Class Engine	275								527.9	5.0

B.1.2. Development and First-Period Production Cost Estimation Logic

The NASA JSC SVLCM and SMAD data provide the building blocks around which a basic logic can be structured to provide cost estimation for development and first-period production, taking into account any pre-existing capabilities from the starting-point configuration. This logic is summarized in Figure 98: Given a transition from one configuration to another (each of which is itself an architecture triplet), checks are first performed to ensure that transition rules are not violated (for details, see Section 6.1.2). Provided that the transition is allowed, the logic proceeds through each of the twenty-five components listed in Table 31, checking first to see if the component is needed in the development architecture of the new configuration. If so, the previous configuration is examined to ascertain whether the component existed in operations memory, or in a just-completed phase of development (if applicable). If so, DDT&E costs need not be incurred since the component has already been developed; only first-period production (TFU) costs are incurred and spread appropriately according to the particular period of development and the distribution in Figure 97. If the component does not exist in previous memory, operations, or just-completed development, it must be developed, and DDT&E costs are incurred as well (and distributed according to Figure 97). If the component does not exist in the new development architecture, no cost is incurred for development or first-period production. All component DDT&E and TFU and costs are drawn directly from the last two columns of Table 44, and production learning effects are assumed to be negligible (in part due to the findings of the validation in Section B.6).

The basic consequence of this development and first-period production cost estimation logic is that any components that exist in operations, memory, or just-completed development of a current configuration need not be re-developed if development in the next period calls for them. As a result, this allows the modeling of cost benefits that may incurred as a consequence of incremental development or the selection of architectures with common components from period to period.

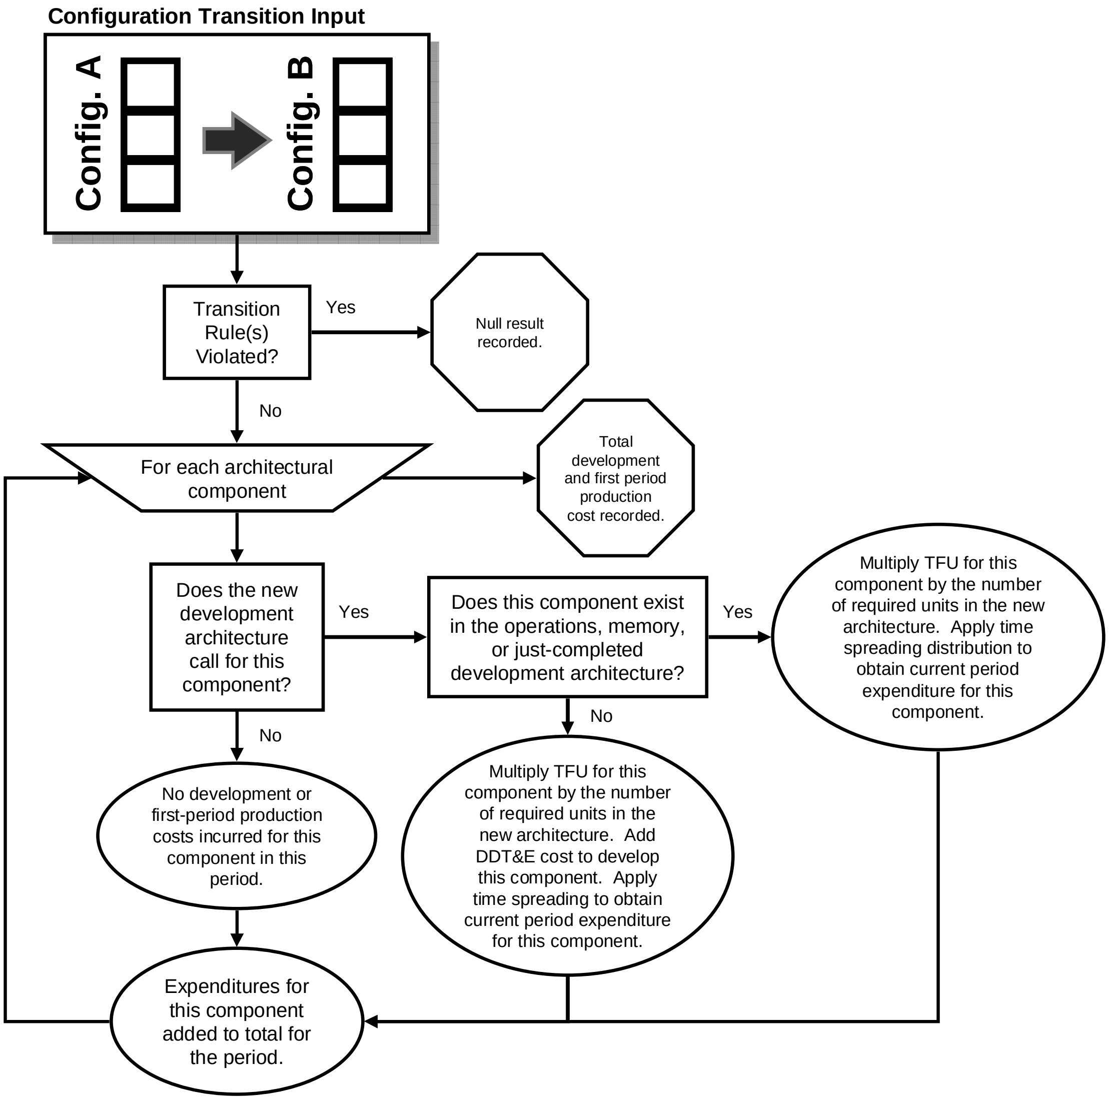
Configuration Transition Input
Config. A
Config. B
Transition Rule(s) Violated?
Yes
Null result recorded.
No
For each architectural component
Total development and first period production cost recorded.
Does the new development architecture call for this component?
Yes
Does this component exist in the operations, memory, or just-completed development architecture?
Yes
Multiply TFU for this component by the number of required units in the new architecture. Apply time spreading distribution to obtain current period expenditure for this component.
No
No
No development or first-period production costs incurred for this component in this period.
Multiply TFU for this component by the number of required units in the new architecture. Add DDT&E cost to develop this component. Apply time spreading to obtain current period expenditure for this component.
Expenditures for this component added to total for the period.

Figure 98. Summary of Development and First-Period Production Cost Estimation Logic.

B.2. Program Management and Systems Engineering Costs

Added to the development effort cost estimates is an estimate for accompanying program management and systems engineering costs based on historical data. This cost estimating relationship, based on the mean of historical data [151], adds 40.8% to the DDT&E and TFU costs incurred in a given period to a system in development and includes the sum of ground support equipment (GSE), integration, assembly, and checkout (IACO), launch and orbital operations support (LOOS), program management (PM), systems engineering and integration (SE&I), and system test operations (STO).

B.3. Recurring Production Costs

Beyond development and first-period production costs, the transition cost model estimates the costs of recurring production of operations architectures. This production is estimated directly from the component TFU costs of Table 44 and is assumed to occur for each component of the next operations architecture (i.e., to allow this architecture to continue operation into the following period, if the decision-maker chooses to do so) unless the development architecture is in the fourth and final phase of development. As in the case of the estimation of first-period production, learning effects are assumed to be negligible.

B.4. Ground and Mission Operations Costs

Also included in the transition cost estimates are the costs of ground and mission operations. These costs are estimated using the parametric NASA JSC Mission Operations Cost Model (MOCM) [152], which takes as an input system investment (DDT&E + TFU) cost c_{inv} and outputs an annual estimate $c_{ops,annual}$ for the sum of ground and mission operations costs. Modified such that it is anchored to the Space Shuttle's $33.9 billion (in FY11 dollars; or $5.97 billion in 1972 dollars [153]) investment cost and average annual $2.43 billion ground and mission operations cost (averaged from the years 2001-2003 and 2007-2010 [154], with production activities excluded from the average), the equation for this model, with inputs and outputs in millions of FY11 dollars, is provided in Eq. (B3).

$$c_{ops,annual} = 0.676 c_{inv}^{0.785} \tag{B3}$$

To apply this model to a given configuration transition, the DDT&E and TFU costs of each component in the new operations architecture are summed and used as the investment cost input to Eq. (B3). The annual operations cost output, shown for each operations architecture in Table 45, is multiplied by the two-year duration of the period and output as a contribution to the total cost in the cost transition matrix.

Table 45. Operations cost model inputs and outputs for each operations architecture.

Operations Architecture	System Investment Cost (c_{inv}, $FY11B)	Annual Operations Cost ($c_{ops,annual}$, $FY11B)	Per-Period Cost, $FY11B
-1- Nothing	0.0	0.00	0.00
-2- LEO	16.6	1.39	2.78
-3- GEO Servicing	38.9	2.71	5.42
-4- Lunar Orbit	28.1	2.10	4.12
-5- Lunar Surface	45.8	3.08	6.16
-6- Sun-Earth L2	40.9	2.82	5.64
-7- Near-Earth Object	38.9	2.71	5.42
-8- Mars Moon	39.4	2.74	5.48
-9- Mars Surface	62.4	3.92	7.85
-10- Deep Space	43.5	2.96	5.92

B.5. Shutdown Costs

The final component of the transition cost model for the NASA application estimates costs in the event that a transition decision is made that involves either the retirement of a current operations architecture or the termination of an architecture program in mid-development. Though in general these costs have seen the least attention in the field of parametric cost modeling, considering them in a decision model helps to simulate the effect of program inertia that is observable (at least anecdotally) in many applications.

B.5.1. Retirement Costs

The first type of shutdown cost is incurred for components in the operations architecture of a current configuration that are not required at all in the operations architecture of a subsequent configuration. In such a situation, it is assumed that a certain retirement cost is incurred in the coming period (e.g., to shut down production lines, safe remaining flight hardware, and close out contracts). With limited data and no parametric models available to NASA cost analysts to address this cost, for the purposes of the present cost model this cost is modeled as a Shuttle-program-derived percentage of the total MOCM operations cost estimate for the components being retired. That is, the total investment cost of all components being retired in a given period are input in Eq. (B3), and based on the total projected post-2011 Space Shuttle expenditures [109] as a percentage of the average Shuttle annual ground and mission operations cost, 29.2% of the MOCM estimate is accounted as the applicable retirement cost.

B.5.2. Termination Liability Costs

The second type of shutdown cost is incurred for development projects that are terminated prior to reaching the final phase of development. This occurs in the case where the development architecture for a present configuration is not in the final phase of development and the development architecture for the next configuration does not require components from the present configuration. In such a case, component development programs have been terminated premature to their completion, and government agencies typically incur termination liability costs. These costs cover contract requirements and damages that accrue from the cancellation of contracts. Although practices vary with specific programs (and again, no parametric models exist), typical rules of thumb are that 10% of to-go program costs are bookkept termination liability costs [155]-[156]. This is the guideline used for the present transition cost model: In the situation where development has been terminated, 10% of the remaining planned development cost-to-go (i.e., the out-years in Figure 97) is incurred in the next period for each terminated component.

B.6. Validation

To validate results of the NASA JSC SVLCM in the present application, the architectural component mass estimates of Table 44 are used to replicate an 11 × 11 cost transition matrix that was manually populated by an experienced NASA JSC cost analyst in August 2010 for 11 architectures (many of which were precursors to those considered in Table 31) on four-year development timelines over four-year period lengths. The only common assumptions between the current transition cost model and the JSC cost analyst's estimation process are the definitions of the names and numbers of the components in each architecture; cost estimating techniques differ (in general, the JSC cost analyst's techniques can be regarded as higher-fidelty and non-parametric), and vehicle-specific mass and other assumptions are independently estimated.

Results from the JSC cost analyst account only for development and production costs and assume no learning effects after production of the first unit; as a result of the latter assumption, the best match in results under the standard learning curve paradigm is found to occur under the assumption of negligible learning effects. Also, the results from the JSC cost analyst are provided in a normalized form, and thus one explicit degree of freedom exists in the scaling factor required to match the dollar-valued transition costs from the present model to the normalized JSC estimates. To accomplish this, the value of the units scaling factor is selected such that the sum of the squared errors between the elements of the model-calculated cost transition matrix and the scaled JSC cost transition matrix is minimized. The resulting element-by-element percent discrepancies (model-predicted minus actual JSC estimate, normalized to the JSC estimate) are shown in Table 46, and the distribution of the absolute values of these errors is shown in Figure 99. Note that nearly 75 percent of these errors fall below 20%, and all fall below 34.9%. Considering the application of this cost model toward conceptual phases of planning (and the level of independence in the formation of these cost estimates), this level of agreement is considered acceptable.

Table 46. Discrepancies, expressed in percent, between the transition cost model with a best-fit single scaling factor and a JSC cost analyst's manual estimate for transition costs among 11 reference architectures. *Positive values indicate model overprediction relative to the prediction of the JSC cost analyst. Both the model and analyst agreed that the first column of the matrix consists of zeros, and thus this is bookkept as 0% error rather than a divide-by-zero error.*

		Validation Architecture										
		V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11
Validation Architecture	V1	0.0%	-7.8%	25.2%	-5.7%	1.3%	2.7%	-3.6%	1.7%	23.0%	22.9%	11.9%
	V2	0.0%	-34.9%	32.7%	-6.3%	-2.1%	0.7%	-4.9%	1.7%	27.8%	27.7%	17.5%
	V3	0.0%	-22.8%	-0.3%	-16.5%	-19.8%	-11.9%	-22.0%	-13.6%	13.5%	1.8%	-0.9%
	V4	0.0%	-29.6%	0.7%	-17.6%	-21.8%	-13.0%	-20.9%	-12.0%	16.8%	17.3%	8.1%
	V5	0.0%	-30.8%	-0.3%	-18.5%	-27.0%	-16.6%	-25.5%	-15.5%	16.0%	17.9%	7.0%
	V6	0.0%	-30.8%	5.4%	-16.8%	-23.4%	-19.5%	-29.5%	-18.4%	18.2%	20.2%	8.0%
	V7	0.0%	-22.8%	5.4%	-15.0%	-19.7%	-16.1%	-29.5%	-18.4%	18.2%	18.5%	8.0%
	V8	0.0%	-22.8%	5.4%	-15.0%	-19.7%	-16.1%	-29.5%	-18.4%	18.2%	18.5%	8.0%
	V9	0.0%	-22.8%	-0.3%	-16.7%	-23.2%	-16.1%	-29.5%	-18.4%	15.4%	15.8%	6.2%
	V10	0.0%	-22.8%	-0.3%	-17.5%	-23.2%	-16.1%	-29.5%	-18.4%	15.4%	16.0%	6.3%
	V11	0.0%	-21.7%	0.7%	-9.2%	-22.6%	-15.6%	-28.9%	-17.9%	16.0%	16.5%	14.9%

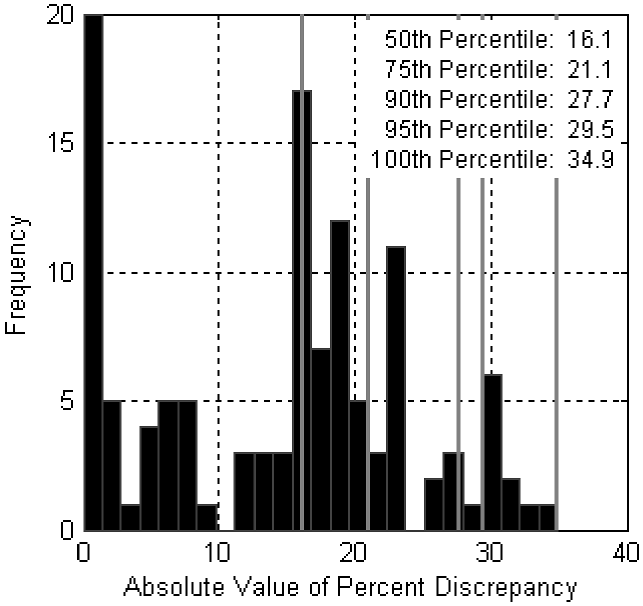

Figure 99. Histogram of absolute value of discrepancies from Table 46. *Vertical gray lines indicate locations of 50th, 75th, 90th, 95th, and 100th percentile errors.*

It should further be noted that an in-depth discussion with the JSC cost analyst who produced the full set of 121 manual estimates indicated that different equipment commonality assumptions had been used in the costing of architectures destined for Mars (in the validation here, this corresponds to architectures V9, V10, and V11). Motivated by this and the fact that the V9, V10, and V11 columns of Table 46 have development and production costs that are systematically overestimated by the present model, it is of some relevance to consider solving for two (instead of just one) best-fit scaling parameters. Applying the first parameter to the first eight columns and the second parameter to the last three columns produces the new percent discrepancies in Table 47 and the absolute value distribution in Figure 100. Note that the new distribution in general has substantially smaller errors; for example, about 95 percent of transition cost matrix elements agree within 20%. This suggests even more strongly that the present transition cost model is appropriate for the conceptual phases of program planning for which it is used in this thesis.

Table 47. Discrepancies, expressed in percent, between the transition cost model with two best-fit scaling factors and a JSC cost analyst's manual estimate for transition costs among 11 reference architectures. *Positive values indicate model overprediction relative to the prediction of the JSC cost analyst. Both the model and analyst agreed that the first column of the matrix consists of zeros, and thus this is bookkept as 0% error rather than a divide-by-zero error.*

		Validation Architecture										
		V1	V2	V3	V4	V5	V6	V7	V8	V9	V10	V11
Validation Architecture	V1	0.0%	4.2%	41.4%	6.6%	14.5%	16.1%	8.9%	14.9%	8.1%	8.0%	-1.7%
	V2	0.0%	-26.5%	49.9%	5.9%	10.6%	13.8%	7.5%	14.9%	12.3%	12.2%	3.2%
	V3	0.0%	-12.8%	12.7%	-5.7%	-9.3%	-0.4%	-11.9%	-2.4%	-0.2%	-10.6%	-13.0%
	V4	0.0%	-20.5%	13.8%	-6.9%	-11.7%	-1.7%	-10.6%	-0.6%	2.7%	3.1%	-5.0%
	V5	0.0%	-21.8%	12.7%	-7.8%	-17.5%	-5.7%	-15.9%	-4.6%	1.9%	3.6%	-6.0%
	V6	0.0%	-21.8%	19.1%	-5.9%	-13.4%	-9.0%	-20.3%	-7.8%	3.9%	5.6%	-5.1%
	V7	0.0%	-12.8%	19.1%	-3.9%	-9.2%	-5.2%	-20.3%	-7.8%	3.9%	4.1%	-5.1%
	V8	0.0%	-12.8%	19.1%	-3.9%	-9.2%	-5.2%	-20.3%	-7.8%	3.9%	4.1%	-5.1%
	V9	0.0%	-12.8%	12.7%	-5.8%	-13.2%	-5.2%	-20.3%	-7.8%	1.4%	1.8%	-6.7%
	V10	0.0%	-12.8%	12.7%	-6.7%	-13.2%	-5.2%	-20.3%	-7.8%	1.4%	2.0%	-6.6%
	V11	0.0%	-11.5%	13.8%	2.6%	-12.6%	-4.7%	-19.7%	-7.2%	1.9%	2.4%	0.9%

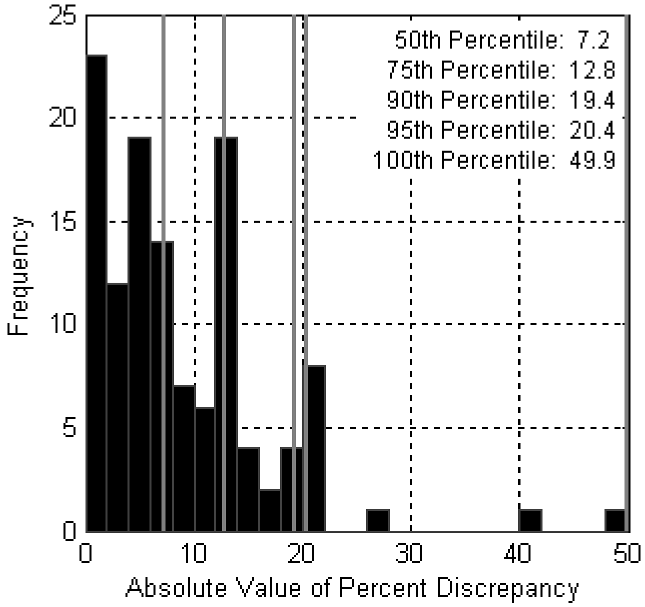

Figure 100. Histogram of absolute value of discrepancies from Table 47. *Vertical gray lines indicate locations of 50th, 75th, 90th, 95th, and 100th percentile errors.*

B.7. Summary

One assumption that this thesis makes is that an appropriate cost model exists via which a cost transition matrix may be populated. However, in many applications, this may not be true, and substantial work may need to occur to create such a model. This appendix has described the construction of a transition cost model for the human space exploration example discussed in Chapter 6 of this thesis. With the purpose of converting an input of two configurations (a "from" configuration and a "to" configuration) to a one-period transition cost, this particular model includes estimation of the costs of development, production, operations, program management and systems engineering, retirement, and termination liability. Also included in this appendix has been a discussion of the results of a validation against independently-generated NASA transition costs.

It should be mentioned that several challenges were encountered in the creation of this model. In particular, while parametric cost models for space vehicle development and production have received substantial attention in the space industry, few such models were found for other relevant costs. For example, recommended future investigations would include reexamination of the NASA MOCM's structure that uses system investment cost as a sole input for the estimation of operations costs. Such a structure does not allow the possibility, for instance, that a system decision-maker can implement design options that result in high development and low operations costs or vice versa. While far more detailed operations cost models do exist (for example, the Exploration Architecture Operations Cost Model developed by the NASA Jet Propulsion Laboratory [157]-[158]), the time, personnel resources, and large number of data inputs required to produce an estimate using a detailed costing tool tend not to be conducive to rapid parametric analysis. A need appears to exist for a fidelity level between these extremes in operations cost modeling. However, perhaps the most obvious gap in cost modeling capability exists for shutdown costs, both in terms of retirement and termination liability costs. In the present cost model, both were estimated based on limited historical data and rules of thumb. Clearly this is an area in need of future work for the cost estimation community, as these high costs (or avoidance of these high costs in favor of alternatives) produce an inertia to continue with current programs that is not negligible in decision-making.

As developments continue in the areas of cost estimation identified above, relevant components of the cost model developed here can be updated and re-applied to the analysis in Chapter 6 of this thesis to continually improve the quality of human space exploration system analysis and decision-making. Overall, the capabilities provided by this transition cost model are intended to permit estimation of the elements of the cost transition matrix required by Step 1 of this thesis' framework to a level of accuracy and comprehensiveness appropriate for conceptual design and program planning.

