

**A Framework for Dynamic Safety and Risk Management
Modeling in Complex Engineering Systems**

by

Nicolas Dulac

B.Eng., Mechanical Engineering, McGill University, 2001
S.M., Aero-Astro Engineering, MIT, 2003

Submitted to the Department of Aeronautical and Astronautical Engineering
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2007

© Nicolas Dulac. All rights reserved.

The author thereby grants to MIT permission to reproduce
and to distribute publicly paper and electronic
copies of this thesis document in whole or in part.

Signature of Author

Department of Aeronautics and Astronautics
February 22nd, 2007

Certified by

Nancy G. Leveson, Professor
Committee Chair, Department of Aeronautics and Astronautics

Certified by

Deborah J. Nightingale, Professor
Department of Aeronautics and Astronautics

Certified by

Nelson P. Repenning, Professor
Sloan School of Management

Accepted by

Jaime Peraire, Professor
Department of Aeronautics and Astronautics
Chair, Committee on Graduate Students

[This page is intentionally left blank]

ABSTRACT

Almost all traditional hazard analysis or risk assessment techniques, such as failure modes and effect analysis (FMEA), fault tree analysis (FTA), and probabilistic risk analysis (PRA) rely on a chain-of-event paradigm of accident causation. Event-based techniques have some limitations for the study of modern engineering systems. Specifically, they are not suited to handle complex software-intensive systems, complex human-machine interactions, and systems-of-systems with distributed decision-making that cut across both physical and organizational boundaries.

STAMP (System-Theoretic Accident Model and Processes) is a comprehensive accident model created by Nancy Leveson that is based on systems theory. It draws on concepts from engineering, mathematics, cognitive and social psychology, organizational theory, political science, and economics. The general notion in STAMP is that accidents result from inadequate enforcement of safety constraints in design, development, and operation. STAMP includes traditional failure-based models as a subset, but goes beyond physical failures to include causal factors involving dysfunctional interactions among non-failing components; software and logic design errors; errors in complex human decision-making; various organizational characteristics such as workforce, safety processes and standards, contracting; and other managerial, social, organizational, and cultural factors.

The main contribution of this thesis is the augmentation of STAMP with a dynamic executable modeling framework in order to further improve safety in the development and operation of complex engineering systems. This executable modeling framework: 1) enables the dynamic analysis of safety-related decision-making in complex systems, 2) assists with the design and testing of non-intuitive policies and processes to better mitigate risks and prevent time-dependent risk increase, and 3) enables the identification of technical and organizational factors to detect and monitor states of increasing risk before an accident occurs.

The modeling framework is created by combining STAMP safety control structures with system dynamic modeling principles. A component-based model-building methodology is proposed to facilitate the building of customized STAMP-based dynamic risk management models and make them accessible to managers and engineers with limited simulation experience. A library of generic executable components is provided as a basis for model creation, refinement, and validation. A toolset is assembled to identify risk increase patterns, analyze time-dependent risks, assist engineers and managers in safety-related decision-making, create and test risk mitigation actions and policies, and monitor the system for states of increasing risk.

The usefulness of the new framework is demonstrated in two independent projects: 1) A risk analysis of the NASA Independent Technical Authority (ITA), an organization mandated by the Columbia Accident Investigation Board (CAIB) to provide independent safety oversight

of space shuttle operations, and 2) A risk management study for the Exploration Systems Mission Directorate (ESMD) at NASA. For these two projects, model refinement, validation and analysis required extensive data collection and interactions with NASA workforce. Over 45 interviews were conducted at five NASA centers (HQ, MSFC, KSC, JSC, and LaRC). Interviewees included representatives from the Office of the Administrator, the Office of the Chief Engineer, the Office of Safety and Mission Assurance, ESMD Directorate Offices, Program/Project Offices, and many others. Among other data sources, 200 pages of interview transcripts were compiled and used for model creation and validation activities. Specific risks analyzed include: 1) NASA workforce and knowledge management issues, 2) the impact of various levels of outsourcing, 3) the impact of safety priority on design, and 4) the impact of requirements change on safety and schedule during development.

ACKNOWLEDGMENTS

First and foremost, I would like to thank Nancy, my academic advisor, thesis supervisor, and friend. Your guidance, support and enthusiasm throughout all these years at MIT made this work possible.

I would also like to thank my committee members (Debbie, John, Nelson, etc.) for their guidance, as well as all the members of the Columbia group (Betty, John, Joel and all the others...), who generated many of the insight upon which this research is based. Thank you to all the people at NASA who agreed to talk to an evil foreign national and whose help with the models was invaluable.

Thanks also to the CSRL crew at MIT (Joe, Steve, Maggie and all the others who come and go...). Special thanks to Brandon for his enthusiasm and help with the interviews and models.

On a less academic note, first thanks go to Martin, for being such a good friend and for bringing me to the gym to ventilate during tough times. Thanks also to the MIT varsity hockey coaches (Marky, Jimmy, Donny and Froggy) for allowing me to extend my career for five years.

Thank you to everyone who participated in making my time at MIT and in Boston interesting, rewarding, and mostly a lot of fun. In no particular order, it would include: Etienne, Yves, the Troquet crew, the hockey teammates, Katie and Minou, my roommates Deb and Allison, Delphine, all the McGill friends who come down once a year for BostonFest (Frank, PM, JS, Alex, Marc and all...), my best buddy J-F, who's always up for taking a ride on the DC-Boston shuttle. It's all been very much appreciated. JennyBess, obviously, because you are awesome. Three cheers for pink cupcakes! I'm forgetting so many people, but that's the way it goes...

And last but not least, the most heart felt thanks go to my family (Mom, Dad, Frank and Marie) for their unconditional love and support.

TABLE OF CONTENTS

Abstract	3
Acknowledgments.....	5
Table of Contents	6
List of tables.....	14
List of figures	15
INTRODUCTION: On the Dynamic Nature of Safety and Risk in Complex Engineering Systems	21
CHAPTER 1: Safety and Risk Dynamics in Complex Systems - Background and Literature Review.....	24
1.1 Risk in Complex Engineering Systems.....	24
1.1.1 Definition	25
1.1.2 Risk in System Safety Engineering.....	27
1.1.3 Origins of Risk: Uncertainty	28
1.1.4 Quantitative Risk Assessment (QRA).....	28
1.2 Safety Definition	31
1.3 Accidents.....	32
1.3.1 Definition	32
1.3.2 Accidents in Complex Socio-Technical Systems	32
1.3.3 Traditional Accident Models.....	34
1.4 Organizational Risk Theories.....	38
1.4.1 Normal Accident Theory (NAT).....	38
1.4.2 High Reliability Organizations (HROs).....	39
1.4.3 Limitations of Organizational Risk Theories	40
1.5 A Systems Theoretic Approach to Safety and Accident Modeling	40
1.6 Definitions and Terminology	43
1.6.1 Accident	43
1.6.2 Incident.....	43
1.6.3 Safety.....	43

1.6.4	Reliability.....	43
1.6.5	Failure	43
1.6.6	Hazard	44
1.6.7	Hazard Analysis	44
1.6.8	Hazard Mitigation	45
1.6.9	Risk	45
1.6.10	Risk Analysis	45
1.6.11	Risk Management.....	45
1.7	Thesis Scope and boundaries	46
1.7.1	Thesis Objective.....	47
1.7.2	Thesis Outline	48
CHAPTER 2: Lifecycle Risk Management Modeling For Complex Socio-Technical		
	Systems	50
2.1	The STAMP Model of Accident Causation.....	51
2.1.1	Safety Constraints	54
2.1.2	Hierarchical Safety Control Structures	54
2.1.3	Control Loops and Process Models	55
2.2	System Dynamics Concepts and Applications.....	58
2.2.1	Review of Common System Dynamics Applications.....	63
2.3	Safety-Centric System Development	64
2.3.1	Loop B1 - Delays Cause Pressure.....	66
2.3.2	Loop R2: The Rework/Burnout Cycle	70
2.3.3	Loop R3: The Burnout Cycle.....	71
2.3.4	Loops R1a/b: Safety and Integration	73
2.3.5	Outer loops: Waivers, cost and resources	77
2.3.6	Impact of System Development on Operations	79
2.4	System Operations, Safety Erosion and Migration Toward High-Risk.....	80
2.5	Example and Summary	85
CHAPTER 3: Guidelines and Criteria for the Creation, Analysis and Monitoring of Static		
	Safety Control Structures	86
3.1	Selecting Control Structure Components.....	88

3.1.1	Listing Relevant Organizational Components	88
3.2	Connecting Control Structure Components	95
3.2.1	Generic Connector Types.....	96
3.2.2	Documenting Organizational Components	107
3.3	Analyzing Safety Control Structures	111
3.3.1	Completeness Criteria	112
3.3.2	Consistency Criteria.....	114
3.4	Formal and Informal structures and COnections	119
3.5	New Control Structures vs. Existing Control Structures	119
3.6	Monitoring Safety Control Structures.....	120
3.6.1	From Component to Individuals: Monitoring using Social Networks.....	121
3.7	Summary	122
CHAPTER 4:	Creating Dynamic Risk Management Models Using Generic Customizable Dynamic Components.....	124
4.1	Model Building Methodology.....	124
4.1.1	Using Generic Structures, Components, and Archetypes	124
4.1.2	Combining STAMP safety control structures with causal loop structures ...	125
4.2	Creating Dynamic Risk Management Models based on STAMP and System Dynamics	127
4.2.1	Step 1: Initial System Characterization.....	131
4.2.2	Step 2: Mapping of Static Safety Control Structure to Generic Dynamic Components	137
4.2.3	Step 3: Refinement of Dynamic Safety Model Structure	140
4.2.4	Step 4: Mapping of Pressures, Influences and Reporting Channels	141
4.2.5	Step 5: Data Collection and Component Calibration.....	146
4.2.6	Step 6: Component Testing and Confidence-Building Activities.....	157
4.2.7	Step 7: Component-Based Model Assembly	161
4.2.8	Step 8: Model Testing and Confidence-Building Activities.....	164
4.3	Moving Forward: Model Analysis	167
CHAPTER 5:	Analysis of Dynamic Safety and Risk Management Models.....	169
5.1	improving Safety/Risk Decision Making.....	169

5.1.1	Management Simulators.....	170
5.1.2	Visualization of Model Structure and Behavior.....	172
5.1.3	Interactive Scenario-Based Learning and Decision-Aid.....	175
5.2	Analyzing identified risk Scenarios	176
5.2.1	Risk Identification.....	176
5.2.2	Scenario Creation and Analysis	179
5.2.3	Recommendations and Policy Development based on Scenario Results.....	182
5.2.4	Policy Robustness and Scenario Sensitivity Analysis	183
5.3	Preventing Risk Increase and improving robustness to safety erosion	183
5.3.1	Preventing Time-Dependent Safety Erosion.....	184
5.3.2	Improving System Robustness to Risk Increase	185
5.4	Improving Risk monitoring in complex systems	186
5.4.1	Identifying Leading Indicators.....	187
5.4.2	Leading Indicator Sensitivity.....	189
5.5	Summary	190
CHAPTER 6: Case Study - Risk Management in the Development of NASA's Space		
	Exploration System	192
6.1	ESMD Model-Building Methodology	192
6.1.1	Step 1: ESMD Initial System Characterization.....	192
6.1.2	Step 2: Mapping of Static Safety Control Structure to Generic Dynamic Components	193
6.1.3	Step 3: Refinement of Dynamic Safety Model Structure	195
6.1.4	Step 4: Mapping of Pressures, Influences and Reporting Channels	196
6.1.5	Step 5: Data Collection and Component Calibration.....	198
6.1.6	Step 6: Component Testing and Confidence-Building Activities.....	205
6.1.7	Step 7: Component-Based Model Assembly	206
6.1.8	Step 8: Model Testing and Confidence-Building Activities	206
6.2	Individual ESMD Component Documentation.....	207
6.2.1	Congress and Executive (White House) Component.....	207
6.2.2	NASA Administration and ESMD Component	207
6.2.3	Exploration Program and Project Management Component	208