APPENDIX C

EXPERT INPUTS FOR HUMAN SPACE EXPLORATION STOCHASTIC SYSTEM DECISION MODELING

Two assumptions inherent to the framework proposed in this thesis are that (1) meaningful quantitative figures of merit (or surrogate figures of merit) exist for the decision problem of interest and (2) it is possible to meaningfully specify the probability transition matrix of the underlying Markov chain for demand environment evolution. While step 1 of the framework relies only upon the relatively mature discipline of cost estimation, the remainder of the framework cannot be applied without these two elements. The following appendix describes a substantial effort undertaken in the course of this work to address and obtain information and estimates for these elements in the context of the NASA human space exploration application.

C.1. Survey Description

The two assumptions above are inherently related to the preferences and perspectives of the decision-making body. For example, there is no single objectively "correct" figure of merit. In the case of the Markov chain probabilities, although they may be considered objective quantities when enough historical data exists, in many applications the lack of sufficient historical data must be substituted with expert estimates.

To address these two assumptions for the NASA exploration application, a survey was distributed to a group of 21 personnel with substantial experience in the field of human space exploration. This group was intended to simulate the opinions, beliefs, and preferences of a senior NASA decision-making body. As detailed in Table 48 and Figure 101, these individuals represented a total of 8 NASA centers plus one external organization; particularly high representation was accorded to Johnson Space Center due to its specialization in human spaceflight activities.

Table 48. Survey Invitee Affiliations.

NASA Center Affiliation	No. of Invitees	Percent of Total
Headquarters (HQ)	2	9.5%
Johnson Space Center (JSC)	9	42.9%
Marshall Space Flight Center (MSFC)	1	4.8%
Langley Research Center (LaRC)	2	9.5%
Glenn Research Center (GRC)	1	4.8%
Goddard Space Flight Center (GSFC)	1	4.8%
Ames Research Center (ARC)	2	9.5%
Jet Propulsion Laboratory (JPL)	2	9.5%
Non-NASA	1	4.8%
TOTAL	21	100.0%

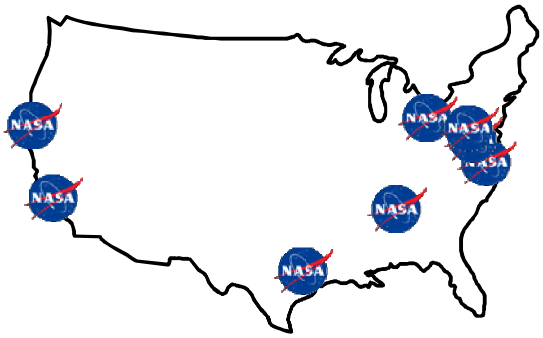

Figure 101. Geographic view of invitee NASA center affiliations.

The survey was distributed to these invitees via a recruitment E-mail that contained a link to a central website (http://www.flexibility.gatech.edu). At this website, invitees found instructions on downloading and later submitting the survey in the form of a Microsoft Excel file. The survey itself asked participants to consider NASA's need to decide which space systems to develop to meet potential future human spaceflight mission demands or expectations. After documenting their consent to voluntarily participate in the study, participants were asked to provide specific inputs on (1) the relevance of various figures of merit and (2) the likely evolution of future mission demands. The estimated time required to complete the survey was 55 minutes, and upon completion the participants were directed to submit their completed Excel file through an anonymous online web form at http://www.flexibility.gatech.edu/submit.php.

As detailed in Table 49, the process of obtaining the appropriate human subjects research training, applying for the needed web domain, and preparing the survey materials began in April 2011. The intended final version of the survey was submitted to the Georgia Institute of Technology Institutional Review Board (IRB) on May 31. The anonymity of the survey permitted the protocol to be approved and classified as exempt from further IRB review on June 23. Survey invitations were sent via E-mail on June 27-28, a reminder E-mail was sent on July 12, and surveys were compiled for analysis on July 15.

Table 49. Survey Activity Timeline.

Date	Event
April 12, 2011	CITI Human Subjects Training Completed
April 19, 2011	Survey Domain Assigned (http://www.flexibility.gatech.edu)
May 15, 2011	First Draft of Survey and Website Completed
May 31, 2011	Survey Protocol Submitted to Georgia Tech IRB
June 23, 2011	Protocol Approved (Protocol No. H11172)
June 27-28, 2011	Survey Invitations Sent
July 12, 2011	Survey Reminder Sent
July 15, 2011	Survey Deadline

The recruitment and reminder E-mails to participants are copied below, as are screenshots from the submission website (Figure 102 and Figure 103) and the survey Excel file (Figure 104 through Figure 107). Also included (Figure 108 and Figure 109) are screenshots of dialog boxes intended to assist participants in filling out the demand evolution section of the survey. The dialog box in Figure 108 is activated upon clicking the “Use Wizard Assistance for Part I” button in Figure 106, and the dialog box in Figure 109 is activated upon clicking the “Use Wizard Assistance for Part II” button in Figure 106. Supporting these dialog boxes and Excel sheets are approximately 700 lines of Visual Basic source code. The survey download and submission website is supported by approximately 600 lines of HTML and PHP source code.

Recruitment E-mail (June 27-28)

Subject: Expert Inputs for Human Space Exploration Stochastic Decision Modeling
From: Jarret Lafleur (jarret.m.lafleur@gatech.edu)

Greetings,

Based on your experience and expertise in NASA human spaceflight program planning and systems engineering, I would like to formally invite you to take part in a survey being conducted by the Space Systems Design Laboratory at the Georgia Institute of Technology. This survey asks you to consider NASA's need to decide which space systems to develop in order to meet potential future human spaceflight mission demands or expectations. In particular, you will be asked about the relevance of various figures of merit as well as the likely evolution of future mission demands. The results collected will be used toward an example application in the development of a new approach to integrating flexibility in space system design decisions.

Personally identifiable information is not collected in the course of the survey, and your identity will not be associated with any results you submit. The total time to complete the survey is estimated at 55 minutes; note, however, that you will be able to save your work for completion among multiple time increments if necessary.

Your participation would be very much appreciated. If you choose to participate, please navigate to the following Internet URL by July 15: http://www.flexibility.gatech.edu/

Thank you in advance. If you have any questions, please feel free to contact me at jarret.m.lafleur@gatech.edu.

Sincerely,

Jarret Lafleur
jarret.m.lafleur@gatech.edu

Ph.D. Candidate
School of Aerospace Engineering
Georgia Institute of Technology

Reminder E-mail (July 12)

Subject: Reminder: July 15 Deadline for Human Space Exploration Decision Modeling Survey
From: Jarret Lafleur (jarret.m.lafleur@gatech.edu)

All,

About two weeks ago, you may have received an invitation to participate in a survey regarding human space exploration figures of merit and the evolution of possible exploration mission demands.

To those who have already submitted survey responses, thank you very much. For those who are interested in participating but have not yet submitted a survey, your response would be much appreciated by the end of the day this Friday, July 15. While the survey submission site will remain up and running, surveys received after this date are not guaranteed to be incorporated into the results.

Personally identifiable information is not collected in the course of the survey, and your identity will not be associated with any results you submit. The total time to complete the survey is estimated at 55 minutes, and you will be able to save your work for completion among multiple time increments if necessary.

If you choose to participate, please navigate by July 15 to: http://www.flexibility.gatech.edu/

Thank you in advance. As before, if you have any questions, please feel free to contact me at jarret.m.lafleur@gatech.edu.

Sincerely,

Jarret Lafleur
jarret.m.lafleur@gatech.edu

Ph.D. Candidate
School of Aerospace Engineering
Georgia Institute of Technology

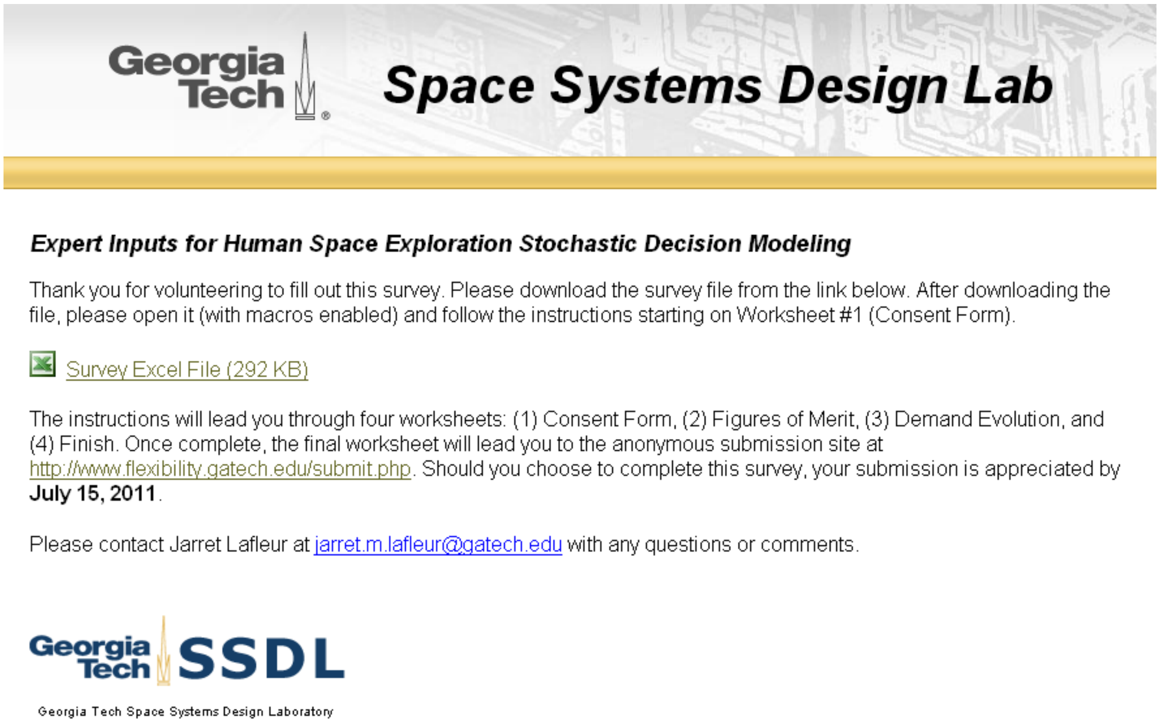

Expert Inputs for Human Space Exploration Stochastic Decision Modeling

Thank you for volunteering to fill out this survey. Please download the survey file from the link below. After downloading the file, please open it (with macros enabled) and follow the instructions starting on Worksheet #1 (Consent Form).

Survey Excel File (292 KB)

The instructions will lead you through four worksheets: (1) Consent Form, (2) Figures of Merit, (3) Demand Evolution, and (4) Finish. Once complete, the final worksheet will lead you to the anonymous submission site at http://www.flexibility.gatech.edu/submit.php. Should you choose to complete this survey, your submission is appreciated by **July 15, 2011**.

Please contact Jarret Lafleur at jarret.m.lafleur@gatech.edu with any questions or comments.

Figure 102. Screen Shot of Main Survey Website (http://www.flexibility.gatech.edu).

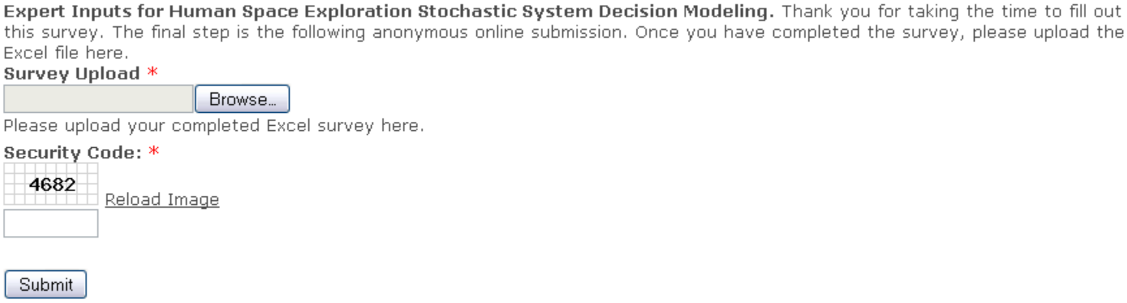

Expert Inputs for Human Space Exploration Stochastic System Decision Modeling. Thank you for taking the time to fill out this survey. The final step is the following anonymous online submission. Once you have completed the survey, please upload the Excel file here.

Survey Upload *

Please upload your completed Excel survey here.

Security Code: *

Figure 103. Screen Shot of Survey Submission Website (http://www.flexibility.gatech.edu/submit.php).

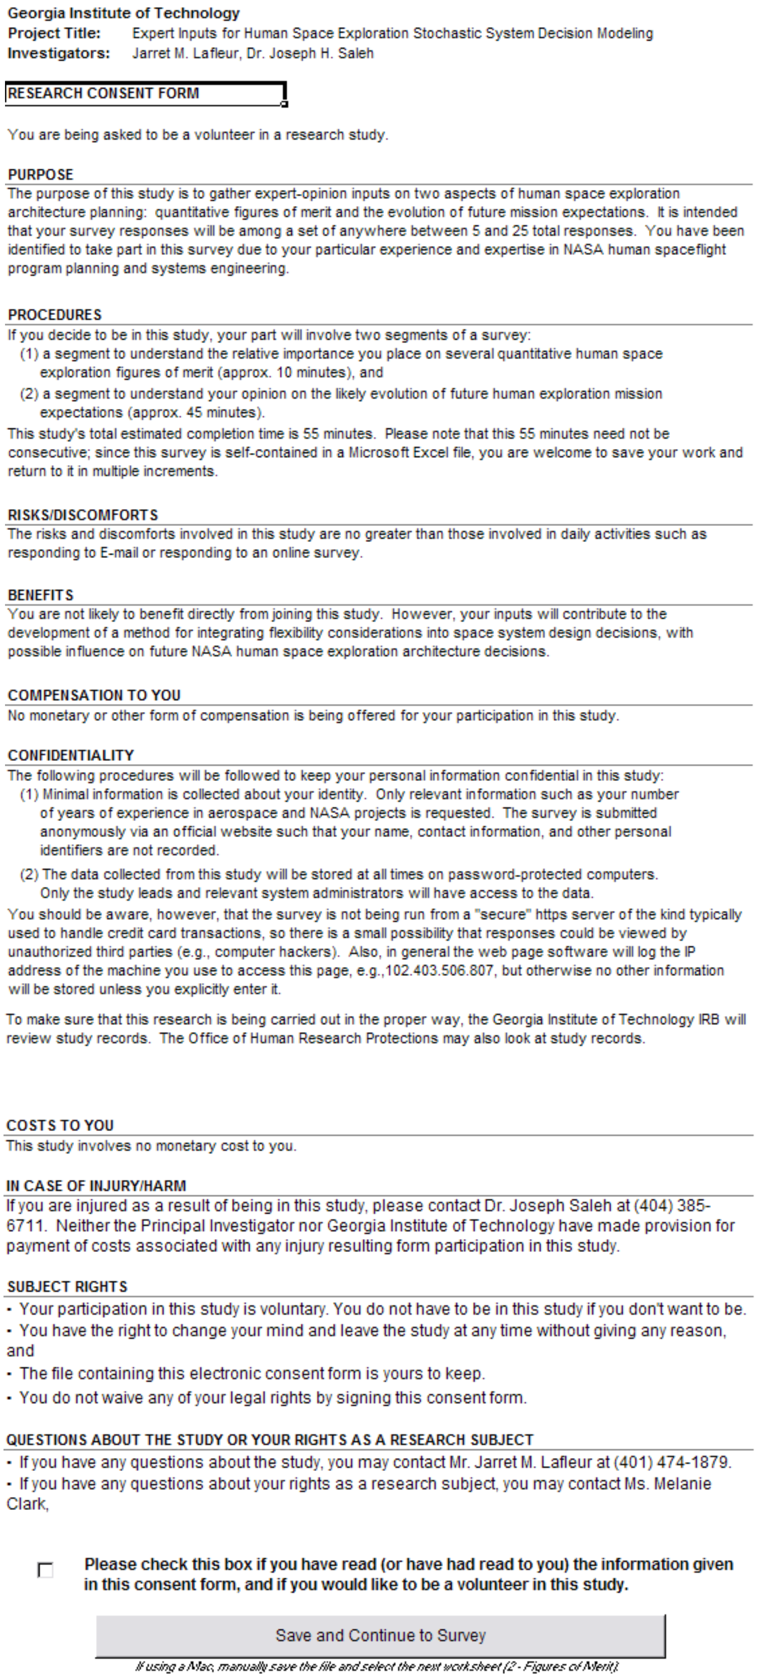

Georgia Institute of Technology
Project Title: Expert Inputs for Human Space Exploration Stochastic System Decision Modeling
Investigators: Jarret M. Lafleur, Dr. Joseph H. Saleh

RESEARCH CONSENT FORM

You are being asked to be a volunteer in a research study.

PURPOSE

The purpose of this study is to gather expert-opinion inputs on two aspects of human space exploration architecture planning: quantitative figures of merit and the evolution of future mission expectations. It is intended that your survey responses will be among a set of anywhere between 5 and 25 total responses. You have been identified to take part in this survey due to your particular experience and expertise in NASA human spaceflight program planning and systems engineering.

PROCEDURES

If you decide to be in this study, your part will involve two segments of a survey:

(1) a segment to understand the relative importance you place on several quantitative human space exploration figures of merit (approx. 10 minutes), and

(2) a segment to understand your opinion on the likely evolution of future human exploration mission expectations (approx. 45 minutes).

This study's total estimated completion time is 55 minutes. Please note that this 55 minutes need not be consecutive; since this survey is self-contained in a Microsoft Excel file, you are welcome to save your work and return to it in multiple increments.

RISKS/DISCOMFORTS

The risks and discomforts involved in this study are no greater than those involved in daily activities such as responding to E-mail or responding to an online survey.

BENEFITS

You are not likely to benefit directly from joining this study. However, your inputs will contribute to the development of a method for integrating flexibility considerations into space system design decisions, with possible influence on future NASA human space exploration architecture decisions.

COMPENSATION TO YOU

No monetary or other form of compensation is being offered for your participation in this study.

CONFIDENTIALITY

The following procedures will be followed to keep your personal information confidential in this study:

(1) Minimal information is collected about your identity. Only relevant information such as your number of years of experience in aerospace and NASA projects is requested. The survey is submitted anonymously via an official website such that your name, contact information, and other personal identifiers are not recorded.

(2) The data collected from this study will be stored at all times on password-protected computers. Only the study leads and relevant system administrators will have access to the data.

You should be aware, however, that the survey is not being run from a "secure" https server of the kind typically used to handle credit card transactions, so there is a small possibility that responses could be viewed by unauthorized third parties (e.g., computer hackers). Also, in general the web page software will log the IP address of the machine you use to access this page, e.g.,102.403.506.807, but otherwise no other information will be stored unless you explicitly enter it.

To make sure that this research is being carried out in the proper way, the Georgia Institute of Technology IRB will review study records. The Office of Human Research Protections may also look at study records.

COSTS TO YOU

This study involves no monetary cost to you.

IN CASE OF INJURY/HARM

If you are injured as a result of being in this study, please contact Dr. Joseph Saleh at (404) 385-6711. Neither the Principal Investigator nor Georgia Institute of Technology have made provision for payment of costs associated with any injury resulting form participation in this study.

SUBJECT RIGHTS

- Your participation in this study is voluntary. You do not have to be in this study if you don't want to be.
- You have the right to change your mind and leave the study at any time without giving any reason, and
- The file containing this electronic consent form is yours to keep.
- You do not waive any of your legal rights by signing this consent form.

QUESTIONS ABOUT THE STUDY OR YOUR RIGHTS AS A RESEARCH SUBJECT

- If you have any questions about the study, you may contact Mr. Jarret M. Lafleur at (401) 474-1879.
- If you have any questions about your rights as a research subject, you may contact Ms. Melanie Clark,

☐ **Please check this box if you have read (or have had read to you) the information given in this consent form, and if you would like to be a volunteer in this study.**

If using a Mac, manually save the file and select the next worksheet (2 - Figures of Merit).

Figure 104. Consent Form (Worksheet #1) from Survey Excel File.

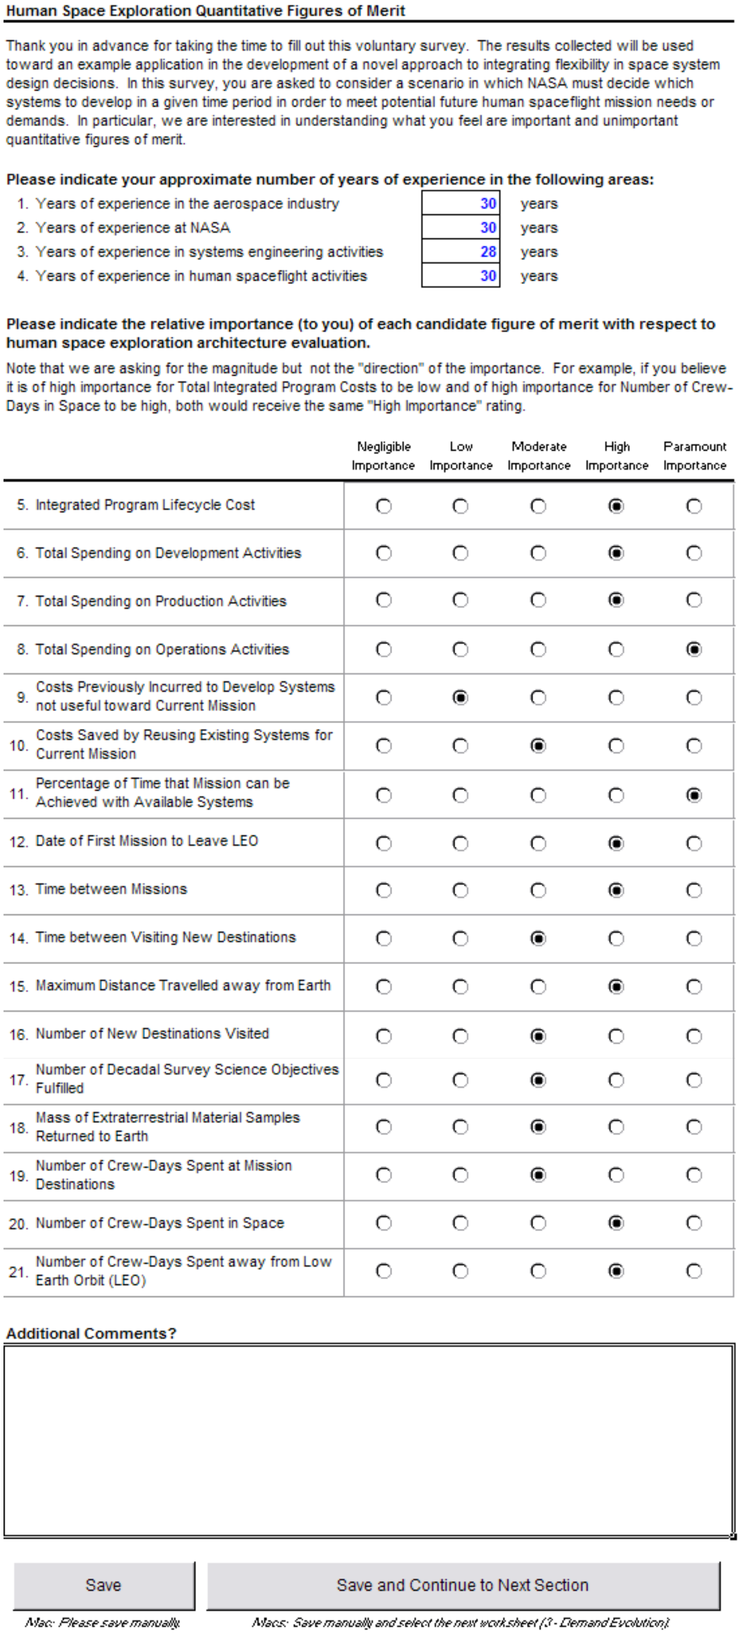

Human Space Exploration Quantitative Figures of Merit

Thank you in advance for taking the time to fill out this voluntary survey. The results collected will be used toward an example application in the development of a novel approach to integrating flexibility in space system design decisions. In this survey, you are asked to consider a scenario in which NASA must decide which systems to develop in a given time period in order to meet potential future human spaceflight mission needs or demands. In particular, we are interested in understanding what you feel are important and unimportant quantitative figures of merit.