6.2.4	System Development Completion and Safety Analyses Component	208
6.2.5	Engineering – Technical Personnel Resources and Experience Component	209
6.2.6	Engineering – Effort and Efficacy of Technical Personnel Component.....	210
6.2.7	Safety and Mission Assurance - Effort and Efficacy of System Safety Analysts (EESSA) Component	210
6.3	Integrated ESMD Model Overview	210
6.4	Model Results and Scenario Analysis.....	213
6.4.1	Scenario 1: Workforce Planning	215
6.4.2	Scenario 2: Investigating Management Reserves	218
6.4.3	Scenario 3: Effect of Schedule Pressure and Safety Priority	221
6.4.4	Scenario 4: Effect of High Safety Influence and Power on Decision- Making	224
6.4.5	Scenario 5: Effect of Changes in Requirements and System Scope	228
6.5	Case Study Summary	230
CHAPTER 7:	Contributions, Future Work and Conclusion	232
7.1	Summary	232
7.2	Contributions.....	232
7.2.1	Static STAMP Control Structures Creation and Analysis Guidelines	232
7.2.2	Model-Building Methodology	233
7.2.3	Model-Based Risk Analysis Techniques	234
7.2.4	Interview-Based Model Validation Protocol	234
7.2.5	Project-Specific Insights	235
7.3	Future Work	235
7.3.1	Model-Building and Analysis Software Tools.....	235
7.3.2	Model Visualization Tools.....	236
7.3.3	Validation in Other Domains	236
7.3.4	Micro-Theories of Risk Increase in Complex Systems	236
7.4	Conclusion	238
APPENDIX A:	Acronyms	239
APPENDIX B:	Types of Uncertainty.....	241
B.1	Ambiguity	241

B.2 Aleatory Uncertainty.....	241
B.3 Epistemic Uncertainty.....	242
B.3.1 Model.....	242
B.3.2 Phenomenological.....	243
B.3.3 Behavioral.....	243
B.4 Interaction.....	244
APPENDIX C: Organizational Risk THEories.....	245
C.1 Normal Accident Theory (NAT).....	245
C.1.1 NAT Limitations.....	246
C.2 High Reliability Organizations (HROs).....	247
C.2.1 HRO Limitations.....	247
C.3 Debate and limitations.....	254
APPENDIX D: Impact of Balancing and Reinforcing Loops: The Case of the NASA	
Independent Technical Authority.....	257
D.1 ITA Context and Modeling.....	258
D.2 Initial ITA Model Analysis.....	260
D.3 ITA Behavior Mode Analysis.....	261
D.3.1 Behavior Mode #1: Successful ITA Implementation.....	262
D.3.2 Behavior Mode #2: Unsuccessful ITA Implementation.....	263
APPENDIX E: Repository of Generic Dynamic Components.....	266
E.1 Research Approach for Component Creation.....	266
E.2 Projects and Case Studies.....	267
E.3 Data Sources.....	268
E.3.1 Group model building and review from safety and organization Theory	
experts.....	268
E.3.2 Books, Accident Reports, and Risk Literature.....	268
E.3.3 Interviews with Domain Experts.....	269
E.3.4 Domain-Specific Quantitative Data Sources.....	269
E.4 Data Analysis.....	269
E.5 Repository of Dynamic Generic Components for System Operation.....	271
E.5.1 Congress and Executive Component.....	271

E.5.2	Regulatory Agency Component.....	274
E.5.3	Company Management Component.....	275
E.5.4	Operations Management Component – Problem Resolution and Learning..	276
E.5.5	Operations Management Component – Production System	277
E.5.6	System Maintenance and Evolution Component	278
E.5.7	System Safety Component - Status, Effort and Efficacy	279
E.5.8	System Safety Component - Knowledge, Skills and Staffing.....	280
E.5.9	Accident Causation and Triggering Component.....	281
E.6	Repository of Generic Dynamic Components for System Development	282
APPENDIX F:	NASA ESMD Project Research Protocol	291
APPENDIX G:	ESMD PROJECT HISTORY AND BACKGROUND	293
APPENDIX H:	ESMD Model Data Sources	295
H.1	Interviews with NASA Officials	295
H.2	Additional Data Sources.....	299
APPENDIX I:	Model Conventions and Assumptions	302
APPENDIX J:	Sample Component Integration.....	306
APPENDIX K:	Model Documentation.....	308
K.1	Congress and Executive (White House) Component.....	309
K.2	NASA Administration and ESMD Component	312
K.3	Exploration Program and Project Management Component	315
K.4	System Development Completion and Safety Analyses Component	318
K.5	Engineering – Technical Personnel Resources and Experience Component.....	321
K.6	Engineering – Effort and Efficacy of Technical Personnel Component.....	324
K.7	Safety and Mission Assurance - Effort and Efficacy of System Safety Analysts (EESSA) Component	326
References	328

LIST OF TABLES

<i>Number</i>	<i>Page</i>
Table 1: List of Development and Operation Generic Components.....	91
Table 2: Summary of mapping between ITA control structure components and generic components	140
Table 3: Shuttle Program/Project Manager Risk [Leveson, 2005]	179
Table 4: Summary of mapping between NASA/ESMD control structure components and generic components.....	195
Table 5: Sample Table for Congress and Executive Component Variables	273
Table 6. Budget and personnel data sources and types of data used in the model.....	300

LIST OF FIGURES

<i>Number</i>	<i>Page</i>
Figure 1: Risk Components (from [Leveson, 1995]).....	27
Figure 2: Uncertainty Taxonomy for Engineering Systems	28
Figure 3: Components of risk (Adapted from Leveson)	45
Figure 4: Generic Safety Control Structure (from [Leveson, 2004]).....	55
Figure 5: A Generic Control Loop (Adapted from [Leveson, 2004]).....	57
Figure 6: A Taxonomy of Generic Control Flaws leading to Hazards (from [Leveson, 2005]).....	58
Figure 7: Three basic feedback structures.....	60
Figure 8: The bathtub analogy of stock and flows.....	61
Figure 9: The stock and flow structure of a classic predator-prey dynamic model	62
Figure 10: Dynamic behavior generated by the predator-prey structure	62
Figure 11: Generic System Development Structure (from [Leveson, 2004]).....	66
Figure 12: Loop B1: Delays cause pressure.....	66
Figure 13: Planned and actual development completion fraction.....	67
Figure 14: Impact of disturbance on completion fraction.....	68
Figure 15: Completion fraction with disturbance and controller with $P>0$ and $I=0$	69
Figure 16: Completion fraction with disturbance and controller with $P>0$ and $I>0$	69
Figure 17: Adding loop R2: Burnout Cycle.....	70
Figure 18: Completion fraction including the burnout cycle (R2).....	70
Figure 19: Adding loop R3: Burnout Rework Cycle	72
Figure 20: Impact of burnout cycle on rework fraction	72
Figure 21: Impact of burnout and rework on completion fraction.....	73
Figure 22: Adding loops R1 and R1b: The impact of safety and integration on rework cycle	74
Figure 23: Impact of loops R1 and R1b on rework fraction	77
Figure 24: Impact of loops R1 and R1b on completion fraction	77
Figure 25: The Impact of Cost and Schedule on Funding and Pressure	79

Figure 26: Simple generic model of safety dynamics during system operation	81
Figure 27: Structure-created oscillatory behavior of model state variables.....	83
Figure 28: Anchoring safety efforts by limiting the impact of budget pressure and complacency.....	84
Figure 29: The STAMP-Based risk analysis process (adapted from [Leveson, 2005]) ...	87
Figure 30: Deming's Plan-Do-Check-Act (PDCA) cycle	89
Figure 31: Technical Authority Flow (Adapted from [NASA, 2005])	90
Figure 32: Structure of NASA Exploration System Mission Directorate (ESMD).....	92
Figure 33: Direct report (authority) link between components.....	97
Figure 34: Direct report structure in the NASA ITA	98
Figure 35: Direct oversight link between components	98
Figure 36: Direct oversight structure in the NASA ITA.....	99
Figure 37: Progress report link between components	99
Figure 38: Progress report structure in the NASA ITA	100
Figure 39: Performance appraisal link between components	100
Figure 40: Performance appraisal structure in the NASA ITA.....	101
Figure 41: Direct resource allocation link between components.....	101
Figure 42: Direct resource allocation structure of the NASA ITA	102
Figure 43: Coordination and technical information exchange link between components	102
Figure 44: Coordination and technical information exchange structure of the NASA ITA	103
Figure 45: Physical co-location link between components.....	103
Figure 46: Using the transitivity connection relation.....	104
Figure 47: Physical co-location structure of the NASA ITA.....	105
Figure 48: Personnel appointment link between components	105
Figure 49: Personnel appointment structure of the NASA ITA.....	106
Figure 50: Procurement link between components.....	106
Figure 51: Superimposing physical co-location and coordination links.....	115
Figure 52: Superimposing performance appraisal and direct report (authority) links	118
Figure 53: The System Dynamics Modeling Cycle [Sterman, 2000]	126

Figure 54: Modeling, Scenario Analysis and Policy Development as a Feedback Process	127
Figure 55: Focus on Step 6: STAMP-Based System Dynamics Model Building	128
Figure 56: Summary of the component-based model building methodology.....	130
Figure 57: Alternative Convergence-Style Flowchart Structure.....	131
Figure 58: Coupling and Interaction Complexity Chart (adapted from [Perrow, 1999])	133
Figure 59: Suggested mishap severity categories (From MIL-STD-882D [DoD, 2000])	134
Figure 60: Step 2 - Mapping of the ITA Static Safety Control Structure to Generic Dynamic Components.....	139
Figure 61: Step 3 - Re-structuring of safety control structure of the ITA.....	141
Figure 62: Mapping of Safety Oversight, Safety Pressure and Performance Pressure on the ITA dynamic control structure.....	144
Figure 63: Mapping of Performance Reports, and Problems, Incidents and Accident Reports on the ITA dynamic control structure.....	145
Figure 64: Mapping of Safety and Operations Resource Pressure on the ITA dynamic control structure	146
Figure 65: Component inputs, state variables, and outputs.	147
Figure 66: Isolating a component by "cutting out" dynamic connectors	148
Figure 67: Input-Output virtual container created using the "Free-Component Technique"	149
Figure 68: Example generic component (Company Management) and associated I/O structure.....	152
Figure 69: Example Human Resource Component Structure	154
Figure 70: Sample relationships used to define intermediate variables.....	155
Figure 71: The Company Management component exhibiting equilibrium behavior....	157
Figure 72: Deactivating the "Pushing the Limits" Loop.....	159
Figure 73: The Company Management component tested under the no-accident condition	160
Figure 74: The Company Management component tested under the accident response condition. An accident occurs at t=60 months.....	161
Figure 75: Connecting components through generic connectors and interfaces.....	163

Figure 76: Example of quasi-equilibrium behavior for integrated model	165
Figure 77 : Sample no-accident behavior for integrated model.....	166
Figure 78: Sample accident response behavior for the integrated model	167
Figure 79: Main interface of the Manned Space Program Risk Management Simulator [Friedenthal, 2006].....	171
Figure 80: Interface showing a 4x3 risk matrix for tracking the risk scores of the leading indicators [Friedenthal, 2006]	172
Figure 81: Context of the ITA subcomponent within the larger ITA model	173
Figure 82: Screen capture of three feedback loops of the ITA model [Friedenthal, 2006]	174
Figure 83: Screen capture of multiple loops of the ITA model [Friedenthal, 2006]	175
Figure 84: Effect of Increased Contracting on Risk Level	180
Figure 85: Effect of Increased Contracting on ITA Effectiveness and Credibility.....	181
Figure 86: Effect of Increased Contracting on Availability of High-Level Technical Personnel.....	182
Figure 87: Possible outcomes of the migration toward high risk process.....	184
Figure 88: Waiver Accumulation Pattern for ITA Risk Increase	188
Figure 89: A Potential Early Indicator: Problems under Investigation.....	189
Figure 90: The Balancing Loop Becomes Reinforcing as Workload Increases	189
Figure 91: Example of a Smaller "Metrics Dynamics" Model	190
Figure 92: Mapping generic components to the NASA/ESMD structure	194
Figure 93: Components and Structure of the NASA/ESMD Model.....	196
Figure 94: Mapping of performance, resource and safety pressures on the ESMD model structure.....	197
Figure 95: Mapping of progress and problem reports on the ESMD model structure....	198
Figure 96: Generic Administration Component for Development Organization.....	200
Figure 97: Initial Technical Personnel Component	203
Figure 98: Technical Personnel Component after two iteration cycles	204
Figure 99: Sample equilibrium behavior for the Administration component.....	205
Figure 100: Schematic of task completion flows.....	209