Please indicate your approximate number of years of experience in the following areas:

1. Years of experience in the aerospace industry	30	years
2. Years of experience at NASA	30	years
3. Years of experience in systems engineering activities	28	years
4. Years of experience in human spaceflight activities	30	years

Please indicate the relative importance (to you) of each candidate figure of merit with respect to human space exploration architecture evaluation.

Note that we are asking for the magnitude but not the "direction" of the importance. For example, if you believe it is of high importance for Total Integrated Program Costs to be low and of high importance for Number of Crew-Days in Space to be high, both would receive the same "High Importance" rating.

	Negligible Importance	Low Importance	Moderate Importance	High Importance	Paramount Importance
5. Integrated Program Lifecycle Cost	○	○	○	●	○
6. Total Spending on Development Activities	○	○	○	●	○
7. Total Spending on Production Activities	○	○	○	●	○
8. Total Spending on Operations Activities	○	○	○	○	●
9. Costs Previously Incurred to Develop Systems not useful toward Current Mission	○	●	○	○	○
10. Costs Saved by Reusing Existing Systems for Current Mission	○	○	●	○	○
11. Percentage of Time that Mission can be Achieved with Available Systems	○	○	○	○	●
12. Date of First Mission to Leave LEO	○	○	○	●	○
13. Time between Missions	○	○	○	●	○
14. Time between Visiting New Destinations	○	○	●	○	○
15. Maximum Distance Travelled away from Earth	○	○	○	●	○
16. Number of New Destinations Visited	○	○	●	○	○
17. Number of Decadal Survey Science Objectives Fulfilled	○	○	●	○	○
18. Mass of Extraterrestrial Material Samples Returned to Earth	○	○	●	○	○
19. Number of Crew-Days Spent at Mission Destinations	○	○	●	○	○
20. Number of Crew-Days Spent in Space	○	○	○	●	○
21. Number of Crew-Days Spent away from Low Earth Orbit (LEO)	○	○	○	●	○

Additional Comments?

Mac: Please save manually.

Macs: Save manually and select the next worksheet (3 - Demand Evolution).

Figure 105. Figures of Merit Section (Worksheet #2) from Survey Excel File.

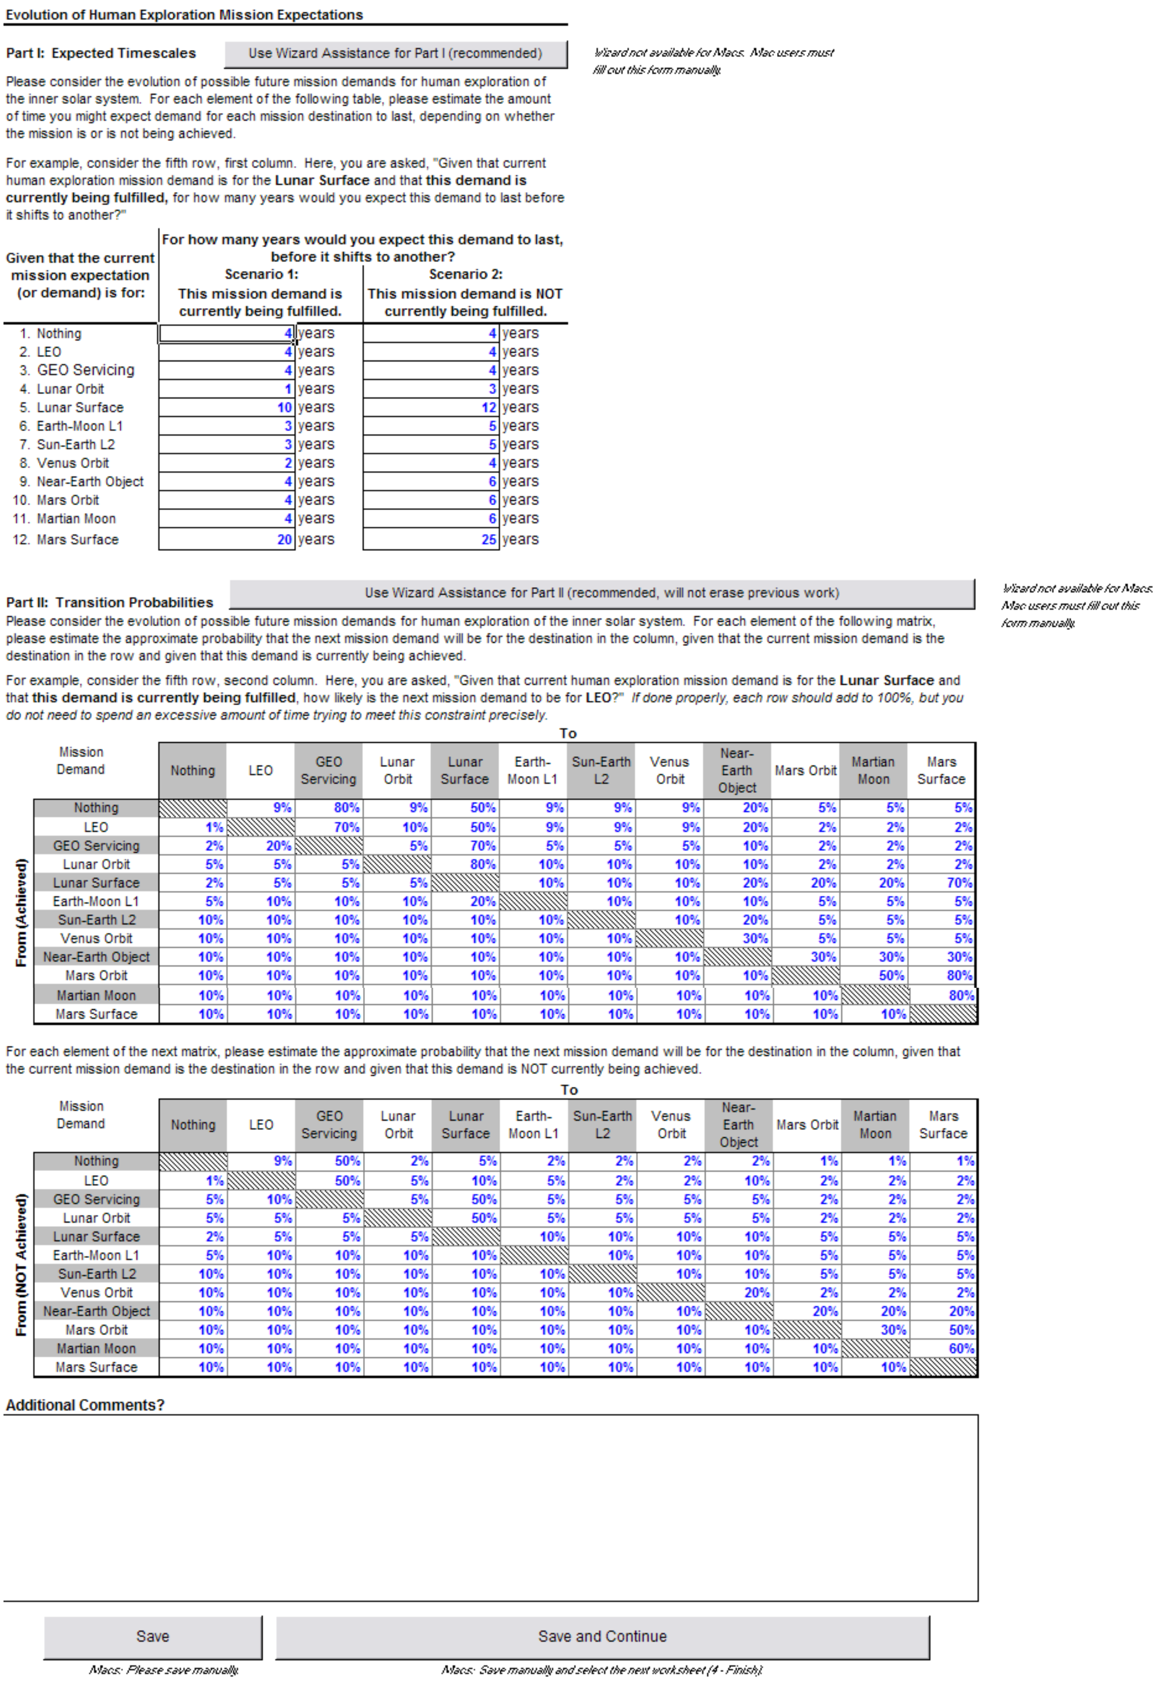

Evolution of Human Exploration Mission Expectations

Part I: Expected Timescales

Wizard not available for Macs. Mac users must fill out this form manually.

Please consider the evolution of possible future mission demands for human exploration of the inner solar system. For each element of the following table, please estimate the amount of time you might expect demand for each mission destination to last, depending on whether the mission is or is not being achieved.

For example, consider the fifth row, first column. Here, you are asked, "Given that current human exploration mission demand is for the **Lunar Surface** and that **this demand is currently being fulfilled,** for how many years would you expect this demand to last before it shifts to another?"

Given that the current mission expectation (or demand) is for:	For how many years would you expect this demand to last, before it shifts to another? Scenario 1: This mission demand is currently being fulfilled.		Scenario 2: This mission demand is NOT currently being fulfilled.	
1. Nothing	4	years	4	years
2. LEO	4	years	4	years
3. GEO Servicing	4	years	4	years
4. Lunar Orbit	1	years	3	years
5. Lunar Surface	10	years	12	years
6. Earth-Moon L1	3	years	5	years
7. Sun-Earth L2	3	years	5	years
8. Venus Orbit	2	years	4	years
9. Near-Earth Object	4	years	6	years
10. Mars Orbit	4	years	6	years
11. Martian Moon	4	years	6	years
12. Mars Surface	20	years	25	years

Part II: Transition Probabilities

Use Wizard Assistance for Part II (recommended, will not erase previous work)

Wizard not available for Macs. Mac users must fill out this form manually.

Please consider the evolution of possible future mission demands for human exploration of the inner solar system. For each element of the following matrix, please estimate the approximate probability that the next mission demand will be for the destination in the column, given that the current mission demand is the destination in the row and given that this demand is currently being achieved.

For example, consider the fifth row, second column. Here, you are asked, "Given that current human exploration mission demand is for the **Lunar Surface** and that **this demand is currently being fulfilled**, how likely is the next mission demand to be for **LEO**?" *If done properly, each row should add to 100%, but you do not need to spend an excessive amount of time trying to meet this constraint precisely.*

To

	Mission Demand	Nothing	LEO	GEO Servicing	Lunar Orbit	Lunar Surface	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Martian Moon	Mars Surface
From (Achieved)	Nothing	▨	9%	80%	9%	50%	9%	9%	9%	20%	5%	5%	5%
	LEO	1%	▨	70%	10%	50%	9%	9%	9%	20%	2%	2%	2%
	GEO Servicing	2%	20%	▨	5%	70%	5%	5%	5%	10%	2%	2%	2%
	Lunar Orbit	5%	5%	5%	▨	80%	10%	10%	10%	10%	2%	2%	2%
	Lunar Surface	2%	5%	5%	5%	▨	10%	10%	10%	20%	20%	20%	70%
	Earth-Moon L1	5%	10%	10%	10%	20%	▨	10%	10%	10%	5%	5%	5%
	Sun-Earth L2	10%	10%	10%	10%	10%	10%	▨	10%	20%	5%	5%	5%
	Venus Orbit	10%	10%	10%	10%	10%	10%	10%	▨	30%	5%	5%	5%
	Near-Earth Object	10%	10%	10%	10%	10%	10%	10%	10%	▨	30%	30%	30%
	Mars Orbit	10%	10%	10%	10%	10%	10%	10%	10%	10%	▨	50%	80%
	Martian Moon	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	▨	80%
	Mars Surface	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	▨

For each element of the next matrix, please estimate the approximate probability that the next mission demand will be for the destination in the column, given that the current mission demand is the destination in the row and given that this demand is NOT currently being achieved.

To

	Mission Demand	Nothing	LEO	GEO Servicing	Lunar Orbit	Lunar Surface	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Martian Moon	Mars Surface
From (NOT Achieved)	Nothing	▨	9%	50%	2%	5%	2%	2%	2%	2%	1%	1%	1%
	LEO	1%	▨	50%	5%	10%	5%	2%	2%	10%	2%	2%	2%
	GEO Servicing	5%	10%	▨	5%	50%	5%	5%	5%	5%	2%	2%	2%
	Lunar Orbit	5%	5%	5%	▨	50%	5%	5%	5%	5%	2%	2%	2%
	Lunar Surface	2%	5%	5%	5%	▨	10%	10%	10%	10%	5%	5%	5%
	Earth-Moon L1	5%	10%	10%	10%	10%	▨	10%	10%	10%	5%	5%	5%
	Sun-Earth L2	10%	10%	10%	10%	10%	10%	▨	10%	10%	5%	5%	5%
	Venus Orbit	10%	10%	10%	10%	10%	10%	10%	▨	20%	2%	2%	2%
	Near-Earth Object	10%	10%	10%	10%	10%	10%	10%	10%	▨	20%	20%	20%
	Mars Orbit	10%	10%	10%	10%	10%	10%	10%	10%	10%	▨	30%	50%
	Martian Moon	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	▨	60%
	Mars Surface	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	▨

Additional Comments?

Macs: Please save manually.

Macs: Save manually and select the next worksheet (4 - Finish).

Figure 106. Demand Evolution Section (Worksheet #3) from Survey Excel File.

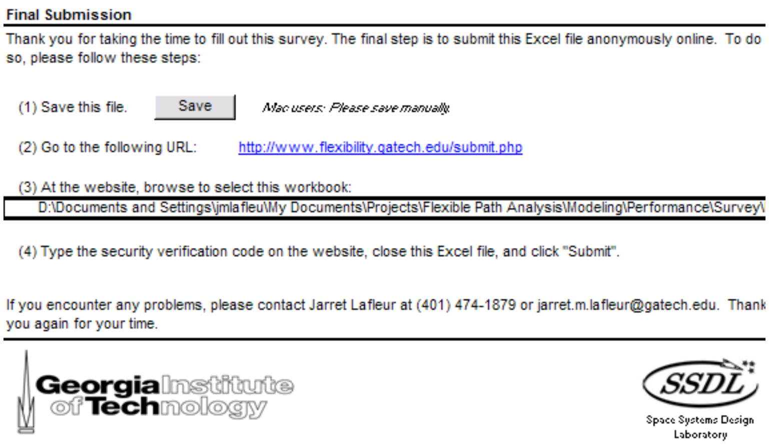

Final Submission

Thank you for taking the time to fill out this survey. The final step is to submit this Excel file anonymously online. To do so, please follow these steps:

(1) Save this file.

Mac users: Please save manually.

(2) Go to the following URL: http://www.flexibility.gatech.edu/submit.php

(3) At the website, browse to select this workbook:

D:\Documents and Settings\jmlafleu\My Documents\Projects\Flexible Path Analysis\Modeling\Performance\Survey\

(4) Type the security verification code on the website, close this Excel file, and click "Submit".

If you encounter any problems, please contact Jarret Lafleur at (401) 474-1879 or jarret.m.lafleur@gatech.edu. Thank you again for your time.

Georgia Institute of Technology
SSDL
Space Systems Design Laboratory

Figure 107. Final Submission Instructions (Worksheet #4) from Survey Excel File

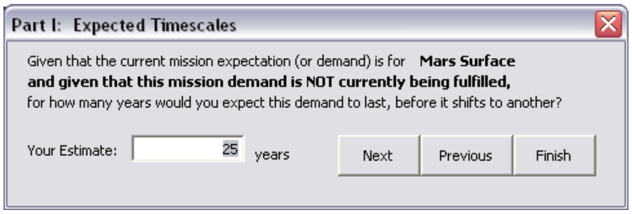

Given that the current mission expectation (or demand) is for **Mars Surface**
and given that this mission demand is NOT currently being fulfilled,
for how many years would you expect this demand to last, before it shifts to another?

Your Estimate: 25 years

Figure 108. Dialog Box for Part I of Worksheet #3.

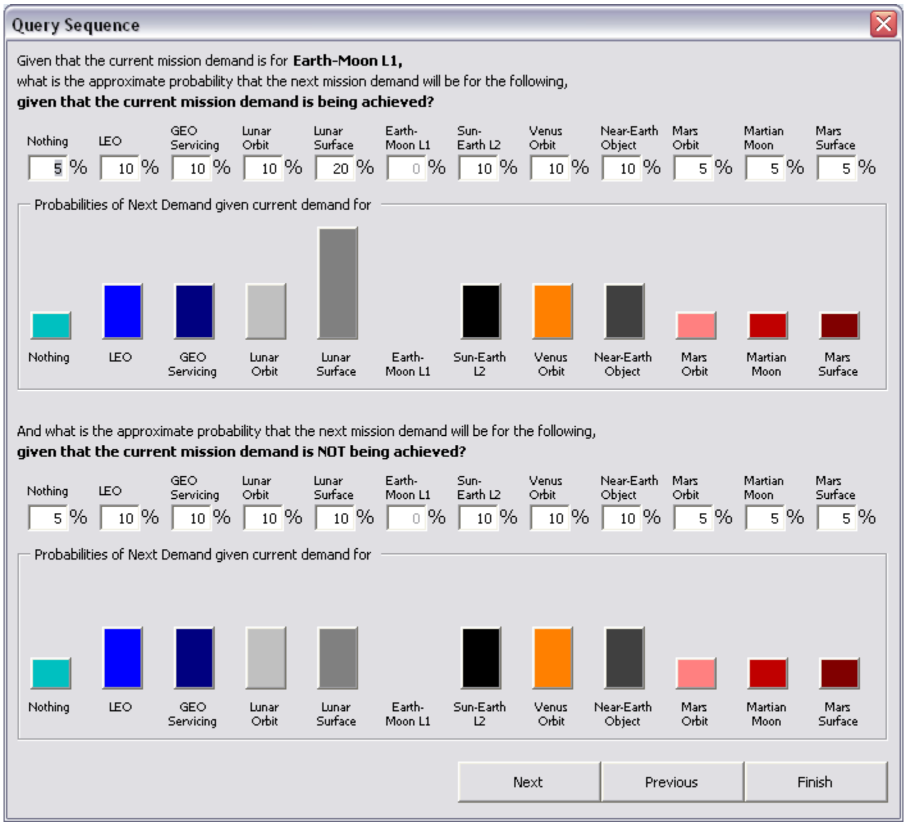

Given that the current mission demand is for **Earth-Moon L1,**
what is the approximate probability that the next mission demand will be for the following,
given that the current mission demand is being achieved?

Nothing	LEO	GEO Servicing	Lunar Orbit	Lunar Surface	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Martian Moon	Mars Surface
5 %	10 %	10 %	10 %	20 %	0 %	10 %	10 %	10 %	5 %	5 %	5 %

And what is the approximate probability that the next mission demand will be for the following,
given that the current mission demand is NOT being achieved?

Nothing	LEO	GEO Servicing	Lunar Orbit	Lunar Surface	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Martian Moon	Mars Surface
5 %	10 %	10 %	10 %	10 %	0 %	10 %	10 %	10 %	5 %	5 %	5 %

Figure 109. Dialog Box for Part II of Worksheet #3.

C.2. Results

In total, nine responses were received. The results, here presented in aggregate form, reveal several interesting insights for human space exploration planning. This section presents the results in the order in which they were requested of the participants, starting with an analysis of participant experience, continuing with an analysis of figure of merit importance, and ending with the analysis of the expert-elicited demand evolution Markov chains.

C.2.1. Participant Experience

To understand the experience level of the survey participants, four questions were asked regarding their years of experience in various areas. Aggregate responses in the form of histograms are shown in Figure 110. Overlaid on each histogram is a box and whisker plot, where the yellow box indicates the interquartile range and the whiskers extend to the minimum and maximum values of the population. The location of the median is indicated on each plot by a vertical black line, and the mean is indicated by a vertical gray line. The maximum discrepancy between the means and medians agreed in all cases within 9%, suggesting symmetry to the distributions. In summary, the average survey participant possessed about 30 years of experience in the aerospace industry, 25 years of experience at NASA and in human spaceflight activities, and 20 years of experience in systems engineering activities[*].

---

[*] Particularly notable variability exists in terms of participants' systems engineering experience, which shows an interquartile range of 18.8 years.

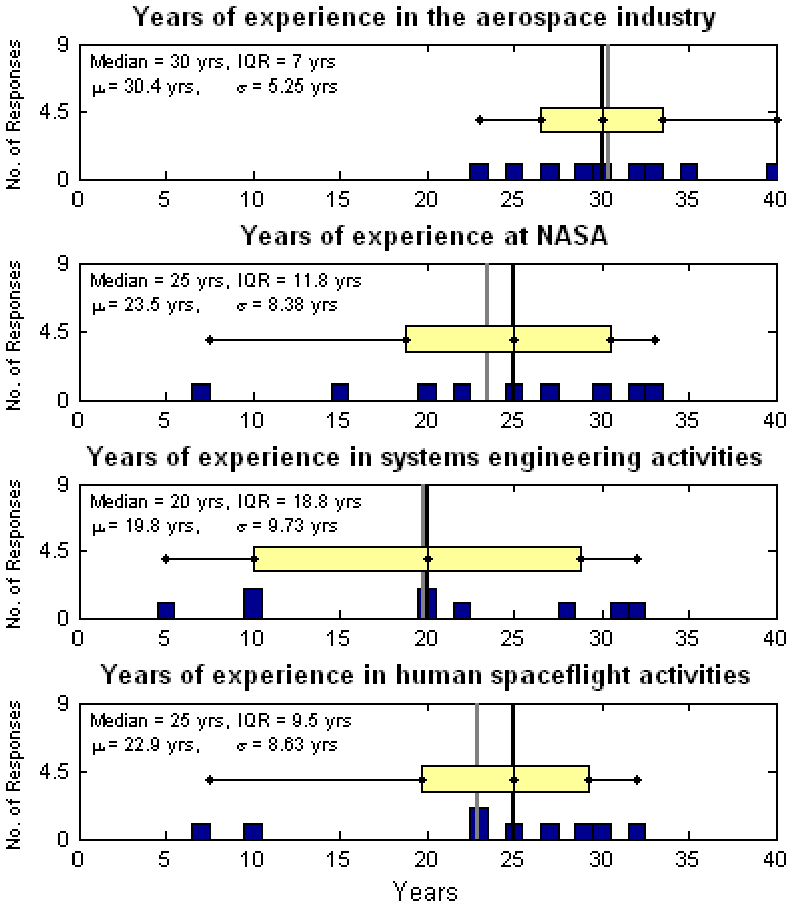

Figure 110. Distributions of Participant Years of Experience.

Figure 111 shows a multivariate plot to illustrate any existing pairwise correlations between participant years of experience in each of the categories in Figure 110. Each dot on each graph indicates one participant's set of experiences, and the gray line indicates the best-fit least-squares regression line through the point. The corresponding linear equation and coefficient of determination (R^2) value is indicated in the upper left corner of each graph. In most cases, correlations are weak though generally positive (as expected). The notable exception is the correlation between years of experience at NASA and years of experience in human spaceflight activities, which has an R^2 value of 0.74 and suggests that the average participant has nearly 9 years of human spaceflight experience for every 10 years of NASA experience. The implication of this correlation is that survey participants within NASA were successfully targeted from within the human spaceflight domains of expertise.

Also shown in Figure 111 are four large dots of different colors which will have bearing on later analysis. If it is accepted that the ideal participant, given the information available from the survey, would have the maximum amount of experience (here, 40 years) in each category, then a set of experts with non-dominated sets of experience can be attained by applying a Pareto filter. In the application of this filter, each participant is compared with each other participant. If, in these comparisons, one participant has fewer years of experience in every category than a second, then the first is filtered out and does not become part of the non-dominated set of experts. The result is a set of four experts, three of whom are part of the set because they possess the maximum years of experience in one or more categories. The fourth represents a non-dominated balance of experience (35 years in the aerospace industry, 20 years at NASA, 5 years in systems engineering, and 23 years in human spaceflight). While the vast majority of the analysis presented will aggregate the inputs of *all* survey participants, a small portion of the analysis will examine differences between the overall results and those from the non-dominated set of experts.

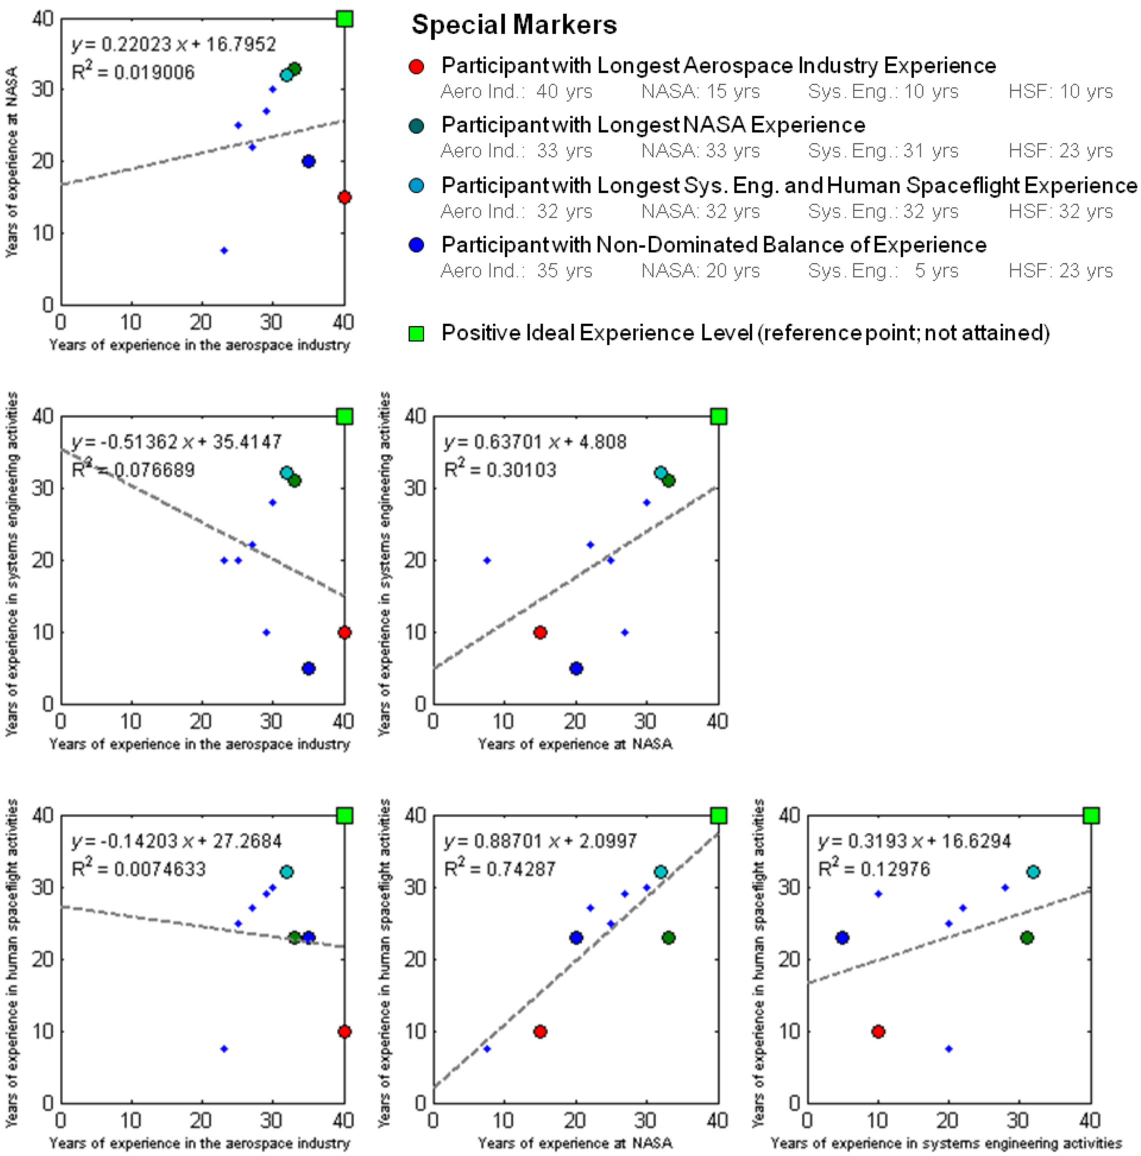

Figure 111. Multivariate Plot Illustrating Years of Experience Correlations.

C.2.2. Figures of Merit

The first major section of the survey requested that participants rate the relative importance of 17 candidate figures of merit for human spaceflight architecture evaluation. For each figure of merit, participants were given the option to rate its importance on a 5-level Likert scale with levels labeled “Negligible”, “Low”, “Medium”, “High”, and “Paramount”. The aggregate results are shown with box-and-whisker plots in Figure 112. In this figure, each figure of merit is associated with a single yellow box and set of black whiskers. Each box represents the interquartile range and each whisker extends to the minimum and maximum range of the responses. The median responses are indicated by vertical red lines, and the means are indicated by vertical gray lines; in cases where no gray line is visible, the mean response is identical to the median.

Several interesting observations can be made regarding these results. First, although the responses for the first four cost metrics possess varying interquartile ranges, all have a median rating of High importance.* The fifth metric (Costs Previously Incurred to Develop Systems not Useful toward Current Mission) is scored consistently at Low Importance, which is reasonable given that this is a sunk cost metric. In this context, however, it is interesting that the conceptually similar sixth metric (Costs Saved by Reusing Existing Systems for Current Mission), scored at High Importance in the median. It is also notable that the final five metrics, themed around quantifiable science performance and crew productivity and time metrics, exhibited substantial variability of expert opinions. Interquartile ranges for all five of these metrics are greater than one rating level; for example, the range of responses the Mass of Extraterrestrial Material Samples Returned to Earth metric spans the entire available range of Negligible to Paramount.

To analyze the figure of merit results further, Figure 113 summarizes the interquartile ranges and median ratings from Figure 112. In this plot, each figure of merit is represented by a single point, and the figure of merit identification number(s)

---

* One puzzling detail regarding the cost metric ratings is one participant's rating of Integrated Program Lifecycle Cost as Low importance and the other three cost metrics as Moderate or High importance. While this participant left no comments explaining his or her rationale, the implication for this response would seem to be that the sum of development, production, and operations costs is less important than any of these individual costs (colloquially, the whole is less important than any of its parts). However, this view appears to be an outlier in the sense that all other participants rated Integrated Program Lifecycle Cost at least as important as the least-important component cost.

corresponding to each point are listed below and to the right of each. Of particular interest in this figure are metrics with high median ratings and low interquartile ranges, as these characteristics indicate metrics that are rated as high importance consistently among the survey participants. Figure 112 shows that the maximum importance rating (4, or "High") and minimum interquartile range (0.25) coincide for four figures of merit: Integrated Program Lifecycle Cost, Total Spending on Production Activities, Date of First Mission to Leave LEO, and Time Between Missions. Thus, in summary these results support the prioritization of these four figures of merit over others within the 17 metrics considered.

To complete the discussion of the figures of merit section, it should be noted that one of the nine participants left remarks in the "Additional Comments" portion of this section (see the bottom of Figure 105). These remarks are reproduced below, unedited. In some respects, this comment well characterizes the goal of this part of the survey to understand which objective or objectives are most consistently agreed upon as high in priority. As the comment suggests, however, the figure of merit results of this survey predominantly reflect the preferences of NASA space system decision-makers.

> "The hierarchy of priorities is not dictated by any one group, but rather a consolidated set of often disparate stakeholder demands (WH, Congress, Industry, Int'l community). Thus, there are often multiple high priorities despite a constraints in resources and capabalities. As such, we are optimizing and preserving a flexible, open, SOS architecture."

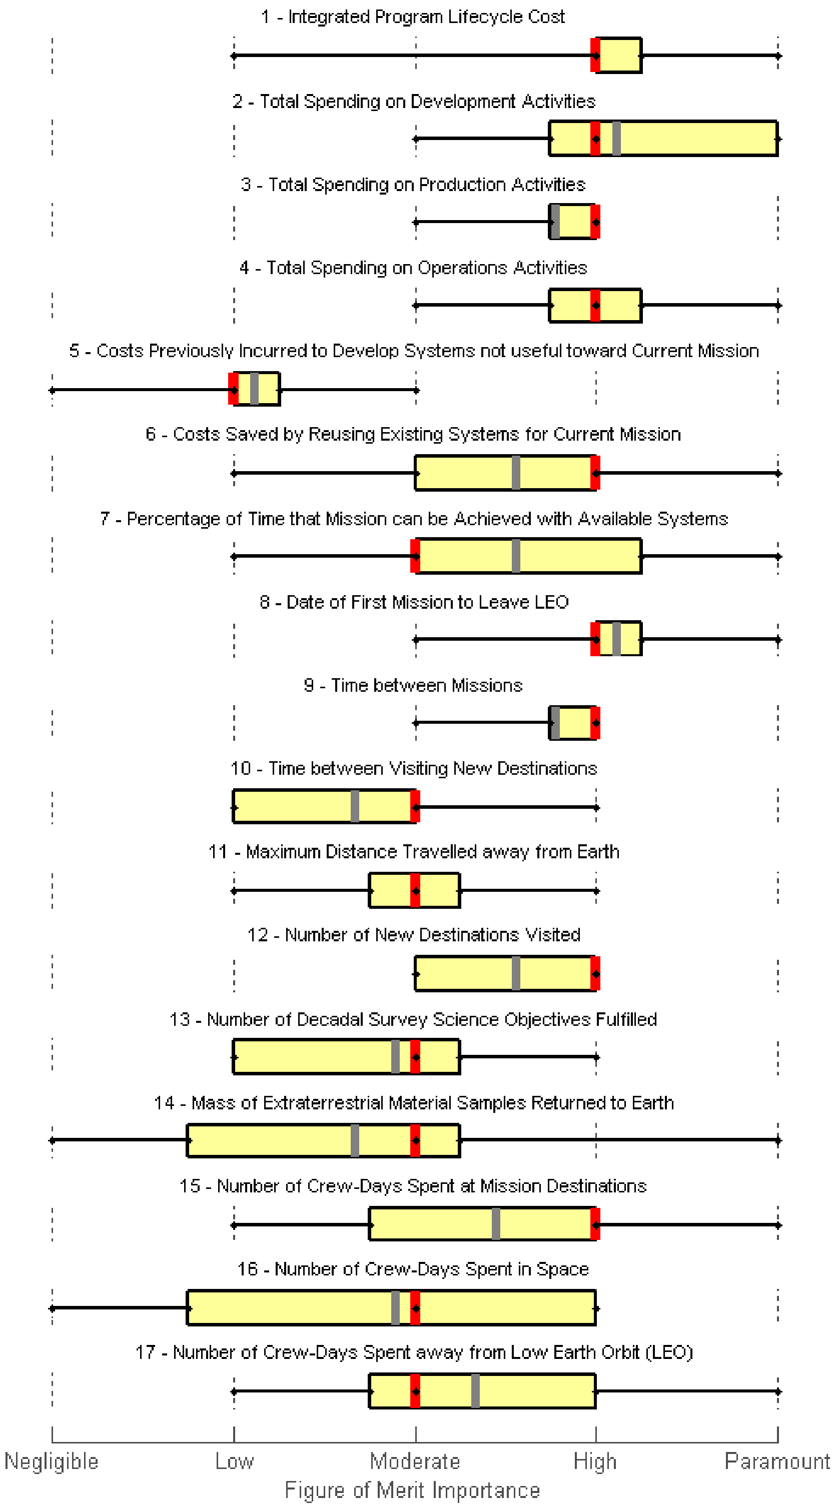

Figure 112. Aggregate Figure of Merit Rating Results. *Yellow boxes represent interquartile ranges and whiskers extend to the minima and maxima. Red lines indicate medians and gray lines indicate means.*

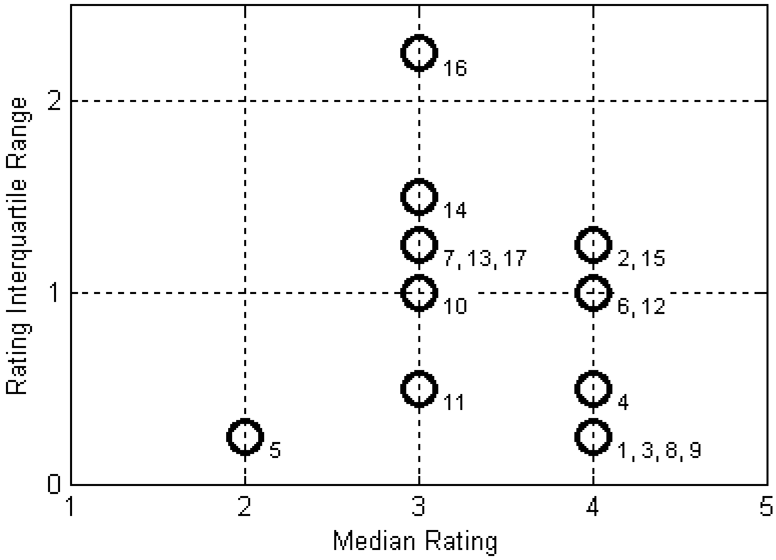

Figure 113. Summary of Figure of Merit Response Interquartile Ranges and Medians. *Corresponding figure of merit identification number(s) are listed to the bottom and to the right of each point.*

C.2.3. Markov Chain Estimates

The second major section of the survey requested that participants provide information regarding the likely evolution of human space exploration mission demands. The section consisted of two parts. In the first part, participants were asked to estimate the amount of time he or she might expect demand for each of twelve mission destinations to last, depending on whether the mission is or is not being achieved at some arbitrary point in the future. The second part is complementary to the first and asked participants to populate two matrices, estimating in each element of each matrix the probability that the next mission demand will be for the destination in the column given that the current mission demand is the destination in the row and given whether or not the current demand is being achieved (one matrix corresponded to each of these binary achievement possibilities). Graphical user interfaces (see Figure 108 and Figure 109) were available to guide the participants through each input. In effect, when a participant had completed these 288 inputs, he or she had populated the transition rates and probability transition matrices of two continuous-time Markov chains describing mission demand evolution, each conditional on whether or not current demand is met.[*]

C.2.3.1. Result Statistics

Aggregated results are shown in Figure 114 and Figure 115. These figures plot the histograms of participant responses to both Parts I and II of this section; the histograms on the diagonal of each figure indicate the expected time responses, and all other histograms indicate the probability responses. The probability responses shown have been normalized for each participant such that the sum of each row of the participant's matrix adds to unity.[†] Thus, the range of all subplot abscissae is zero to unity, except for subplots on the diagonal, which have a range of 0 to 30 years. The color of each subplot indicates the relative amount of variability in the responses, as measured by the interquartile range, with red being high and green being low. The difference between Figure 114 and Figure 115 is that Figure 114 is associated with the condition that current mission demand is fulfilled, whereas Figure 115 is associated with the condition that the current mission demand is not being fulfilled.

---

[*] Given that Markov chains are not covered in the curricula of many engineering degree programs and that the survey participants were not likely to be familiar with them, the term "Markov chain" was not used in the survey.

[†] As Figure 106 indicates, participants were told that each row of probabilities should add to 100%, but that they were not required to spend excessive time attempting to meet this constraint precisely. As a result, normalization was required in post-processing for most rows of participant probability inputs.

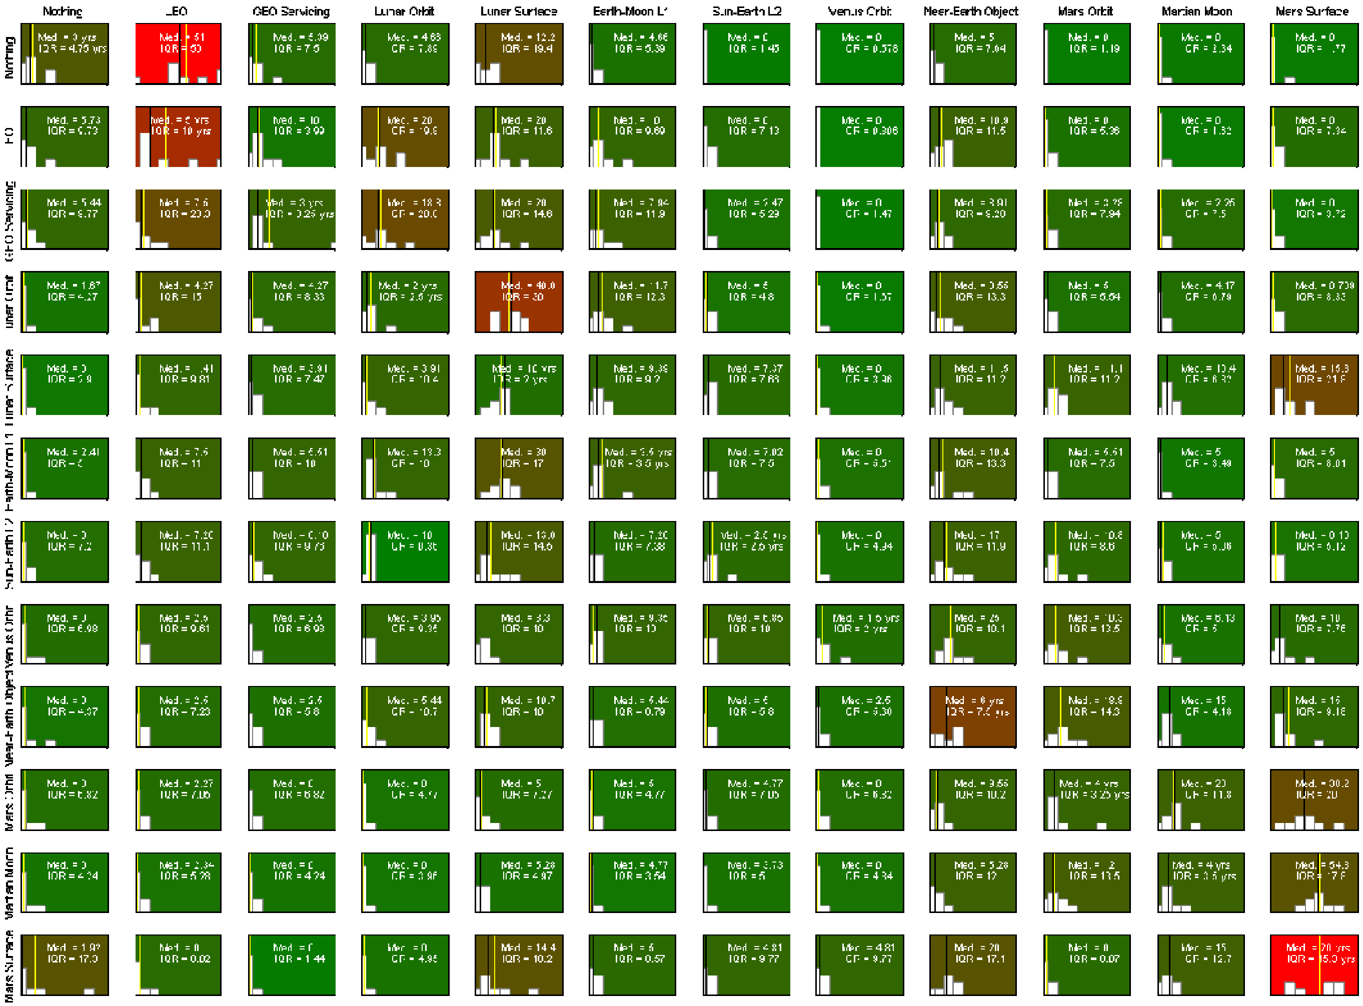

Figure 114. Summary of probability transition matrix and expected time responses, conditioned on current demand being fulfilled. *The range of all subplot abscissae is zero to unity, except for subplots on the diagonal, which have a range of 0 to 30 years. The color of each subplot indicates the relative amount of variability in the responses, as measured by the interquartile range, with red being high and green being low.*