Figure 101: Management pressure feedback loops reproduced through component interactions	211
Figure 102: Interplay of schedule pressure and system safety.....	212
Figure 103: Increasing demand for ESMD technical civil servants (Fixed civil servant to contractor ratio).....	217
Figure 104: Ratio of productive civil servants to support contractors	217
Figure 105: Fraction of work remaining for the baseline and envelope scenarios	220
Figure 106: Relative safety of the final system for the baseline and envelope scenarios	220
Figure 107: Outcome distributions (completion time and safety) for the sensitivity analysis	221
Figure 108: Outcome (Safety, Schedule, Cost) as a function of schedule and safety priority (low, high).....	223
Figure 109: Estimated relative cost as a continuous function of schedule pressure and safety priority	224
Figure 110: Impact of safety influence and power on project dynamics	226
Figure 111: Impact of high-level technical leaders on project dynamics	227
Figure 112: Effect of scope and requirements change on project dynamics and outcome	229
Figure 113: Mass-Spring-Damper Simplified Model	243
Figure 114: The Interactions in the Zeebrugge Ferry Accident (from [Rasmussen, 1997])	250
Figure 115: The Nine Subsystem Models and their Interactions.....	258
Figure 116: ITA Model Structure	260
Figure 117: ITA Sensitivity Analysis Trace Results	261
Figure 118: Behavior Mode #1 Representing a Successful ITA Program Implementation	263
Figure 119: Behavior Mode #2 Representing an Unsuccessful ITA Program Implementation	264
Figure 120: Congress and Executive Component.....	272
Figure 121: Regulatory Agency Component	274
Figure 122: Company Management Component	275

Figure 123: Operations Management Component - Problem Resolution and Learning.	276
Figure 124: Operations Management Component - Production System	277
Figure 125: System Maintenance and Evolution Component	278
Figure 126: System Safety Component – Status, Effort and Efficacy.....	279
Figure 127: System Safety Component – Knowledge, Skills and Staffing	280
Figure 128: Operating Process Component - System Technical Risk and Safety Dynamics	281
Figure 129: Congress and Executive Component (Development).....	282
Figure 130: Safety Regulatory Agency (Development).....	283
Figure 131: Company Administration (Development)	284
Figure 132: Project and Program Management (Development)	285
Figure 133: Engineering - System Development Completion and Safety Analyses	286
Figure 134: Technical Personnel Resources and Experience (Development)	287
Figure 135: Engineering - Effort and Efficacy of System Safety Analysts	288
Figure 136: Engineering - Efforts and Efficacy of Non-Safety Technical Personnel (Development).....	289
Figure 137: Safety and Mission Assurance Personnel and Experience	290
Figure 138: Exploration-Centric NASA Structure.....	297
Figure 139: Model components and interviews	299
Figure 140: Databases on the NASA people website.....	301
Figure 141. Assumed Breakdown of Procurement Efforts [MSFC, 2002].....	303
Figure 142: Congress and Executive Component Structure	309
Figure 143: NASA Administration and ESMD Component Structure.....	312
Figure 144: Exploration Program and Project Management Component Structure	315
Figure 145: System Development Completion and Safety Analyses Component Structure	318
Figure 146: Engineering – Technical Personnel Resources and Experience Structure ..	321
Figure 147: Engineering – Effort and Efficacy of Technical Personnel Component Structure	324
Figure 148: Safety and Mission Assurance - (EESSA) Component Structure	326

INTRODUCTION: ON THE DYNAMIC NATURE OF SAFETY AND RISK IN COMPLEX ENGINEERING SYSTEMS

Since the beginning of the industrial revolution in the late 18th century, the cause of many serious accidents has shifted from natural causes to human and technology-related causes. While natural disasters still account for a significant amount of human and material losses, man-made disasters are responsible for an increasingly large portion of the toll. In addition, the boundary between natural and man-made disasters becomes ever more blurry as humans increasingly tamper (intentionally or not) with their natural environment. Natural disasters such as hurricane Katrina that was responsible for over 1400 human casualties and over 75 billion dollars of damage in the New Orleans area in 2005 cannot be entirely prevented. However, the man-made systems created to mitigate their effect may exacerbate the problem by providing a false sense of security. During hurricane Katrina, the levees protecting the city of New Orleans arguably hindered evacuation procedures because citizens and authorities believed they had ample time on their hand before the surge caused the levees to topple off [Davis, 2006]. When the poorly designed and implemented levees breached, the water level increased so rapidly that emergency response became overwhelmed and ineffective.

Examples abound where the interaction between the environment and man-made systems, safety devices or policies have increased the consequences of normally benign events and disturbances. In the 19th century, an increasing number of levies were constructed along the Mississippi to protect villages against surges and to increase farmland area by drying up marshes. This upstream reduction in marshes and wetland area caused a decline in the natural surge damping and absorption capacity of the river banks, effectively moving the problem downstream and exacerbating the net amplitude of seasonal surges.

Indirect interactions between humans and their environment may also contribute to risk. The hurricane season of 2005 was the worst in history. At the same time, the levels of greenhouse gases in the atmosphere are at the highest level in 650,000 years and the five warmest years of the last century occurred in the past 7 years. Scientists agree that this increase in greenhouse gases is at least partially due to human activity and contributes to global warming [Cicerone,

2001]. While a causality link has not been officially established, accumulating scientific evidence also links global warming to the late increase in the incidence and severity of natural cataclysms. Many eminent scientists such as James Lovelock, who proposed the “Gaia Hypothesis” [Lovelock, 1979] where the earth is seen as a self-regulating system, that is a system governed by a negative feedback loop, believe that a tipping point will soon be reached where the polarity of the loop will be reversed, making the system unstable and leading to deforestation, dramatic increases in carbon dioxide, further warming and catastrophic floods. Unless we are able to appreciate the system-level effects of human actions and policies, we are bound to repeat the cycle of errors that contribute to ever more dangerous man-made systems with their associated consequences.

Many similar feedback loops are active at the level of complex engineering systems and organizations. Increasing emphasis on low cost and performance foster the creation of systems with very little built-in “slack”. These tightly coupled systems, although usually more efficient, operate closer to the safety boundary, making them much more vulnerable to small disturbances that could escalate into major catastrophes. Very often, systems are initially designed and built with enough safety margins for sustained safe operation. However, as the system operates successfully over time, multiple feedback processes including performance and economic pressures cause an incremental erosion of safety margins. Dekker illustrates this erosion process in a detailed timeline of the changes in maintenance and operating requirements of the MD-11 elevator that led to the Alaska Airlines accident [Dekker, 2005].