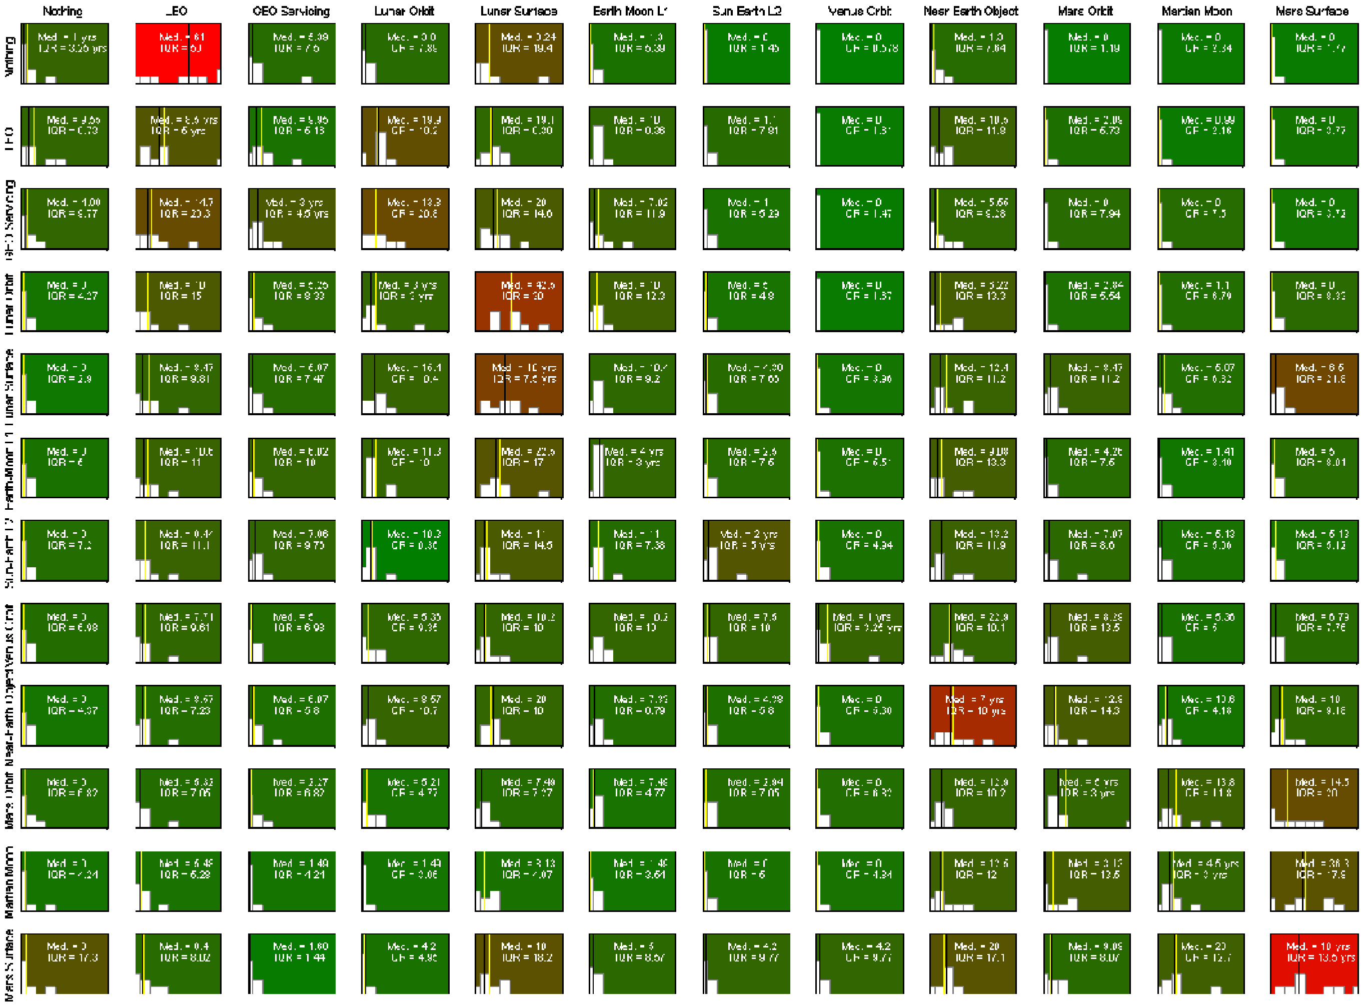

Figure 115. Summary of probability transition matrix and expected time responses, conditioned on current demand not being fulfilled. *The range of all subplot abscissae is zero to unity, except for subplots on the diagonal, which have a 0-30 year range. The color of each subplot indicates the relative amount of variability in the responses, as measured by the interquartile range, with red being high and green being low.*

As the colors of the subplots in Figure 114 and Figure 115 indicate, response agreement is generally high. Figure 116 quantifies this: The histograms on the left in Figure 116 indicate that among all the non-diagonal subplots of Figure 114 and Figure 115, the median interquartile range is 7.7% and the 90th percentile interquartile range is 15.6%. The histograms on the right in Figure 116 indicate that among all the diagonal subplots of Figure 114 and Figure 115, the median interquartile range is 3-4 years and the 90th percentile interquartile range is 11-12 years. Consistent between Figure 114 and Figure 115 is that the elements with the highest variability are the "Nothing" to "LEO" mission destination transition probability and the expected duration of the "Mars Surface" mission demand.

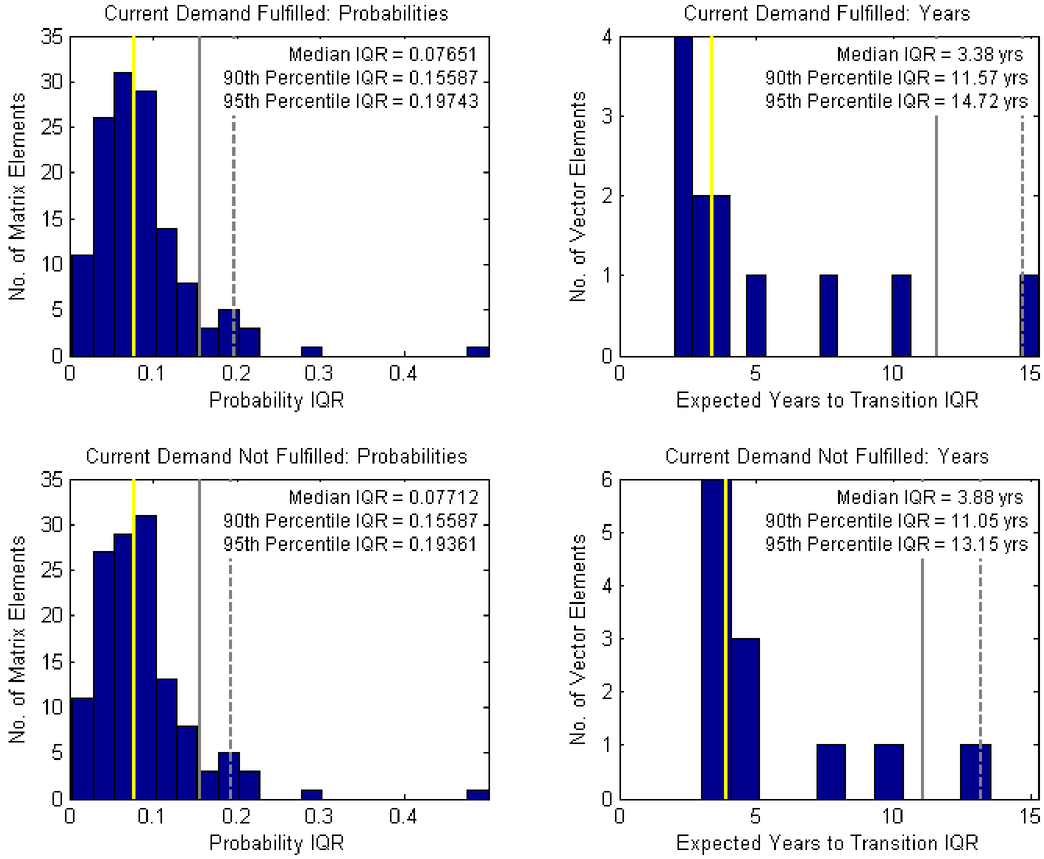

Figure 116. Aggregated Interquartile Range Statistics from Figure 114 and Figure 115. *Yellow lines indicate medians, solid gray lines indicate 90th percentile statistics, and dashed gray lines indicate 95th percentile statistics.*

C.2.3.2. Conversion to a Markov Chain Model

As suggested at the start of this appendix, the primary aim in gathering the data summarized in Figure 114 and Figure 115 is the population of a probability transition matrix for a Markov chain describing the evolution of the demand environment for the NASA human space exploration application. In particular, the Markov chain required by the framework proposed by the current thesis is a discrete-time Markov chain. Thus, two challenges exist given the data in Figure 114 and Figure 115: First, the data represents multiple expert opinions and must be reduced into a single representative model. Second, recalling that the diagonal of Figure 114 and Figure 115 represents expected time responses, the data is in the form of a continuous-time Markov chain rather than a discrete-time Markov chain, and a conversion must be made.

Treating the second challenge first, the conversion of a continuous-time to a discrete-time Markov chain is known as *uniformization* [86] and has a known solution given by Eqs. (C1) and (C2). In Eq. (C1), P_{CTMC} and v_{CTMC} are the probability transition matrix and transition rate vector for the continuous-time Markov chain. In this application, P_{CTMC} and v_{CTMC} are gathered directly from the data provided from the survey; in the case of the rate vector, it is the inverse of the vector of expected times between transitions. The number v in principle can be any rate such that $v \geq v_i \; \forall i$. Since the rates of interest in this application are central tendencies and no mean or median numbers of expected years fell below 1 year (and thus $v_i \geq 1$ yr^{-1} $\forall i$), v for this application is selected as $v = 1$ yr^{-1}. The intermediate matrix P^* is then converted through Eq. (C2) to an equivalent discrete-time Markov chain probability transition matrix P referenced to any desired time step Δt that is longer in duration than $1/v$. Although Eq. (C2) technically requires an infinite sum, in this application acceptable results (specifically, all rows of the resulting transition matrix adding to within 10^{-6} or less of unity) were observed by summing to $k = 10$.

$$P_{ij}^{*} = \begin{cases} 1 - \frac{v_{CTMC,i}}{v}, & j = i \\ \frac{v_{CTMC,i}}{v} P_{CTMC,ij}, & j \neq i \end{cases} \tag{C1}$$

$$P(\Delta t) = \sum_{k=0}^{\infty} P^{*k} e^{-v\Delta t} \frac{(v\Delta t)^{k}}{k!} \tag{C2}$$

With a conversion from a continuous-time to a discrete-time Markov chain now available, the question remains: Which continuous-time Markov chain should be converted? Since many experts contributed to the results of Figure 114 and Figure 115, at least one representative model must be selected to carry forward. In this thesis, results are presented for two different models that represent two different sets of the sample population. The first, on which the primary results of the thesis are based, is based upon a central-tendency model for entire population of expert participants. The second, which is treated in a sensitivity study in Section 6.5, considers a central-tendency model only for the non-dominated experts discussed in Section C.2.1.

C.2.3.2.1. Central Tendency Model for All Experts

In seeking a model to describe the central tendency of the probabilities and expected times to transition in Figure 114 and Figure 115, the two most obvious metrics to consider are median and mean. While there is no objectively correct choice to describe the central tendency of the expert opinions, the median (50th percentile) measure possesses a certain property of fairness that the mean does not. That is, while an exaggerated individual input might highly skew the results of a mean measurement (especially in the case of a small sample), the same is not true for a median measurement. In using the median, each participant is given an equal influence on the determination of the central tendency measure, and for this reason it is selected as the central tendency measurement of choice for this application. For the sake of comparison, however, this section will also show some results that would have been obtained if the mean had been chosen in lieu of the median.

Taking the median values each result represented by a histogram in Figure 114 and Figure 115 and normalizing the probabilities by row forms two versions of P_{CTMC} and v_{CTMC} , one version which corresponds to the case in which current demand is fulfilled and the other which corresponds to the case in which current demand is not fulfilled. Applying Δt = 2 years in Eq. (C2) yields the transition matrices in Table 34 and Table 35.

Table 34. Discrete-time Markov chain probability transition matrix for median expert inputs and Δt = 2 years, for the condition that current mission demand is fulfilled.

		To											
	Mission Demand	Noth.	LEO	GEO Serv.	Lunar Orbit	Lunar Surf.	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Mars Moon	Mars Surf.
From	Nothing	0.5180	0.2447	0.0301	0.0311	0.0928	0.0308	0.0027	0.0001	0.0372	0.0043	0.0037	0.0045
	LEO	0.0192	0.6784	0.0340	0.0577	0.1028	0.0395	0.0039	0.0002	0.0475	0.0060	0.0052	0.0057
	GEO Servicing	0.0261	0.0489	0.5192	0.0776	0.1483	0.0479	0.0157	0.0002	0.0598	0.0246	0.0190	0.0126
	Lunar Orbit	0.0101	0.0326	0.0266	0.3771	0.2868	0.0709	0.0295	0.0003	0.0664	0.0389	0.0338	0.0270
	Lunar Surface	0.0005	0.0046	0.0079	0.0080	0.8261	0.0195	0.0136	0.0002	0.0278	0.0240	0.0231	0.0447
	Earth-Moon L1	0.0095	0.0346	0.0223	0.0435	0.1491	0.5733	0.0259	0.0003	0.0522	0.0278	0.0255	0.0360
	Sun-Earth L2	0.0022	0.0439	0.0325	0.0466	0.1089	0.0448	0.4550	0.0005	0.1057	0.0637	0.0363	0.0598
	Venus Orbit	0.0018	0.0248	0.0201	0.0290	0.0957	0.0690	0.0447	0.2647	0.1950	0.0826	0.0568	0.1157
	Near-Earth Object	0.0006	0.0094	0.0076	0.0138	0.0431	0.0181	0.0141	0.0047	0.7242	0.0540	0.0453	0.0651
	Mars Orbit	0.0005	0.0106	0.0014	0.0024	0.0295	0.0207	0.0171	0.0006	0.0442	0.6123	0.0760	0.1846
	Mars Moon	0.0006	0.0100	0.0012	0.0020	0.0290	0.0187	0.0133	0.0007	0.0273	0.0439	0.6133	0.2400
	Mars Surface	0.0021	0.0011	0.0006	0.0009	0.0213	0.0068	0.0057	0.0039	0.0267	0.0025	0.0185	0.9099

Table 35. Discrete-time Markov chain probability transition matrix for median expert inputs and Δt = 2 years, for the condition that current mission demand is *not* fulfilled.

		To											
	Mission Demand	Noth.	LEO	GEO Serv.	Lunar Orbit	Lunar Surf.	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Mars Moon	Mars Surf.
From	Nothing	0.1417	0.5730	0.0495	0.0515	0.1171	0.0256	0.0029	0.0000	0.0284	0.0050	0.0029	0.0024
	LEO	0.0104	0.8055	0.0209	0.0405	0.0575	0.0234	0.0031	0.0000	0.0270	0.0063	0.0034	0.0019
	GEO Servicing	0.0151	0.1183	0.5203	0.0799	0.1572	0.0471	0.0072	0.0000	0.0436	0.0048	0.0030	0.0035
	Lunar Orbit	0.0009	0.0617	0.0267	0.5249	0.2376	0.0523	0.0197	0.0000	0.0417	0.0189	0.0093	0.0062
	Lunar Surface	0.0003	0.0221	0.0120	0.0291	0.8297	0.0220	0.0074	0.0000	0.0292	0.0180	0.0131	0.0170
	Earth-Moon L1	0.0009	0.0589	0.0269	0.0508	0.1273	0.6136	0.0104	0.0001	0.0501	0.0226	0.0098	0.0287
	Sun-Earth L2	0.0015	0.0813	0.0497	0.0687	0.1126	0.0757	0.3716	0.0001	0.1022	0.0509	0.0380	0.0475
	Venus Orbit	0.0015	0.0879	0.0429	0.0542	0.1325	0.0882	0.0441	0.1356	0.2076	0.0752	0.0521	0.0783
	Near-Earth Object	0.0004	0.0267	0.0145	0.0217	0.0608	0.0199	0.0088	0.0001	0.7579	0.0315	0.0267	0.0310
	Mars Orbit	0.0004	0.0281	0.0103	0.0219	0.0439	0.0305	0.0097	0.0002	0.0584	0.6743	0.0537	0.0688
	Mars Moon	0.0003	0.0297	0.0074	0.0094	0.0459	0.0094	0.0016	0.0005	0.0619	0.0168	0.6470	0.1701
	Mars Surface	0.0002	0.0192	0.0044	0.0093	0.0258	0.0110	0.0065	0.0037	0.0413	0.0179	0.0356	0.8252

In order to visualize the conditional Markov chains in Table 34 and Table 35 as is done for the Markov chains in the examples of Chapters 4-5 of this thesis, it will be helpful to extend them over more than one two-year time increment. Note that the probabilities on the diagonals of these matrices tend to quite high due to this short time step (naturally, as the time step of becomes smaller and smaller, the probability in remaining in a particular state would be expected to approach closer and closer to unity), and thus a visualization of the Markov chain on the two-year step would only reveal the obvious tendency for the system to stay in its current demand state over the coming period. Extending the time increment to an eight-year step for the purposes of visualization yields the diagrams in Figure 73 and Figure 74. In these figures, as in those depicting Markov chains in Chapters 4-5, high-probability transitions are represented as thick dark links and low-probability transitions are represented as thin light links. Also, from each demand state, a green link identifies the highest-probability transition; and if different from the green link, a red link identifies the highest probability transition given departure from a given demand state.

Thus, for example, several differences can be noticed between Figure 73 and Figure 74, which themselves represent the difference in demand evolution experts believed would exist if demand itself were fulfilled (in the case of Figure 73) versus not fulfilled (in Figure 74). Whereas the most likely transition from LEO is to a Lunar Surface demand if LEO demand is fulfilled, it is to remain in LEO if that demand is not fulfilled. Whereas the most likely transition from a Venus Orbit demand is to Mars Surface if demand is fulfilled, it is to the less ambitious Lunar Surface mission if that demand is not fulfilled; and similarly, if Mars Orbit demand is not fulfilled, the most likely demand is to continue Mars Orbit missions rather than progress to Mars Surface missions. It might also be noticed that the red link from the Lunar Surface mission (the second most likely next demand) leads to a Near-Earth Object mission rather than a Mars Surface mission in the event that the Lunar Surface demand is not being met in the current period. These are a few examples that illustrate the general characteristic of the model that the condition of demand being fulfilled favors progression of demand toward missions aimed at more ambitious destinations that are generally farther away from Earth; conversely, the condition of demand not being fulfilled tends to favor a constancy or sometimes regression of demand toward less ambitious destinations closer to Earth.

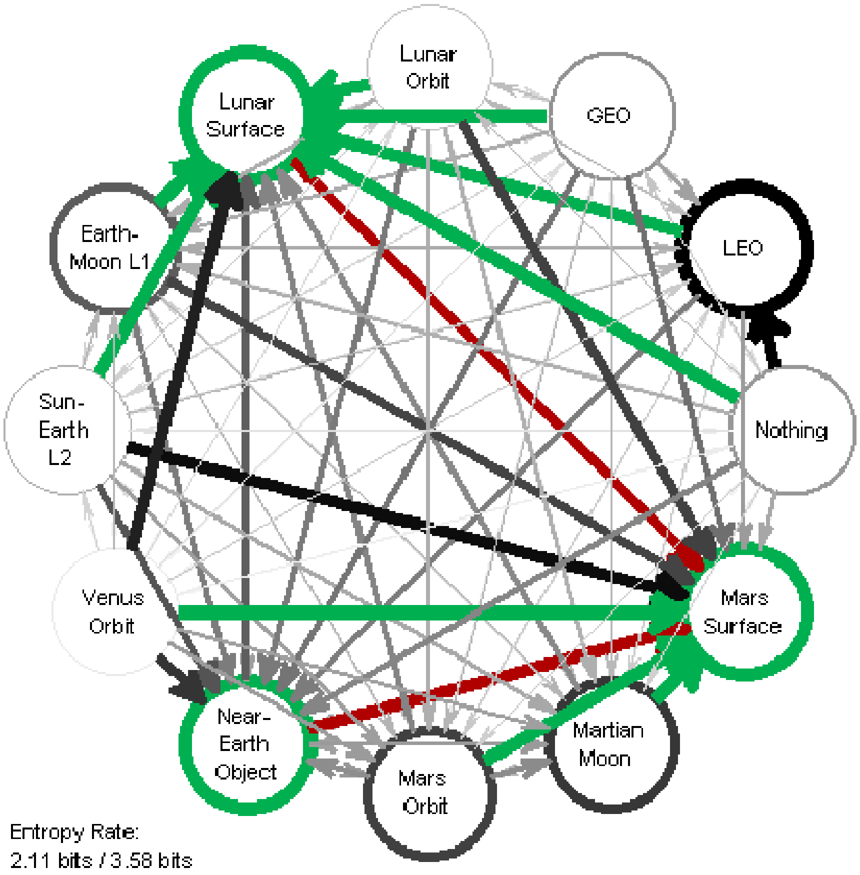

Figure 73. Visualization of the Markov chain of median expert inputs for the condition that current mission demand is fulfilled, with Δt = 8 years. *High-probability transitions are represented as thick dark links and low-probability transitions are represented as thin light links. From each state, a green link identifies the highest-probability transition. If different from the green link, a red link identifies the highest probability transition given departure from that state.*

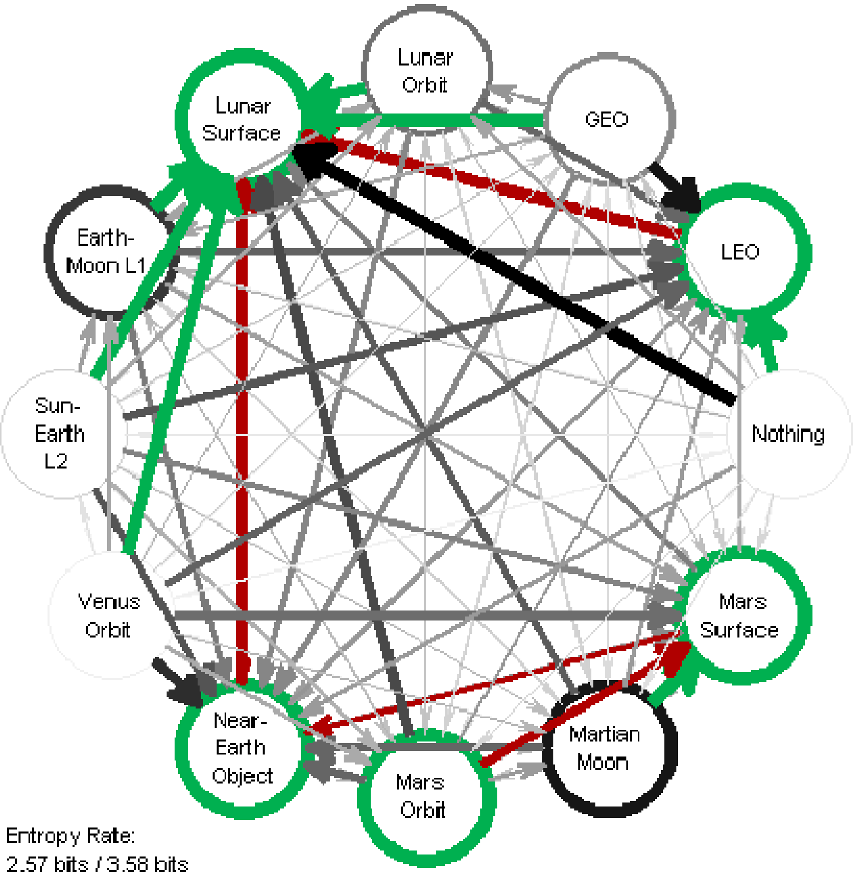

Figure 74. Visualization of the Markov chain of median expert inputs for the condition that current mission demand is *not* fulfilled, with Δt = 8 years. *High-probability transitions are represented as thick dark links and low-probability transitions are represented as thin light links. From each state, a green link identifies the highest-probability transition. If different from the green link, a red link identifies the highest probability transition given departure from that state.*

For comparison, Figure 117 and Figure 118 show the eight-year visualizations of the Markov chains that would have resulted had the mean (instead of the median) been used as the measure of central tendency. While there exist some differences in comparison with Figure 73 and Figure 74, the models share many similarities. In particular, the most-likely and second-most-likely (green and red) transitions in the figures are nearly identical. In the case that demand is fulfilled (i.e., comparing Figure 73 and Figure 117), the main exceptions are the Nothing, LEO, and GEO demand states. In the case of the mean, the most likely next-period demand is for the LEO and GEO states to remain in LEO and GEO, respectively; however, the second-most-likely links from these states match exactly the most likely links from the median case. In the case that demand is not fulfilled (i.e., comparing Figure 74 and Figure 118), the main exceptions are that in the mean, the most likely transition from the Martian Moon mission demand is to remain in the same demand (with transition to Mars Surface demand ranking second, instead of first as in the median case), and that in the mean, the second most likely transition from the Mars Surface mission demand is to the Lunar Surface mission instead of to the Near-Earth Object mission.