In effect, complex socio-technical systems have a tendency to slowly migrate from a safe state toward a higher-risk state, where they are highly vulnerable to small disturbances. Once the system operates in this high-risk state, any number of different seemingly inconsequential events can lead to an accident. If one event does not trigger the loss, another one will. The Bhopal accident provides a good example of a system operating in a high-risk state [Leveson, 2006]. The release of methyl isocyanate from the Union Carbide chemical plant in Bhopal, India, in 1984 caused 2000 human casualties, 10,000 permanent disabilities, and over 200,00 injuries, arguably making it the worst industrial disaster in history [Shrivastava, 1992; Leveson, 1995]. The accident was officially blamed on human error as the worker assigned to

wash out some pipes and filters in the plant did not insert a safety disk as required. Without the safety disk, wash water leaked through a faulty valve and came in contact with methyl isocyanate. The resulted chemical reaction increased the temperature and pressure in the tank until the relief valve opened, releasing highly toxic chemicals in the atmosphere, which were then carried by the wind to populated areas. A more careful observation of the context in which the accident took place uncovers dozens of irregularities, disabled safety equipment, management negligence and regulatory deficiencies that all contributed to the accident. The Bhopal Union Carbide plant was a disaster waiting to happen. If the worker had inserted the safety disk on that day of December 1984, another small event or mistake would have eventually triggered an accident. Rasmussen [Rasmussen, 1997] explains this migration process:

“The stage for an accidental course of events very likely is prepared through time by the normal efforts of many actors in their respective daily work context, responding to the standing request to be more productive and less costly. Ultimately, a quite normal variation in somebody’s behavior can then release an accident. Had this ‘root cause’ been avoided by some additional safety measure, the accident would very likely be released by another cause at another point in time. In other words, an explanation of the accident in terms of events, acts, and errors is not very useful for design of improved systems [Rasmussen, 1997].”

The final objective of this dissertation is to design and operate engineering systems that will remain safe during their entire lifecycle. In order to achieve this objective, it is necessary to understand and define the dynamic feedback processes that may cause risk to increase over time in complex socio-technical systems. This dissertation introduces a framework to model some critical aspects of safety in complex systems. Subsequently, new tools and methods based on a system-theoretic accident model are proposed to create customized dynamic risk management system to help decision-makers in managing risks, and avoid an eventual migration of systems toward a state of high risk.

CHAPTER 1: SAFETY AND RISK DYNAMICS IN COMPLEX SYSTEMS - BACKGROUND AND LITERATURE REVIEW

In this thesis, a new, more comprehensive accident model created by Leveson [Leveson, 2004; Leveson, 2006] is extended with a dynamic modeling framework that enables the modeling and analysis of safety-related decision-making in complex socio-technical systems. The hypothesis for this research is:

System Theoretic Accident Models and Processes (STAMP) can be extended with a dynamic modeling framework to further improve current risk management techniques. The framework supports safety-related decision-making and assists with the design and testing of non-intuitive policies and processes to better mitigate risks and prevent time-dependent risk increase. Additionally, this new framework enables the identification of technical and organizational factors to detect and monitor states of increasing risk before an accident occurs.

This chapter provides a review of existing research in the area of safety and risk in complex engineering systems and provides the foundation upon which the rest of this work is based.

1.1 RISK IN COMPLEX ENGINEERING SYSTEMS

Risk takes many forms. Financial risk is associated with variability in returns on equity and cash flow resulting from financing. Investment risk is associated with variations in valuation and performance of individual businesses, sectors and the economy at large. Project risk is associated with the likelihood and consequence of not achieving objectives such as schedule, cost and performance. These are often called, respectively, schedule, cost and performance risks. Security risk refers to the risk of losses associated with external hostile intent. Occupational risk refers to health and safety hazards in the workplace. While the tools and methodologies developed in this thesis may be useful to mitigate many types of risk, the focus is on risks associated with the development and operation of complex safety-critical

engineering systems where undesired events during operation can lead to major losses. However, in most instances, safety cannot be extracted and treated in isolation from other programmatic concerns such as financial, schedule, and performance risks. That is, the various components of risk are not orthogonal, and financial, schedule and performance risks have an impact on system safety. As an example, programmatic and political concerns, among other things, prevented the inclusion of an escape system on the space shuttle that could have prevented the human losses associated with shuttle accidents [Rogers, 1986; McCurdy, 1994; Gehman, 2003].

1.1.1 DEFINITION

Risk is traditionally defined as a combination of the probability (or likelihood) and the consequence of a negative outcome or loss. Combining these components leads to the expected value of risk.

$$\text{Risk E(Loss)} = \text{Probability(Loss)} \times \text{Consequence(Loss)}$$

This simple formulation allows the calculation of expected losses associated with an event. For example, consider a lottery with one chance out of ten of losing ten dollars and another lottery with one chance out of a million of losing a million dollars. The rational expected loss, or risk, associated with each lottery is the same, at one dollar. A perfectly rational individual should be indifferent as to which lottery is chosen. However, it is likely that real-world players would have strong personal preferences for choosing one lottery over another (given that they necessarily have to choose one).

The expected value formulation is only useful in the most simple of cases. It assumes that the probability of a loss is perfectly known and that the consequence can be estimated. The problem with this simple formulation is somewhat analogous to Eisenberg's principle, that is, except for the simplest artificially created cases, it is impossible to know both the exact probability and the exact consequence of a real-world event. In financial risk analysis, the consequences of an event are usually well defined in monetary units, but the probabilities are not precisely known. In safety-critical systems, both the outcome (e.g. loss of human lives, damage to the environment) and the likelihood are usually impossible to estimate precisely.

One of the difficulties in estimating the outcome of loss events in safety-critical systems is the lack of objective units to quantify outcomes. Dollars may be appropriate as a measurement unit for financial losses, but may be limited in other areas such as systems with a potential for loss of life, permanent damage to the environment, or even damage to the reputation of a company and/or its products.

For safety-critical systems, the outcome of a loss event is often highly dependent on environmental conditions. For example, in the Bhopal case, where toxic chemicals were released in the atmosphere, the loss was exacerbated because of strong winds that carried the chemicals toward populated areas. Had the wind conditions been different, the accident, while still causing casualties, would not have been so deadly. While it is not always possible to exactly quantify the outcome of loss events, many different mitigation strategies exist to reduce the consequence of loss events. Some environment conditions can be partially controlled, for example, by locating nuclear power plants away from high-density population areas.

The second part of the equation, the likelihood of loss events, is usually much more difficult to estimate. While in some cases the outcome of a potential loss event is well known, there is often very high uncertainty associated with the likelihood of that event occurring. For example, if the space shuttle loses attitude control during supersonic descent, the outcome is clear: total loss of the vehicle and the crew. On the other hand, evaluating the likelihood of an attitude control loss is much less straightforward.

While much research is still based on the assumption that decision-makers are perfectly rational, especially in some economics and operations research fields, Herbert Simon [Simon, 1957; Simon, 1976] proposed an alternative to the “perfectly rational” paradigm for conducting research. His idea was that studying decision-making in isolation from the environment in which it is taking place, as well as without regard for the biases and limitations of decision-makers, would only provide mitigated or weak impact on real-life decision-making. Consequently, he proposed a research agenda that is consistent with the “bounded rationality” of decision-makers, which is related to the concept of “local rationality” and naturalistic decision-making used in the newest human factors and human-computer

interaction research [Rasmussen, 1994; Vicente, 1999]. The research presented in this thesis embraces the bounded/local rationality and naturalistic paradigms because accidents (and the human behaviors that contribute to their occurrence) cannot be explained or analyzed in isolation from the context in which they happen.

1.1.2 RISK IN SYSTEM SAFETY ENGINEERING

In system safety terminology, the definition of risk is extended to divide the likelihood of a negative outcome into the likelihood of a hazard occurring and the likelihood of that hazard leading to a loss or accident. The definition of risk becomes the hazard level (hazard severity and likelihood) combined with the likelihood of the hazard leading to an accident and the hazard exposure (See Figure 1 adapted from [Leveson, 1995]).

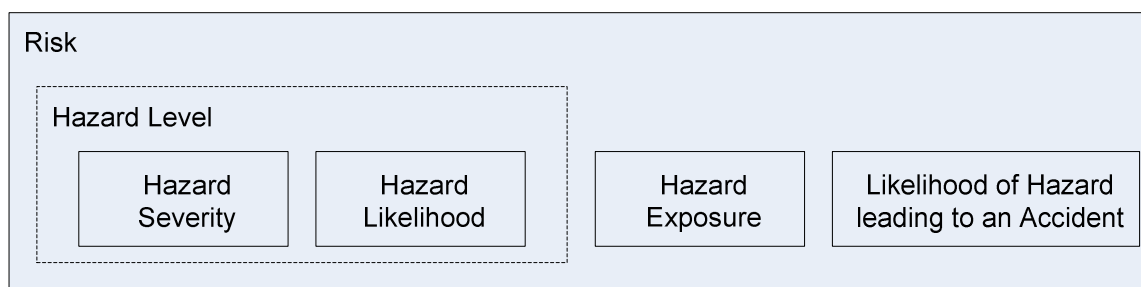


Figure 1: Risk Components (from [Leveson, 1995])

The hazard exposure or duration is a component of risk that accounts for the coincidence of conditions necessary for an accident to occur. The longer the hazard state exists, the greater chance of these conditions occurring. Even if there is a low probability of the individual conditions for occurrence, if the hazard exposure is long enough, the probability of occurrence can be dramatically increased. This concept of exposure is relevant to the concept of socio-technical systems migrating toward states of high risk. If the whole socio-technical system has migrated toward high risk as in the case of the Bhopal plant, the longer the system operates in this hazardous regime, the higher the likelihood of occurrence of the conditions necessary for an accident. Systems may operate in a hazardous state for a period of time without losses, but continuous operation in this state will eventually lead to a loss. Sometimes, a near-miss event or a threat to employee jobs and system survival will be sufficient to bring the system back to a lower risk level that will allow sustainable safe

operations. Carroll documented such a case [Carroll, 2002] at the Millstone nuclear power station in Connecticut, where threats to plant survival created external pressures, change in leadership, and self-generated change in safety culture among employees that contributed to a significant decrease in risk and allowed continued plant operation. This concept will be further discussed in Chapters 2 and 4.

1.1.3 ORIGINS OF RISK: UNCERTAINTY

The origins of risk stem from uncertainty in outcomes. An event that is certain to occur has no associated risk, but a certain definite outcome. If it were possible to predict exactly the future behavior of a system, risk would effectively disappear. There are many sources of uncertainty in complex systems. Every disciplinary area uses its own definition of uncertainty, but to study complex engineering systems, it is sufficient to divide uncertainty into four different types as proposed by Hastings [Hastings, 2004] and illustrated in Figure 2. Ambiguity refers to the imprecision associated with the terms and expressions used for human communication. Aleatory uncertainty is associated with the variations inherently associated with a physical system. Epistemic uncertainty is related to a lack of information about some characteristics of the system. Interaction uncertainty arises at the intersection between components of a system, or discipline areas, when the behavior of individual components or disciplines is well-understood, but the interactions between them are not. This taxonomy stresses that uncertainty arises from a lack of information about the future behavior of a system. Appendix B provides more detailed definitions and examples to illustrate each type of uncertainty.




Figure 2: Uncertainty Taxonomy for Engineering Systems

1.1.4 QUANTITATIVE RISK ASSESSMENT (QRA)

The objective of quantitative risk assessment is to assign quantitative figures to the likelihood and consequence of an identified risk. Most of the time, the focus will be on assessing the

probability of risk occurrence. Many different types of quantitative risk assessment methods are available, all with their strengths and weaknesses. While the author does not believe that quantitative risk assessment by itself is the best way to approach risk management in modern complex socio-technical systems, it can be effective to solve some well-defined problems in systems exhibiting strong characteristics of typically random failures and/or unorganized complexity [Weaver, 1958; Weinberg, 1975], where statistical analysis is an appropriate approach to investigate systems. Common QRA methods include Failure Modes and Effect Analysis (FMEA) and Failure Modes and Criticality Analysis (FMECA), actuarial approaches and Probabilistic Risk Assessment (PRA).