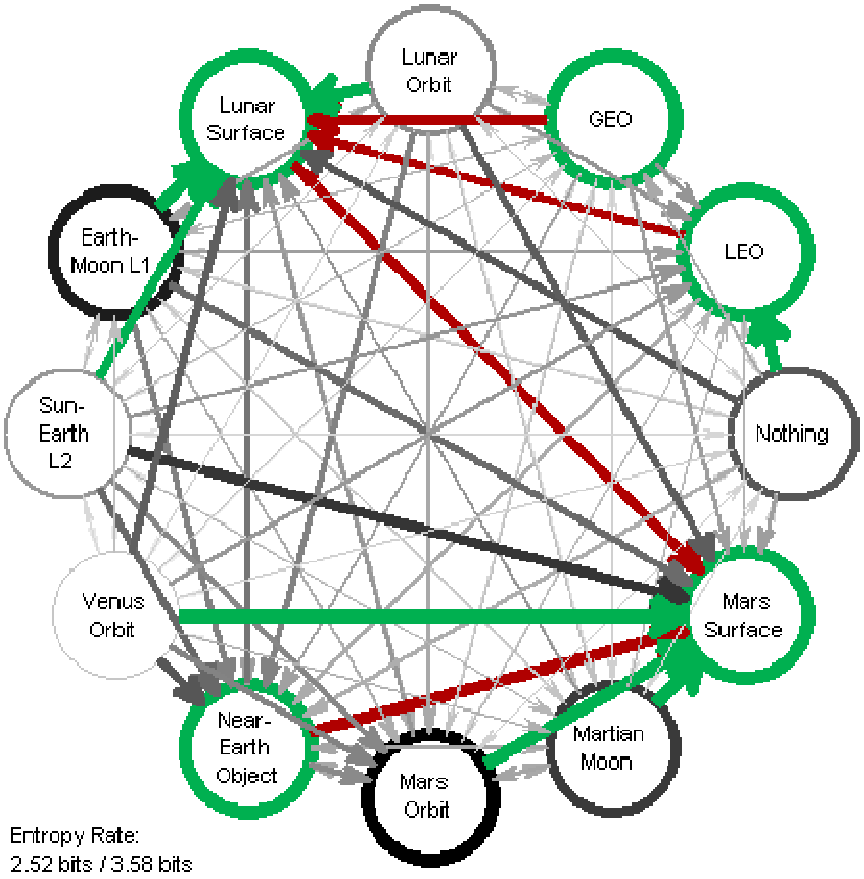

Figure 117. Visualization of the Markov chain of mean expert inputs for the condition that current mission demand is fulfilled, with Δt = 8 years. *High-probability transitions are represented as thick dark links and low-probability transitions are represented as thin light links. From each state, a green link identifies the highest-probability transition. If different from the green link, a red link identifies the highest probability transition given departure from that state.*

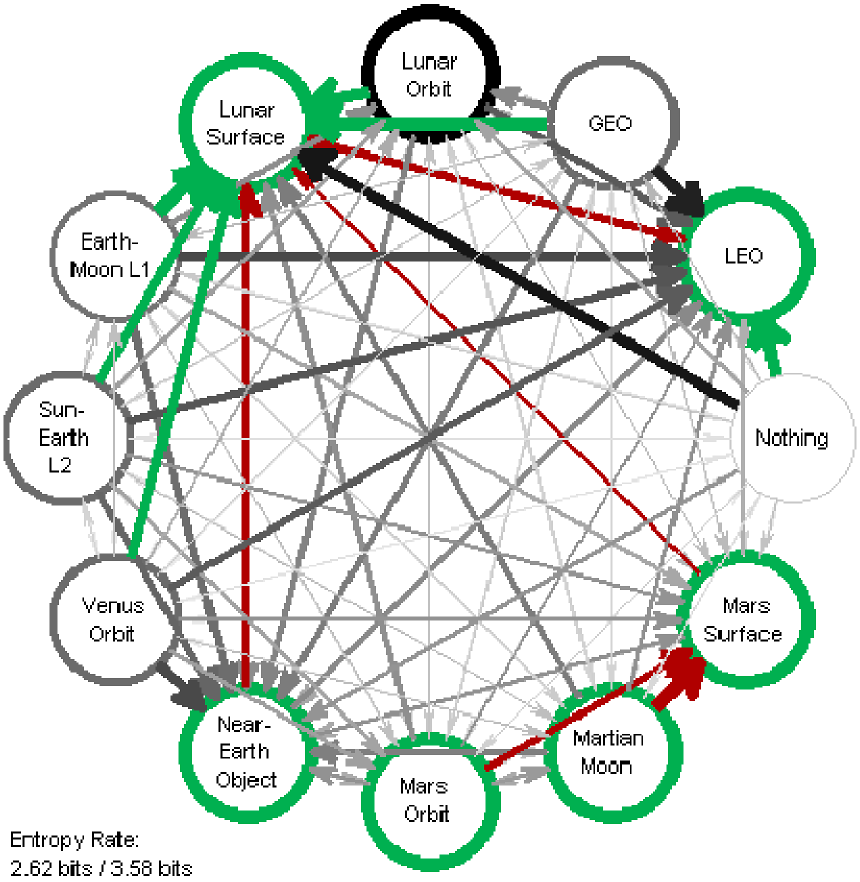

Figure 118. Visualization of the Markov chain of mean expert inputs for the condition that current mission demand is *not* fulfilled, with Δt = 8 years. *High-probability transitions are represented as thick dark links and low-probability transitions are represented as thin light links. From each state, a green link identifies the highest-probability transition. If different from the green link, a red link identifies the highest probability transition given departure from that state.*

As discussed in Section 4.2.2, the stationary distribution of a Markov chain can provide the analyst helpful intuition regarding the direction toward which the demand will eventually tend as a consequence of the probability transition matrix. Toward this end, Figure 119 and Figure 120 provide the stationary distributions for the Markov chains in Table 34 and Table 35, respectively. Each figure compares the result of using the median central tendency measure (in red) to using the mean central tendency measure (in blue); note the general agreement. The most distinctive difference between the two central tendency measures in the long-term stationary sense is that the mean models a somewhat lower probability of running Mars Surface missions and a higher probability of running LEO missions; in this sense it is recognized that the median model is somewhat more optimistic about the demand for Mars Surface missions conditioned on other missions' success. However, conditioned on lack of mission success, Figure 120 shows that the models agree quite well.

The most distinctive difference between the stationary distributions of either central tendency model in Figure 119 and Figure 120 is the much lower Mars Surface mission demand probability in Figure 120. This, a result of the experts' judgements regarding the consequences of not fulfilling mission demand, is accompanied by rises in the probabilities of mission demand for Lunar Surface and LEO missions. In general, these three mission destinations of LEO, Lunar Surface, and Mars Surface, and to a somewhat lesser degree Near-Earth Objects, can be seen to form a set of long-term "sinks" for mission demand in the opinion of the expert participants. Not only do these destinations have long-term demand probabilities significantly higher than others, but Figure 73, Figure 74, Figure 117, and Figure 118 tend to show these destinations as states with consistently high-probabiliy incoming transitions and consistently high probabilities of remaining in their present demand state. In contrast, mission demands like Venus Orbit, Sun-Earth L2, and Nothing tend to act almost as transient states for which demand is rare and, when it does exist, is fleeting.

Before concluding this discussion of stationary probabilities, it should be emphasized that each of the Markov chains in Table 34 and Table 35 is conditioned on mission achievement. Thus, in a sequence of events it is unlikely mission demand will be always fulfilled or never fulfilled, and the true stationary distribution (which could in theory be obtained once a decision policy is defined) will fall between the extremes of Figure 119 and Figure 120.

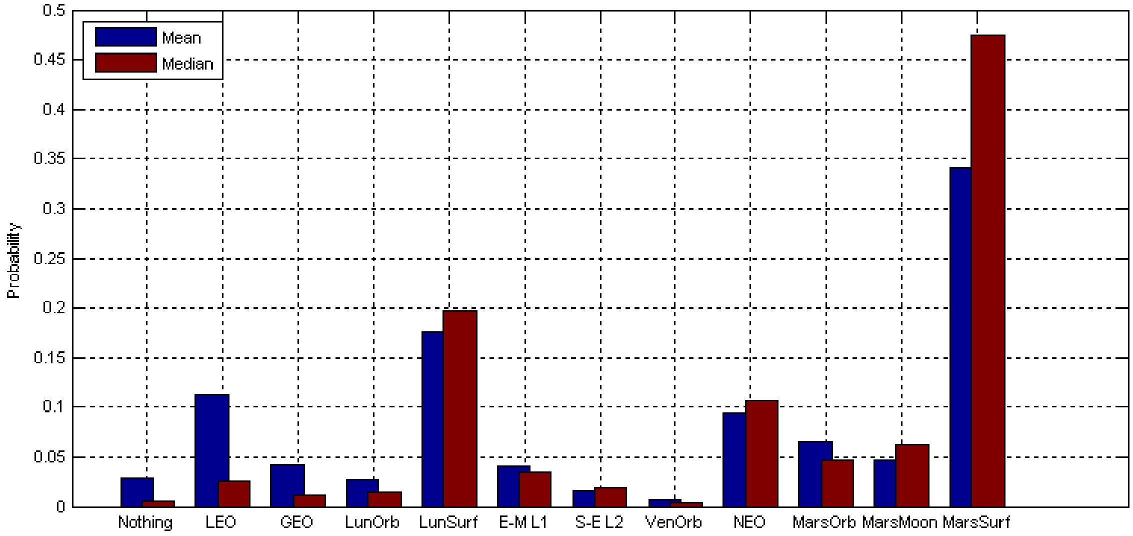

Figure 119. Stationary distribution of Markov chain model (both median and mean versions compared), for the condition that current mission demand is *always* fulfilled.

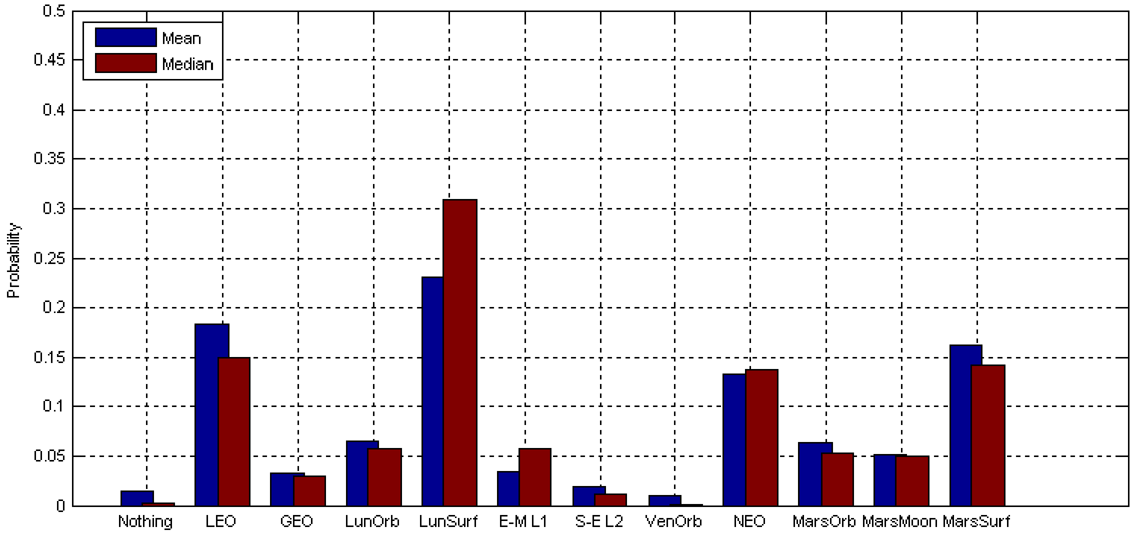

Figure 120. Stationary distribution of Markov chain model (both median and mean versions compared), for the condition that current mission demand is *never* fulfilled.

C.2.3.2.2. Central Tendency Model for Non-Dominated Experts

Also considered in this thesis is a Markov chain demand model derived from a subset of the expert participants of the survey. Described in Section C.2.1 as a set of four experts who are non-dominated based on their number of years of experience in the relevant experience metrics of interest, this subset of results is used to produce two conditional median Markov chains in the same manner as those described in Section C.2.3.2.1. The result for the *Δt* = 2 year time step is shown in Table 50 and Table 51. Visualizations of these chains over *Δt* = 8 year time steps are shown in Figure 121 and Figure 122, and the stationary distributions are shown in Figure 123 and Figure 124. Note that there exist relatively few qualitative differences between this model and that of Section C.2.3.2.1: For example, the relative strengths of the links in the Markov chain diagrams are largely the same (the green and red links are nearly all identical), and the probabilities of the stationary distributions all agree within 5.5%. The most significant difference, which is likely responsible for differences in the results observed when this model is applied to assess sensitivity of the results, is that the non-dominated experts assign a noticeably higher probability of continuing demand for missions to LEO in the event that current mission demand is fulfilled (85.7% in Table 50 vs. 67.8% in Table 34 for the same two-year time increment). Adoption of this model over the general model of Section C.2.3.2.1 will tend to encourage an optimal decision-maker to adopt a policy that, at least in the short term, develops systems oriented more toward this high-likelihood and easy-to-fulfill LEO objective.

Table 50. Discrete-time Markov chain transition matrix for median non-dominated expert inputs and *Δt* = 2 years, for the condition that current mission demand is fulfilled.

		To											
	Mission Demand	Noth.	LEO	GEO Serv.	Lunar Orbit	Lunar Surf.	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Mars Moon	Mars Surf.
From	Nothing	0.3716	0.4635	0.0262	0.0342	0.0495	0.0323	0.0008	0.0000	0.0119	0.0046	0.0026	0.0030
	LEO	0.0087	0.8568	0.0131	0.0280	0.0390	0.0161	0.0006	0.0000	0.0260	0.0067	0.0022	0.0027
	GEO Servicing	0.0307	0.0324	0.5145	0.1039	0.1465	0.0465	0.0025	0.0000	0.0670	0.0233	0.0208	0.0120
	Lunar Orbit	0.0001	0.0143	0.0001	0.5165	0.2248	0.0371	0.0119	0.0002	0.0681	0.0437	0.0407	0.0425
	Lunar Surface	0.0000	0.0065	0.0001	0.0057	0.8232	0.0177	0.0098	0.0003	0.0269	0.0354	0.0247	0.0498
	Earth-Moon L1	0.0001	0.0110	0.0001	0.0409	0.1228	0.6741	0.0091	0.0002	0.0467	0.0414	0.0229	0.0309
	Sun-Earth L2	0.0001	0.0160	0.0001	0.0448	0.1001	0.0174	0.5158	0.0004	0.0945	0.0884	0.0460	0.0764
	Venus Orbit	0.0001	0.0161	0.0001	0.0029	0.0501	0.0291	0.0129	0.5139	0.1321	0.1118	0.0517	0.0793
	Near-Earth Object	0.0000	0.0058	0.0000	0.0127	0.0249	0.0151	0.0048	0.0002	0.8224	0.0532	0.0261	0.0347
	Mars Orbit	0.0000	0.0083	0.0001	0.0013	0.0264	0.0157	0.0073	0.0008	0.0411	0.7202	0.0449	0.1340
	Mars Moon	0.0000	0.0081	0.0001	0.0010	0.0188	0.0155	0.0074	0.0011	0.0121	0.0298	0.7200	0.1859
	Mars Surface	0.0000	0.0005	0.0000	0.0009	0.0139	0.0113	0.0096	0.0092	0.0252	0.0031	0.0175	0.9087

Table 51. Discrete-time Markov chain transition matrix for median non-dominated expert inputs and Δt = 2 years, for the condition that current mission demand is *not* fulfilled.

		To											
	Mission Demand	Noth.	LEO	GEO Serv.	Lunar Orbit	Lunar Surf.	Earth-Moon L1	Sun-Earth L2	Venus Orbit	Near-Earth Object	Mars Orbit	Mars Moon	Mars Surf.
From	Nothing	0.1398	0.6512	0.0459	0.0488	0.0769	0.0116	0.0009	0.0000	0.0182	0.0044	0.0015	0.0009
	LEO	0.0064	0.8306	0.0235	0.0310	0.0531	0.0179	0.0006	0.0000	0.0290	0.0057	0.0012	0.0011
	GEO Servicing	0.0158	0.1262	0.6106	0.0612	0.1119	0.0252	0.0012	0.0000	0.0407	0.0037	0.0021	0.0015
	Lunar Orbit	0.0006	0.0696	0.0151	0.5221	0.2104	0.0453	0.0134	0.0000	0.0672	0.0325	0.0175	0.0062
	Lunar Surface	0.0002	0.0286	0.0046	0.0235	0.8592	0.0139	0.0041	0.0000	0.0214	0.0184	0.0139	0.0122
	Earth-Moon L1	0.0004	0.0579	0.0110	0.0725	0.1121	0.6752	0.0014	0.0000	0.0411	0.0133	0.0027	0.0124
	Sun-Earth L2	0.0005	0.0642	0.0118	0.0363	0.0776	0.0491	0.5659	0.0002	0.0685	0.0494	0.0347	0.0418
	Venus Orbit	0.0007	0.0828	0.0175	0.0240	0.1084	0.0688	0.0295	0.3681	0.1307	0.0633	0.0433	0.0630
	Near-Earth Object	0.0002	0.0270	0.0058	0.0166	0.0525	0.0226	0.0098	0.0001	0.8052	0.0228	0.0172	0.0204
	Mars Orbit	0.0002	0.0329	0.0010	0.0096	0.0337	0.0178	0.0010	0.0003	0.0674	0.6987	0.0514	0.0859
	Mars Moon	0.0002	0.0346	0.0009	0.0105	0.0275	0.0200	0.0010	0.0005	0.0345	0.0525	0.6748	0.1430
	Mars Surface	0.0002	0.0423	0.0011	0.0090	0.0374	0.0164	0.0071	0.0054	0.0342	0.0026	0.0224	0.8221

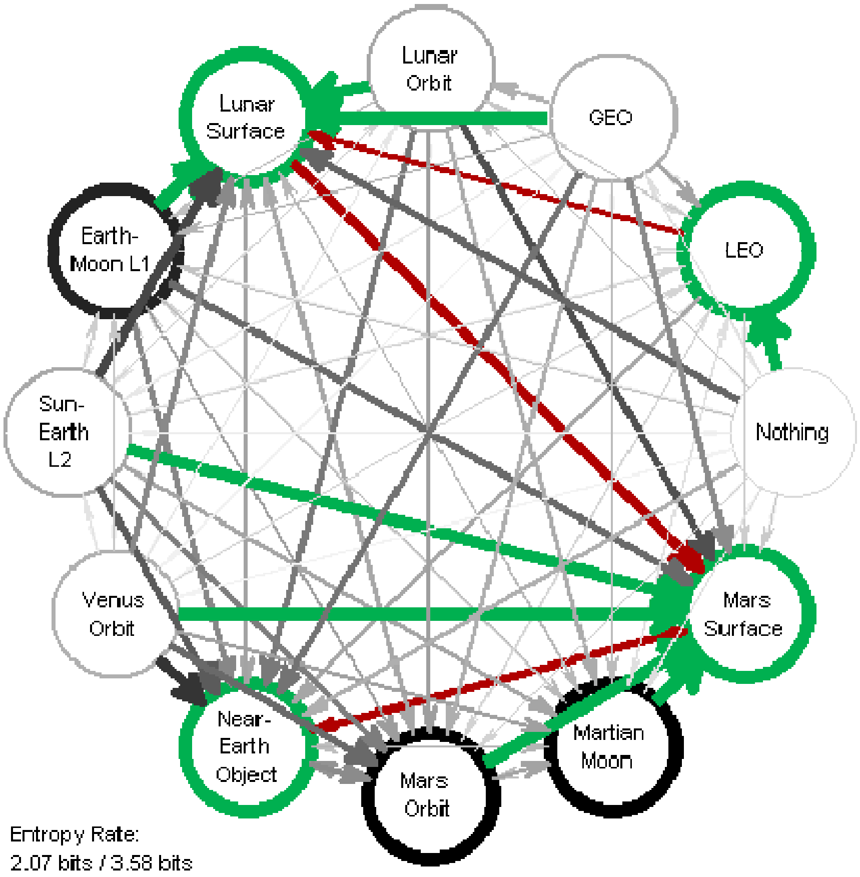

Figure 121. Visualization of the Markov chain of median non-dominated expert inputs for the condition that current mission demand is fulfilled, with Δt = 8 years. *High-probability transitions are represented as thick dark links and low-probability transitions are representedas thin light links. From each state, a green link identifies the highest-probability transition. If different from the green link, a red link identifies the highest probability transition given departure from that state.*

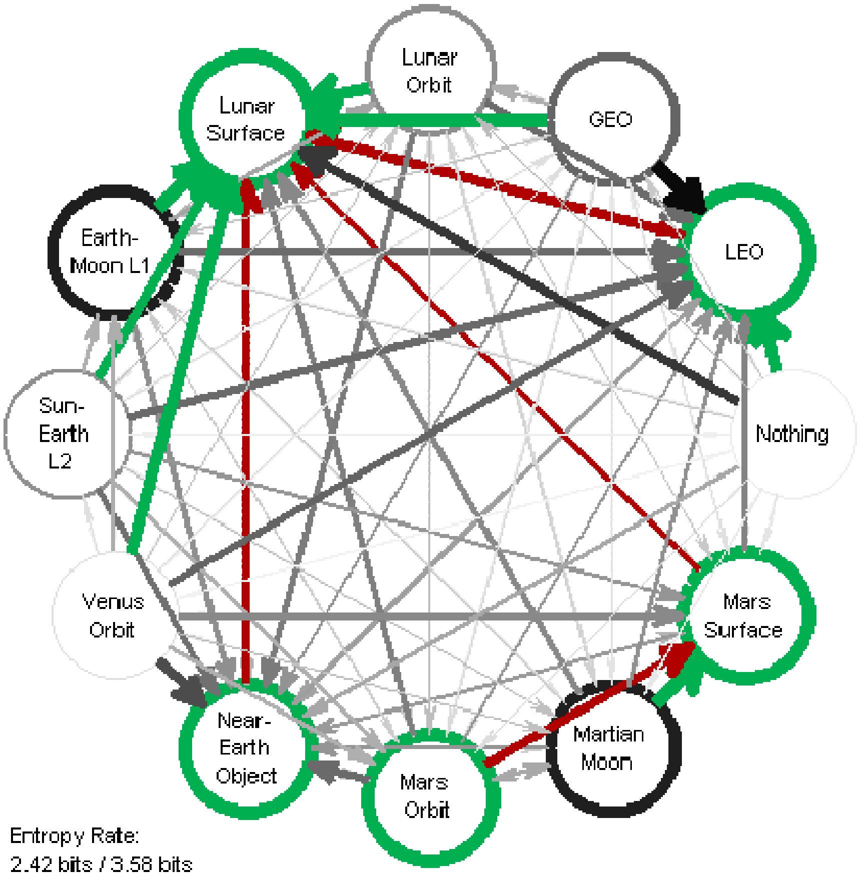

Figure 122. Visualization of the Markov chain of median non-dominated expert inputs for the condition that current mission demand is *not* fulfilled, with Δt = 8 years. *High-probability transitions are represented as thick dark links and low-probability transitions are representedas thin light links. From each state, a green link identifies the highest-probability transition. If different from the green link, a red link identifies the highest probability transition given departure from that state.*

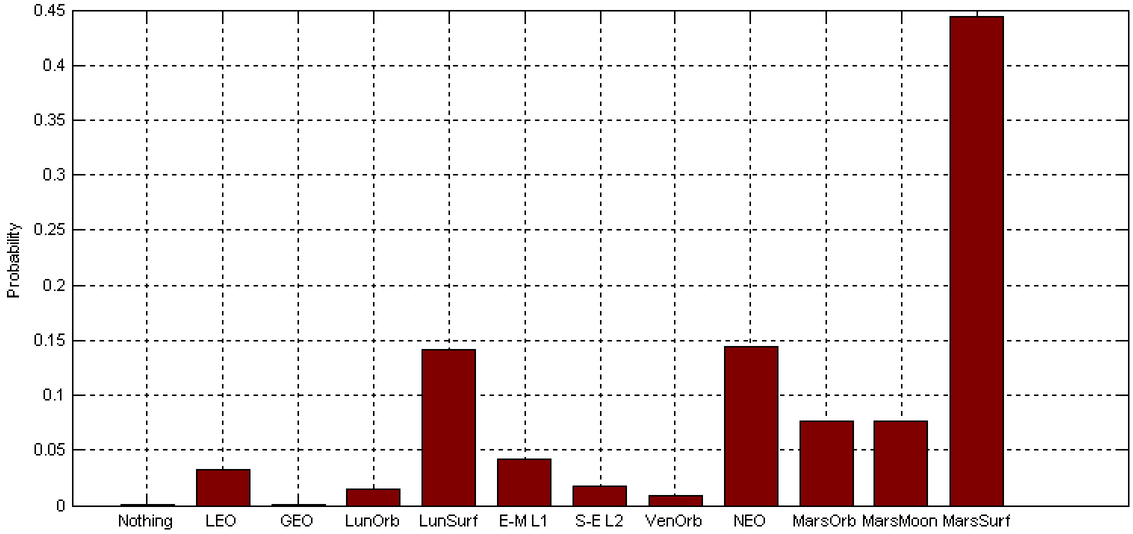

Figure 123. Stationary distribution of median Markov chain model for non-dominated expert inputs and the condition that current mission demand is *always* fulfilled.

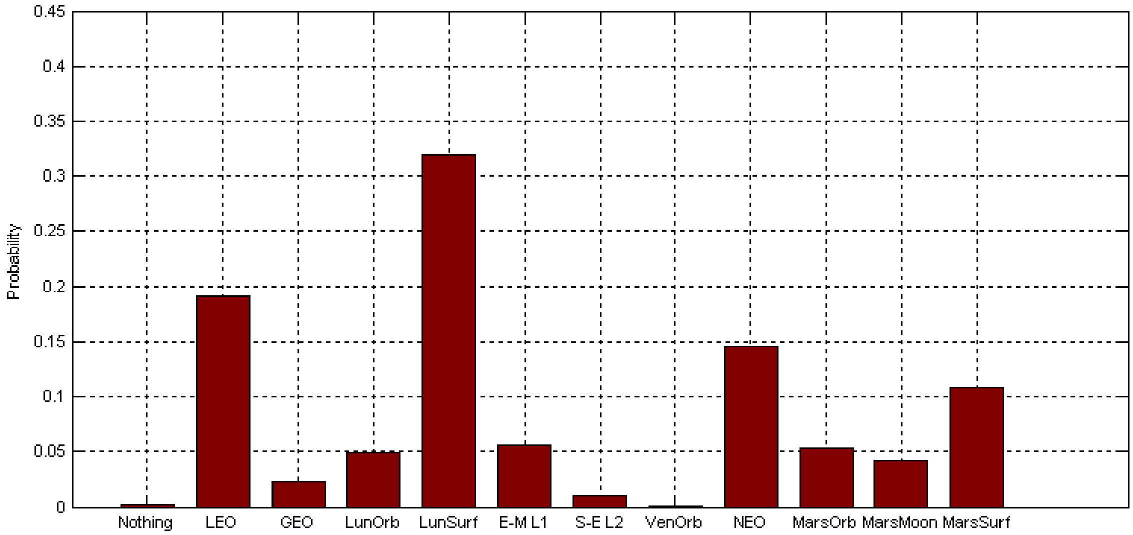

Figure 124. Stationary distribution of median Markov chain model for non-dominated expert inputs and the condition that current mission demand is *never* fulfilled.

C.2.3.3. Participant Comments

To complete the discussion of the Markov chain estimates section, it should be noted that three of the nine participants left remarks in the "Additional Comments" portion of this section (see the bottom of Figure 106). These remarks are reproduced below, unedited. Overall, these comments convey opinions on various topics. All three contain reflections on the survey process itself, with two of the comments explicitly suggesting the approach of eliciting probabilities was overcomplicated and one explicitly suggesting the approach was overly simplistic. One of the participants who felt the survey was overcomplicated suggests that it would have been less complicated if more extensive information were added for the participants to consider. Another topic covered is what the participant believes is a misplaced emphasis within the agency, government, or society on achieving "firsts". The third comment also suggests considering alternate tools for data collection, the fact that the results of the survey will be variable, and that demand evolves as a function of multiple time-varying factors.

> "I think you are making this way more complicated than needed. Be careful to not over think it…………you can get any answer you want by doing that."

> "This seems to be a very overcomplicated way to predict the interest of various destinations. l lost interest quickly in answering your questions because there isnt enough information provided to answer in a way that i feel comforatable will return data worth basing any decisions on. For exploration, I believe that much will depend on what we plan to do at each destination. If we go to plant a flag/say we have been there, very little interest will be created. We need to consider how we can deliver benefits at each destination we visit and be sure we can deliver it, i.e design systems accordingly and plan stay times accordingly. I believe this will create interest in teh next destination. If we think exploration is about "firsts" and we must keep delivering "firsts" then we are not delivering enough real benefits to justify the expense. I think the interest in delivering "firsts" is seriously misplaced."

> "The concept of demand and the pairwise comparisions are probably overly simplistic and all comparisons are likely not suitable. I would have chosen a different tool or setup. The results may not be clear due to variability in interpretation. The demand is very depenedenty upon multiple factors that are always dynamic."

C.3. Summary

This appendix has documented the approach and results associated with an extensive survey distributed to 21 personnel with substantial experience in the field of human space exploration in June 2011. The aim of the survey was to gather information to support (1) the identification of meaningful quantitative figures of merit (or surrogate figures of merit) and (2) the specification of the probability transition matrix of the underlying Markov chain for demand environment evolution for the decision problem in Chapter 6 of this thesis.

The first aim of the survey yielded an objective tie among four figures of merit, which earned the highest median score and lowest score interquartile range: Integrated Program Lifecycle Cost, Total Spending on Production Activities, Date of First Mission to Leave LEO, and Time Between Missions. In short, these results support the prioritization of these four figures of merit over others within the 17 metrics considered.

The second aim of the survey yielded a primary model for the Markovian evolution of mission destination demand for human space exploration missions, converted from the continuous-time Markov chain input of the expert participants and aggregated (with acceptable agreement in advance) into a central tendency via the median statistic to permit each participant an equal influence on the results. A secondary model for sensitivity studies was generated using a non-dominated subset of experts based on years of experience in different categories of interest. Both models suggested (as expected) that the expert participants felt that fulfillment of current mission demands tended to result in the progression of next-period demand toward more ambitious destinations away from Earth, while failure to fulfill current demand would tend to result in either constancy or regression of demand toward less ambitious destinations closer to Earth. Both models also suggested that certain destinations, such as LEO, Lunar Surface, Mars Surface, and Near-Earth Objects, are ultimately highly likely to be demanded of NASA, while other destinations, such as Sun-Earth L2, are unlikely to be demanded or would be transient if demanded. These observations and others throughout this appendix agree with many expectations that an engineer in the industry might have for the demands placed on NASA and help to provide some additional validity to the model that has been developed here to enable a stochastic description of human space exploration mission demand.

REFERENCES

[1] *Aeronautics and Space Report of the President: Fiscal Year 2007 Activities,* NASA NP-2008-05-513-HQ, Appendix D-1A, Available: http://history.nasa.gov/NP-2008-05-513-HQ.pdf [9 Sept. 2010].

[2] Logsdon, J.M., "The decision to develop the Space Shuttle," *Space Policy,* Vol. 2, No. 2, May 1986, pp. 103-119.

[3] Ertel, I.D., Newkirk, R.W., and Brooks, C.G., *The Apollo Spacecraft: A Chronology,* Vol. 4, NASA SP-4009, NASA, Washington, 1978, pp. 325-326.

[4] "Apollo Missions Extended to '74," *The New York Times*, 5 Jan. 1970, p. 10.

[5] *Aeronautics and Astronautics,* NASA SP-4015, NASA, Washington, 1972, pp. 284-285.

[6] Compton, W.D. and Benson, C.D., *Living and Working in Space: A History of Skylab,* NASA SP-4208, NASA, Washington, 1983, chap. 1.

[7] Shayler, D.J., *Skylab: America's Space Station,* Springer, London, 2001, p. 19.

[8] *Skylab News Reference,* NASA Office of Public Affairs, March 1973, pp. iii-3-iii-5, I-1.

[9] Belew, L.F., Ed., *Skylab, Our First Space Station,* NASA SP-400, NASA, Washington, 1977, chap. 2.

[10] White, M. and Bongst, O., Eds., *JSC Digital Image Collection* [online database], 13 May 2008, Available: http://images.jsc.nasa.gov/index.html [22 Sept. 2010].

[11] *Manned Lunar Landing Program Mode Comparison,* Office of Manned Space Flight, NASA TM-X-74929, 30 July1962.

[12] Brown, O.C. and Eremenko, P., "Acquiring Space Systems in an Uncertain Future: The Introduction of Value-Centric Acquisition," *United States Air Force Space Command High Frontier: The Journal for Space, Cyberspace, and Missile Professionals,* Vol. 6, No. 1, Nov. 2009, pp. 37-43.

[13] Mish, F.C., Ed., “Flexible,” *Merriam-Webster Collegiate Dictionary,* 11th ed., Merriam-Webster, Springfield, 2003, pp. 478-479.

[14] Saleh, J.H., Hastings, D.E., and Newman, D.J., “Flexibility in system design and implications for aerospace systems,” *Acta Astronautica,* Vol. 53, No. 12, Dec. 2003, pp. 927-944.

[15] Chen, W. and Lewis, K., “Robust Design Approach for Achieving Flexibility in Multidisciplinary Design,” *AIAA Journal,* Vol. 37, No. 8, Aug. 1999, pp. 982-989.

[16] Porter, A.L. and Read, W.H., Eds., *The Information Revolution: Current and Future Consequences,* Ablex, Greenwich, 1998, chap. 3.

[17] Kirby, M.R., “A Methodology for Technology Identification, Evaluation, and Selection in Conceptual and Preliminary Aircraft Design,” Ph.D. Thesis, Georgia Institute of Technology, March 2001.

[18] Mavris, D.N., DeLaurentis, D.A., Bandte, O., and Hale, M.A., “A Stochastic Approach to Multi-disciplinary Aircraft Analysis and Design,” AIAA 1998-0912, 36th Aerospace Sciences Meeting and Exhibit, Reno, 12-15 Jan. 1998.

[19] Haberfellner, R. and de Weck, O.L., “Agile SYSTEMS ENGINEERING versus AGILE SYSTEMS engineering,” INCOSE 2005 Annual International Symposium, Rochester, 10-15 July 2005.

[20] *Quality Engineering using Design of Experiments: A Supplier Symposium on Taguchi Methods Sponsored by the Ford Motor Company,* American Supplier Institute, Romulus, 1984.

[21] Taguchi, G., *Introduction to Quality Engineering,* Asian Productivity Organization, Tokyo, 1986.

[22] Taguchi, G., *System of Experimental Design,* Quality Resources, White Plains, 1987, Vols. 1-2.

[23] Taguchi, G. and Clausing, D., “Robust Quality,” *Harvard Business Review,* Vol. 68, No. 1, Jan. 1990, pp. 65-75.

[24] Park, G., Lee, T., Lee, K.H., and Hwang, K., "Robust Design: An Overview," *AIAA Journal*, Vol. 44, No. 1, Jan. 2006, pp. 181-191.

[25] Park, S.H., *Robust Design and Analysis for Quality Engineering*, Chapman & Hall, London, 1996.

[26] Launius, R.D., "Assessing the legacy of the Space Shuttle," *Space Policy*, Vol. 22, No. 4, Nov. 2006, pp. 226-234.

[27] Lafleur, J.M. and Saleh, J.H., "Survey of mission evolution and flexibility in the Space Shuttle program," *Space Policy,* Vol. 26, No. 4, Nov. 2010, pp. 236-245.

[28] Lafleur, J. and Saleh, J., "Survey of Intra- and Inter-Mission Flexibility in Space Exploration Systems," *Acta Astronautica,* Vol. 67, No. 1, July 2010, pp. 97-107.

[29] Stanley, D., Cook, S., Connolly, J., et. al., *NASA's Exploration Systems Architecture Study: Final Report*, NASA-TM-2005-214062, Nov. 2005, preface, chaps. 4, 6, 14.

[30] Maciuca, D., Chow, J., Siddiqi, A., de Weck, O., Alban, S., Dewell, L., et. al., "A Modular, High-Fidelity Tool to Model the Utility of Fractionated Space Systems," AIAA 2009-6765, Space 2009 Conference and Exposition, Pasadena, 14-17 Sept. 2009.

[31] Choi, S., Westley, D., Eichenberg-Bicknell, E., and Wisniewski, M., "Using a Value-Centric Tool to Optimize Lifecycle Cost, Value and Risk of Spacecraft Architectures," AIAA 2009-6766, Space 2009 Conference and Exposition, Pasadena, 14-17 Sept. 2009.

[32] McCormick, D., Barrett, B., and Burnside-Clapp, M., "Analyzing Fractionated Satellite Architectures Using RAFTIMATE: A Boeing Tool for Value-Centric Design," AIAA 2009-6767, Space 2009 Conference and Exposition, Pasadena, 14-17 Sept. 2009.

[33] O'Neill, M.G., Yue, H., Nag, S., Grogan, P., and de Weck, O.L., "Comparing and Optimizing the DARPA System F6 Program Value-Centric Design Methodologies," AIAA 2010-8828, Space 2010 Conference and Exposition, Anaheim, 30 Aug. – 2 Sept. 2010.

[34] Augustine, N.R., Austin, W.M., Chyba, C.C., Kennel, C.F., Bejmuk, B.I., Crawley, E.F., et. al., *Seeking a Human Spaceflight Program Worthy of a Great Nation,* Review of U.S. Human Spaceflight Plans Committee, Oct. 2009, Exec. Summary, Chap. 3.

[35] Knight, F.H., "Cost of Production and Price over Long and Short Periods," *The Journal of Political Economy,* Vol. 29, No. 4, April 1921, pp. 304-335.

[36] Hart, A.G., "Anticipations, Business Planning, and the Cycle," The *Quarterly Journal of Economics,* Vol. 51, No. 2, Feb. 1937, pp. 237-297.

[37] Stigler, G., "Production and Distribution in the Short Run," *The Journal of Political Economy,* Vol. 47, No. 3, June 1939, pp. 305-327.

[38] Koopmans, T.C., "On flexibility of future preference," *Human Judgments and Optimality,* Wiley, New York, 1964, pp. 243-254.

[39] Raiffa, H., *Decision Analysis: Introductory Lectures on Choices under Uncertainty,* Addison-Wesley, Reading, 1968.

[40] Howard, R.A., "Decision Analysis: Applied Decision Theory," *Proceedings of the Fourth International Conference on Operational Research,* Wiley-Interscience, New York, 1966, pp. 55-71.

[41] Edwards, W., Miles, R.F. Jr., Von Winterfeldt, D., *Advances in Decision Analysis: From Foundations to Applications,* Cambridge University Press, Cambridge, 2007.

[42] Schlaifer, R., *Probability and Statistics for Business Decisions*, McGraw-Hill, New York, 1959.

[43] Raiffa, H. and Schlaifer, R., *Applied Statistical Decision Theory,* Colonial Press, Clinton, 1961.

[44] Pratt, J.W., Raiffa, H., and Schlaifer, R., *Introduction to Statistical Decision Theory,* McGraw-Hill, New York, 1965.

[45] Daniel, W.W., *Decision Trees for Management Decision Making: An Annotated Bibliography,* Public Administration Series Bibliography P-254, Vance Bibliographies, Monticello, 1979.

[46] Kirkwood, C.W., *Decision Tree Primer,* Arizona State University, Tempe, 2002, Available: http://www.public.asu.edu/~kirkwood/DAStuff/decisiontrees/index.html [2 Oct. 2010].

[47] Saleh, J.H., Lamassoure, E.S., Hastings, D.E., and Newman, D.J., "Flexibility and the Value of On-Orbit Servicing: New Customer-Centric Perspective," *Journal of Spacecraft and Rockets*, Vol. 40, No. 2, March 2003, pp. 279-291.

[48] Saleh, J.H., "Weaving Time into System Architecture: New Perspectives on Flexibility, Spacecraft Design Lifetime, and On-Orbit Servicing," Ph.D. Thesis, Massachusetts Institute of Technology, Feb. 2002.

[49] Saleh, J.H., Mark, G., and Jordan, N.C., "Flexibility: A Multi-Disciplinary Literature Review and a Research Agenda for Designing Flexible Engineering Systems," *Journal of Engineering Design,* Vol. 20, No. 3, June 2009, pp. 307-323.

[50] Bodie, Z., Merton, R.C., and Samuelson, W.F., "Labor supply flexibility and portfolio choice in a life cycle model," *Journal of Economic Dynamics and Control,* Vol. 16, No. 3, July 1992, pp. 427-449.

[51] Jones, R.A. and Ostroy, J.M., "Flexibility and Uncertainty," *Review of Economic Studies,* Vol. 51, No. 1, Jan. 1984, pp. 13-32.

[52] Christian, J.A. III, "A Quantitative Approach to Assessing System Evolvability," NASA Johnson Space Center, 2004.

[53] Christian, J.A. III and Olds, J.R., "A Quantitative Methodology for Identifying Evolvable Space Systems," AIAA 2005-2543, 1st Space Exploration Conference: Continuing the Voyage of Discovery, Orlando, 30 Jan. – 1 Feb. 2005.

[54] Gupta, S.K. and Rosenhead, J., "Robustness in Sequential Investment Decisions," *Management Science,* Vol. 15, No. 2, Oct. 1968, pp. B18-B29.

[55] Baykasoğlu, A., "Quantifying machine flexibility," *International Journal of Production Research,* Vol. 47, No. 15, Aug. 2009, pp. 4109-4123.

[56] Silver, M.R. and de Weck, O.L., "Time-Expanded Decision Network: A New Framework for Designing Evolvable Complex Systems," AIAA 2006-6964, 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Portsmouth, 6-8 Sept. 2006.

[57] Silver, M.R. and de Weck, O.L., "Time-Expanded Decision Networks: A Framework for Designing Evolvable Complex Systems," *Systems Engineering*, Vol. 10, No. 2, April 2007, pp. 167-186.

[58] Mandelbaum, M. and Buzacott, J., "Flexibility and decision making," *European Journal of Operational Research*, Vol. 44, No. 1, Jan. 1990, pp. 17-27.

[59] Tetzlaff, U.A.W., *Optimal Design of Flexible Manufacturing Systems*, Physica-Verlag, Heidelberg, 1990, chaps. 2, 8, 9, Appendix C.

[60] Van Looveren, A.J., Gelders, L.F., and Van Wassenhove, L.N., "A Review of FMS Planning Models," *Modelling and Design of Flexible Manufacturing Systems*, Elsevier, Amsterdam, 1986, chap. 1.

[61] Tempelmeier, H. and Kuhn, H., *Flexible Manufacturing Systems*, Wiley, New York, 1993, chaps. 1, 4.

[62] Mohamed, Z.M., *Flexible Manufacturing Systems: Planning Issues and Solutions*, Garland Publishing, New York, 1994, chap. 2.

[63] Gupta, Y.P. and Goyal, S., "Flexibility of manufacturing systems: Concepts and measurements," *European Journal of Operational Research*, Vol. 43, No. 2, Nov. 1989, pp. 119-135.

[64] Raouf, A. and Anjum, M.F., "Manufacturing Systems: Flexibility Assessment," *Flexible Manufacturing Systems: Recent Developments*, Elsevier, Amsterdam, 1995, pp. 69-84.

[65] Muramatsu, R., Ishi, K., and Takahashi, K., "Flexibility in Pull and Push Type Production Ordering Systems – Some Ways to Increase Flexibility in Manufacturing Systems," *Flexible Manufacturing Systems: Recent Developments*, Elsevier, Amsterdam, 1995, pp. 95-109.

[66] Cox, T. Jr., "Toward the Measurement of Manufacturing Flexibility," *Production and Inventory Management Journal*, Vol. 30, No. 1, 1989, pp. 68-72.

[67] Viscito, L. and Ross, A.M., "Quantifying Flexibility in Tradespace Exploration: Value Weighted Filtered Outdegree," AIAA 2009-6561, Space 2009 Conference and Exposition, Pasadena, 14-17 Sept. 2009.

[68] Ross, A.M. and Rhodes, D.H., "Using Natural Value-Centric Time Scales for Conceptualizing System Timelines through Epoch-Era Analysis," INCOSE International Symposium 2008, Utrecht, June 2008.

[69] Olthoff, C.T., Cunio, P.M., Hoffman, J.A., and Cohanim, "Incorporation of Flexibility into the Avionics Subsystem for the TALARIS Small Advanced Prototype Vehicle," AIAA 2010-8903, Space 2010 Conference and Exposition, Anaheim, 30 Aug. – 2 Sept. 2010.

[70] Mark, G.T., "Incorporating Flexibility into System Design: A Novel Framework and Illustrated Developments," M.S. Thesis, Massachusetts Institute of Technology, June 2005.

[71] Nilchiani, R., "Measuring Space Systems Flexibility: A Comprehensive Six-Element Framework," Ph.D. Thesis, Massachusetts Institute of Technology, Sept. 2005.

[72] Lim, D., "A Systematic Approach to Design for Lifelong Aircraft Evolution," Ph.D. Thesis, Georgia Institute of Technology, May 2009.

[73] Lim, D.W. and Mavris, D.N., "A Design Methodology for Lifelong Aircraft Evolution," AIAA 2010-0283, 48th Aerospace Sciences Meeting, Orlando, 4-7 Jan. 2010.

[74] Daniels, M., Tracey, B., Irvine, J., Schram, W., and Paté-Cornell, M.E., "Probabilistic Simulation of Multi-Stage Decisions for Operation of a Fractionated Satellite Mission," IEEEAC Paper #1323, IEEE Aerospace Conference, Big Sky, 5-12 March 2011.

[75] Bernanke, B.S., "Irreversibility, Uncertainty, and Cyclical Investment," *The Quarterly Journal of Economics*, Vol. 98, No. 1, Feb. 1983, pp. 85-106.

[76] Twiss, B.C., *Forecasting for Technologists and Engineers: A Practical Guide for Better Decisions,* Peter Peregrinus Ltd., London, 1992, chap. 5.

[77] Lafleur, J.M., Lantoine, G., Hensley, A.L., Retaureau, G.J., Kranzusch, K.M., Hickman, J.W., Wilson, M.N., and Schrage, D.P., "A Systematic Concept Exploration Methodology Applied to Venus In Situ Explorer," 6th International Planetary Probe Workshop, Atlanta, 23-27 June 2008.

[78] Dieter, G.E., *Engineering Design,* 3rd ed., McGraw-Hill, Boston, 2000, chap. 5.

[79] Engler, W., Biltgen, P., and Mavris, D., "Concept Selection Using an Interactive Reconfigurable Matrix of Alternatives (IRMA)," AIAA 2007-1194, 45th AIAA Aerospace Sciences Meeting and Exhibit, Reno, 8-11 Jan. 2007.

[80] Goldberg, M.S. and Touw, A.E., *Statistical Methods for Learning Curves and Cost Analysis,* INFORMS, Linthicum, 2003, chap. 2.

[81] Jordan, C., *Calculus of Finite Differences*, Chelsea Publishing Co., New York, 1965, chap. 1.

[82] Dahlquist, G. and Björck, Å., *Numerical Methods*, Prentice-Hall, Englewood Cliffs, 1974, chap. 7.

[83] Ames, W., *Numerical Methods for Partial Differential Equations*, Barnes and Noble, New York, 1969, chap. 1.

[84] Hildebrand, F.B., *Finite-Difference Equations and Simulations*, Prentice-Hall, Englewood Cliffs, 1968, chap. 1.

[85] Richardson, C.H., *An Introduction to the Calculus of Finite Differences,* D. Van Nostrand, Princeton, 1954, chap. 1.

[86] Ross, S.M., *Introduction to Probability Models,* 10th ed., Elsevier, Amsterdam, 2010, chaps. 4, 6.

[87] Winston, W.L., *Operations Research: Applications and Algorithms,* 4th ed., Duxbury Press, Pacific Grove, 2003, chaps. 11, 17, 19.

[88] Brito, M.P. and Griffiths, G., “A Markov Chain State Transition Approach to Establishing Critical Phases for AUV Reliability,” *IEEE Journal of Oceanic Engineering*, Vol. 36, No. 1, Jan. 2011, pp. 139-149.

[89] Turns. S.R., *An Introduction to Combustion: Concepts and Applications*, 2nd ed., McGraw-Hill, Boston, 2006, chap. 5.

[90] Cover, T.M. and Thomas, J.A., *Elements of Information Theory*, 2nd ed., Wiley-Interscience, Hoboken, 2006, chap. 4.

[91] Sen, P. and Yang, J., *Multiple Criteria Decision Support in Engineering Design*, Springer, London, 1998, chap. 4.

[92] Puterman, M.L., *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley-Interscience, Hoboken, 2005, chaps. 2, 4.

[93] Bellman, R., *Dynamic Programming*, Princeton University Press, Princeton, 1957, preface, chap. 3.

[94] “Broad Agency Announcement: System F6,” DARPA Tactical Technology Office, BAA07-31, 16 July 2007, Available: http://www.darpa.mil/TTO/solicit/BAA07-31/F6_BAA_Final_07-16-07.doc [21 Nov. 2008].

[95] Lafleur, J.M. and Saleh, J.H., “Exploring the F6 Fractionated Spacecraft Trade Space with GT-FAST,” AIAA 2009-6802, AIAA Space 2009 Conference and Exposition, Pasadena, 14-17 Sept. 2009.

[96] Eremenko, P., “System F6,” DARPA Tactical Technology Office Presentation, Feb. 2010, Available: http://www.darpa.mil/WorkArea/DownloadAsset.aspx?id=2648 [28 June 2011].

[97] Brown, O., “Program of Record Analysis.” DARPA/TTO Memorandum, 6 May 2008.

[98] Larson, W.J., and Wertz, J.R. (ed.), *Space Mission Analysis and Design*, 3rd ed., Microcosm Press and Kluwer Academic Publishers, El Segundo, 1999.

[99] NASA Goddard Space Flight Center, *National Space Science Data Center Spacecraft Query* [online database], Available: http://nssdc.gsfc.nasa.gov/nmc/SpacecraftQuery.jsp [13 Feb. 2009].

[100] Martin, D.H., Anderson, P.R., and Bartamian, L., *Communication Satellites*, 5th ed., The Aerospace Press, El Segundo, 2007.

[101] "SSTL VHRI 250," Surrey Satellite Technology, Available: http://www.sst-us.com/Downloads/Datasheets/Datasheet_VHRI250_OL [6 Aug. 2011].

[102] Carrel, A.R., Cawthorne, A.D., Richardson, G., and Gomes, L.M., "NigeriaSat-2: Advances in GN&C that Enable a High Performance Small Satellite Earth Observation Mission," AAS 10-033, 33rd AAS Rocket Mountain Guidance and Control Conference, Breckenridge, 5-10 Feb. 2010.

[103] Lafleur, J.M. and Saleh, J.H., "GT-FAST: A Point Design Tool for Rapid Fractionated Spacecraft Sizing and Synthesis," AIAA 2009-6563, AIAA Space 2009 Conference and Exposition, Pasadena, 14-17 Sept. 2009.

[104] NASA Johnson Space Center, Mission Operations Cost Model, Online Software, 25 May 2007, Available: http://cost.jsc.nasa.gov/MOCM.html [25 July 2011].

[105] Garber, S., Ed., "Biographical and other Personnel Information," NASA History Division, 17 July 2009, Available: http://history.nasa.gov/prsnnl.htm [11 Oct. 2010].

[106] Dumoulin, J., NSTS 1988 News Reference Manual [online database], Available: http://science.ksc.nasa.gov/shuttle/technology/sts-newsref/stsref-toc.html [17 Oct. 2011].

[107] Ertel, I.D. and Morse, M.L., *The Apollo Spacecraft: A Chronology*, Vol. 1, NASA SP-4009, NASA, Washington, 1969, pp. 127-128.

[108] Thompson, F.L., Borman, F., Van Dolah, R.W., Faget, M.A., et. al., *Report of Apollo 204 Review Board*, NASA-TM-84105, NASA, Washington, 1967, parts 4, 6.

[109] Bolden, C., *Fiscal Year 2012 Budget Estimates*, NASA, Available: http://www.nasa.gov/pdf/516675main_NASAFY12_Budget_Estimates-508.pdf [20 Oct 2011].

[110] Dubos, G.F., Saleh, J.H., and Braun, R., “Technology Readiness Level, Schedule Risk, and Slippage in Spacecraft Design,” *Journal of Spacecraft and Rockets*, Vol. 45, No. 4, 2008, pp. 836-842.

[111] United States Government Accountability Office, “NASA: Assessments of Selected Large-Scale Projects,” Report to Congressional Committees, GAO-11-239SP, Washington, 2011, p. 15.

[112] Mattson, C.A., Mullur, A.A., and Messac, A., “Minimal Representation of Multiobjective Design Space using a Smart Pareto Filter,” AIAA 2002-5458, 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Atlanta, 4-6 Sept. 2002.

[113] Grant, M. and Mendeck, G., “Mars Science Laboratory Entry Optimization Using Particle Swarm Methodology.” AIAA 2007-6393, Atmospheric Flight Mechanics Conference and Exhibit, Hilton Head, 20-23 Aug. 2007.

[114] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T., “A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II,” *IEEE Transactions on Evolutionary Computation*, Vol. 6, No. 2, April 2002, pp. 182-197.

[115] Daskilewicz, M.J. and German, B.J., “RAVE: A Graphically Driven Framework for Agile Design-Decision Support,” AIAA 2010-9033, 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, Fort Worth, 13-15 Sept. 2010.

[116] Saaty, T.L., *The Analytic Hierarchy Process*, McGraw-Hill, New York, 1980.

[117] Osman, M.S., Abo-Sinna, M.A., and Mousa, A.A., “An effective genetic algorithm approach to multiobjective resource allocation problems (MORAPs),” *Applied Mathematics and Computation*, Vol. 163, No. 3, 2005, pp. 755-768.

[118] Beltrami, E., Katehakis, M., and Durinovic, S., “Multiobjective Markov Decisions in Urban Modelling,” *Mathematical Modelling*, Vol. 6, No. 4, 1985, pp. 333-338.

[119] Zadeh, L.A., “Optimality and Non-Scalar-Valued Performance Criteria,” *IEEE Transactions on Automatic Control*, Vol. 8, No. 1, 1963, pp. 59-60.

[120] Saramago, S.F.P. and Steffen, Jr., V., “Optimization of the Trajectory Planning of Robot Manipulators Taking into Account the Dynamics of the System,” *Mechanism and Machine Theory*, Vol. 33, No. 7, 1998, pp. 883-894.

[121] Satak, N., Majji, M., and Hurtado, J.E., “Sequential Shape Estimation of an Asteroid Using Vision,” AAS 11-106, 21st Space Flight Mechanics Meeting, New Orleans, 13-17 Feb. 2011.

[122] Dutta, A., “On-Orbit Servicing of Satellites in Circular Constellation Using a Single Service Vehicle,” AAS 11-227, 21st Space Flight Mechanics Meeting, New Orleans, 13-17 Feb. 2011.

[123] Messac, A., Sundararaj, G.J., Tappeta, R.V., and Renaud, J.E., “Ability of Objective Functions to Generate Points on Nonconvex Pareto Frontiers,” *AIAA Journal*, Vol. 38, No. 6, June 2000, pp. 1084-1091.

[124] Messac, A. and Ismail-Yahaya, A., “Required Relationship Between Objective Function and Pareto Frontier Orders: Practical Implications,” *AIAA Journal*, Vol. 39, No. 11, Nov. 2001, pp. 2168-2174.

[125] Koski, J., “Defectiveness of Weighting Method in Multicriterion Optimization of Structures,” *Communications in Applied Numerical Methods*, Vol. 1, No. 6, Nov. 1985, pp. 333-337.

[126] Das, I. and Dennis, J.E., “A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems,” *Structural Optimization*, Vol. 14, No. 1, Aug. 1997, pp. 63-69.

[127] Chen, W., Wiecek, M.M., and Zhang, J., “Quality Utility – A Compromise Programming Approach to Robust Design,” *Journal of Mechanical Design*, Vol. 121, No. 2, June 1999, pp. 179-187.

[128] Bowman, V.J., “On the Relationship of the Tchebycheff Norm and the Efficient Frontier of Multiple-Criteria Objectives,” *Multiple Criteria Decision Making: Proceedings of a Conference*, Springer-Verlag, Berlin, 1976, pp. 76-85.

[129] Marler, R.T. and Arora, J.S., “Function-transformation methods for multi-objective optimization,” *Engineering Optimization*, Vol. 37, No. 6, Sept. 2005, pp. 551-570.

[130] Marler, R.T. and Arora, J.S., “The weighted sum method for multi-objective optimization: new insights,” *Structural and Multidisciplinary Optimization*, Vol. 41, No. 6, Dec. 2009, pp. 853-862.

[131] Bellman, R. and Dreyfus, S.E., *Applied Dynamic Programming*, Princeton University Press, Princeton, 1962.

[132] Wynholds, H.W. and Skratt, J.P., “Weapon System Parametric Life Cycle Cost Analysis,” *Proceedings of the 1977 Annual Reliability and Maintainability Symposium,* Philadelphia, 18-20 Jan. 1977, pp. 303-309.

[133] Cyr, K., Ed., *Spacecraft/Vehicle Level Cost Model* [online database], NASA Johnson Space Center, 25 May 2007, Available: http://cost.jsc.nasa.gov/SVLCM.html [20 Oct. 2011].

[134] National Aeronautics and Space Administration, “HAT Cycle 2011-B Element Baseball Cards,” Presentation Slides, 17 May 2011.

[135] National Aeronautics and Space Administration, “Lander Module Baseball Cards,” Presentation Slides, 20 May 2011.

[136] Drake, B.G., NASA Johnson Space Center / Constellation Program Systems Engineering and Integration Office, E-mail Communication, 13 Jan. 2011.

[137] Connolly, J.F., NASA Johnson Space Center / Constellation Program Altair Project Office, E-mail Communication, 17 June 2011.

[138] Drake, B.G., Ed., *Human Exploration of Mars Design Reference Architecture 5.0*, NASA SP-2009-566, NASA, Washington, July 2009, section 5.

[139] Petty, J.I., Ed., “Multi-Purpose Logistics Modules,” National Aeronautics and Space Administration, 23 Oct. 2010, Available: http://www.nasa.gov/mission_pages/station/structure/elements/mplm.html [22 Oct. 2011].

[140] Joosten, B.K. and Guerra, L.A., “Early Lunar Resource Utilization: A Key to Human Exploration,” AIAA 1993-4874, AIAA Space Programs and Technologies Conference and Exhibit, Huntsville, 21-23 Sept. 1993.

[141] NASA Jet Propulsion Laboratory, “NASA’s 2009 Mars Science Laboratory,” Presentation Slides, Available: http://marsoweb.nas.nasa.gov/landingsites/msl/memoranda/MSL_overview_LS.pdf [22 Oct. 2011].

[142] Wise, M.A., Lafleur, J.M., and Saleh, J.H., “Regression Analysis of Launch Vehicle Payload Capability for Interplanetary Missions,” IAC-10.D2.P06, 61st International Astronautical Congress, Prague, 27 Sept. – 1 Oct. 2010.

[143] Carr, C.E. and McGee, J., “The Apollo Number: Space Suits, Self-Support, and the Walk-Run Transition,” *PLoS ONE*, Vol. 4, No. 8, August 2009, e6614.

[144] Pratt & Whitney Rocketdyne, “RS-68 Propulsion System,” Canoga Park, 2005, Available: http://www.pw.utc.com/products/pwr/assets/pwr_rs-68.pdf [22 Oct. 2011].

[145] Pratt & Whitney Rocketdyne, “J-2X Propulsion System,” Canoga Park, 2007, Available: http://www.pw.utc.com/products/pwr/assets/pwr_j2x.pdf [22 Oct. 2011].

[146] Andrews Space & Technology, *Rocket Engines* [online database], 2001, Available: http://www.spaceandtech.com/spacedata/engines/engines.shtml [22 Oct. 2011].

[147] Purdue School of Aeronautics and Astronautics, *Liquid Rocket Engines* [online database], 1998, Available: https://engineering.purdue.edu/~propulsi/propulsion/rockets/liquids.html [22 Oct. 2011].

[148] Turnbough, L., Ed., “Commercial Crew and Cargo,” National Aeronautics and Space Administration, 8 Sept. 2011, Available: http://www.nasa.gov/offices/c3po/home/index.html [22 Oct. 2011].

[149] Space Exploration Technologies Corp., “Falcon 9 Overview,” Available: http://www.spacex.com/falcon9.php [24 Feb. 2011].

[150] Space Exploration Technologies Corp., “Dragon Overview,” Available: http://www.spacex.com/dragon.php [24 Feb. 2011].

[151] Stelly, J.M., NASA Johnson Space Center / Chief Financial Officer Cost Estimating and Assessment Office, E-mail Communication, 30 June 2011.

[152] Cyr, K., Ed., *Mission Operations Cost Model* [online database], NASA Johnson Space Center, 25 May 2007, Available: http://cost.jsc.nasa.gov/MOCM.html [25 Oct. 2011].

[153] Launius, R.D., “Assessing the legacy of the Space Shuttle,” *Space Policy*, Vol. 22, No. 4, Nov. 2006, pp. 226-234.

[154] Wilson, J., Ed., *NASA Budget Documents, Strategic Plans, and Performance Reports* [online database], National Aeronautics and Space Administration, 5 April 2011, Available: http://www.nasa.gov/news/budget/index.html [2 Nov. 2011].

[155] Rendon, R.G. and Mutty, J., “DoD Contract Termination Liability: An Analysis of the Special Termination Cost Clause (STCC),” Naval Postgraduate School Acquisition Research Sponsored Report, NPS-CM-06-042, 30 Sept. 2006.

[156] Stelly, J.M., NASA Johnson Space Center / Chief Financial Officer Cost Estimating and Assessment Office, Personal Communication, 30 June 2011.

[157] Shishko, R., “The Application of Architecture Frameworks to Modelling Exploration Operations Costs,” 16th Annual International Symposium of the International Council on Systems Engineering, Orlando, 8-14 July 2006.

[158] Shishko, R., “Building an O&S Cost Model on DoDAF Concepts,” *Parametric World*, Vol. 28, No. 1, 2009, pp. 4-6, 9.

VĪTA

Jarret M. Lafleur was born in 1984 and grew up in the villages of Nasonville and Mapleville within the town of Burrillville, Rhode Island. A student of Burrillville High School, he graduated as valedictorian in 2002. Jarret enrolled at the Georgia Institute of Technology in 2002 and graduated with the institute's Highest Honor and a B.S. in Aerospace Engineering in 2007. As an undergraduate, Jarret engaged in work experiences at the Naval Undersea Warfare Center, NASA White Sands Test Facility, and NASA Johnson Space Center. He also conducted research in Georgia Tech's Space Systems Design Lab (SSDL) under Dr. John Olds and Dr. Robert Braun, contributing in part to his 2005 selection as an Astronaut Scholar.

Jarret continued within the SSDL as a graduate student in 2007, beginning studies on flexibility in space system design under the advisement of Dr. Joseph Saleh. From 2007-2009 he supported work on the DARPA System F6 fractionated spacecraft program. During 2009-2010, he was selected as one of ten Sam Nunn Security Program fellows at the institute, focusing on policy aspects of space infrastructure security. Jarret additionally continued to contribute to the Flight Mechanics and Trajectory Design Branch at NASA Johnson Space Center during summers, often supporting Mars aerocapture, entry, descent, and landing trajectory assessments for strategic planning of future human-class missions. He earned his M.S. in Aerospace Engineering in 2009, and was supported in his studies by a National Defense Science and Engineering Graduate Fellowship and National Science Foundation Graduate Research Fellowship. During his time at Georgia Tech, Jarret authored over 30 conference and journal publications.

In spare time, Jarret is an avid musician, having earned a Music Minor and at various times served as principal flutist Georgia Tech's Symphony Orchestra, Symphonic Band, Concert Band, and Marching Band. He also enjoys playing piano and, when time allows or when Star Trek re-runs are not on the air, composing for small ensembles.