1.1.4.1 Failure Modes and Effect (and Criticality) Analysis

The objective of FMEAs and FMECAs is to examine quantitatively each potential component failure and decide which components should be the focus of reliability improvement efforts in order to “balance” risk as much as possible. FMEAs and FMECAs are reliability engineering tools and have their uses in that area. However, their effectiveness is very limited for systems with high degrees of redundancy or systems where common mode failure is a problem. Another limitation is that since FMEAs and FMECAs are bottom-up failure-based techniques, every possible identified failure mode is documented, regardless of consequence, which requires a large amount of documentation. FMEAs and FMECAs should be used to inform well chosen problems in the development and operations of complex systems, but like other bottom-up reliability approaches, they should not be confused with system-level analyses, and they are usually of limited value in examining and ensuring the safety of complex socio-technical systems.

1.1.4.2 Actuarial Approaches

Actuarial methods focus on extrapolating accident probabilities based on past operational history. This approach requires the availability of extensive operational data and experience. It is useful for nuclear reactors and commercial air transportation, where several millions of data points are available, but it has extremely limited usefulness in modern complex systems where data points are few and apart, or completely nonexistent. Even in areas with lots of

operational history, the introduction of new technologies or products may limit the usefulness of actuarial quantitative assessments and forecasts.

1.1.4.3 Probabilistic Risk Assessment (PRA)

Probabilistic risk assessment works by breaking a system down into subsystems and components, until a decomposition level is achieved where reliability data for the subsystem or component can be estimated. The reliability data is then re-aggregated, using a system model, such as a fault tree or event tree, to estimate the overall probability of accidents for the entire system. The main advantage of PRA is that it does not require extensive system-level operational experience, which is critical for most new engineering systems. An assumption underlying PRA analyses is that if reliability data is available for every component in a new system, then it is possible to achieve accurate reliability figures for the system-as-a-whole. PRA does not only provide estimates of system failure rates, but can also theoretically help modify designs and allocate resources more optimally to specific components in order to improve system-level reliability.

The advantages of PRA are undeniable, but there are many great difficulties associated with the validity and accuracy of resulting probability estimates. Uncertainty in component-level reliability figures is a very important limitation. Even small errors in the assessment of component failure probabilities can have a large impact when propagated to a system-level reliability estimate [Freudenburg, 1988]. Even if sufficient historical data is available for a specific component in a specific system, there is no proof that the reliability figures will translate to a new system where the component is used in a different way. Similarly, very slight changes to a component or its environment may result in a disproportionate change in reliability figures, so unless the environment and utilization are exactly similar, the numbers cannot be used with confidence. Assumptions about the use of components are rarely made explicit, so if components are used in a different design and a different operational environment, then assumptions will most likely be violated, thus invalidating reliability estimates, but this invalidation will not be noticed nor addressed. This inability to obtain accurate reliability figures is a very important limitation in using PRA for complex systems built with newly developed technologies and components and for software-intensive systems.

In fact, software is a textbook case of where the use of reliability figures breaks down. Not only is software inherently deterministic (in most cases), and not amenable to a probabilistic analysis, but slight changes in the software or its environment completely invalidate any sort of shaky “reliability” estimates that could have been obtained. For example, software modules that were created and successfully used for the Ariane 4 launcher were reused in Ariane 5 and were the direct cause of the Ariane 5 accident of 1996. The rationale was “*Unless proven necessary, it was not wise to make changes in software which worked well on Ariane 4 [Lions, 1996]*”, which proved to be immensely costly.

Other difficulties arise from the limitations of the models used as a basis for the aggregation of reliability figures. For example, the subjectivity associated with the “stopping rule” used to decide what will constitute an initiating event in a fault tree, as well as which branches will be included (or left out) undermines the validity of probability estimates. Other limitations include the difficulty to deal with human factors or organizational factors where probability estimates may not be obtainable. Additionally, PRA analyses do not typically consider system accidents resulting from dysfunctional interactions between components rather than random component failure. Finally, a detailed design must be available to obtain aggregate probability estimate because the only way to obtain accurate probability combination rules is to know how the components interact together in the first place. Once these design decisions are made, it may be too late to have a significant impact on the safety of the system. Further discussion of the benefits and limitations of PRA can be found in [Freudenburg, 1988; Bedford, 2001; Ayyub, 2003; Apostolakis, 2004; Marais, 2005].

1.2 SAFETY DEFINITION

The relation between risk and safety is defined differently depending on the application domain. As mentioned previously, risk can take many forms including project risk (cost, schedule, performance) as well as financial risk and risk to human lives. Safety is sometimes associated with risk to human lives only. In this thesis, safety is defined in absolute terms as the absence of losses due to an undesired event (usually an accident) [Leveson, 1995]. This definition thus takes an extended view of safety and includes losses such as human losses, mission or goal losses, equipment or material losses and environmental losses.

1.3 ACCIDENTS

To manage risk in complex engineering systems, it is necessary to understand how accidents happen. In order to do this, the use of an appropriate model of accident causation is critical. This section discusses the changing nature of accident in modern engineering systems and presents existing models and accident causation theories. Finally, a new accident causation model based on systems theory called STAMP [Leveson, 2004; Leveson, 2006] is described that will be used throughout the remainder of this thesis.

1.3.1 DEFINITION

Leveson defines an accident as a loss associated with an undesired or unplanned (but not necessarily unexpected) event [Leveson, 1995]. As such, a near-miss or incident does not fit the definition but could be defined as an undesired or unplanned event with no associated loss. Very often, the difference between an incident and an accident will only lie in different environment conditions. For example, consider a driver losing control of a car on an icy road. If there is no incoming traffic, the driver waits for the car to stop, then continues on his/her way. This is considered an incident. If there is incoming traffic and the car gets hit as a result of the loss of control, an accident has occurred.

1.3.2 ACCIDENTS IN COMPLEX SOCIO-TECHNICAL SYSTEMS

As the complexity of engineered systems increases, new types of accidents have started to emerge that result from dysfunctional interactions between system components [Leveson, 2004]. These accidents result from unplanned or unexpected interactions between different components of a system, rather than single (or multiple) component failure. For example, the loss of the Mars Polar Lander occurred because the designers did not take into account a particular interaction between the thruster's software controller and the mechanical leg deployment. When the legs deployed, a spurious signal was interpreted by the controller as a sign that the lander had reached the Martian surface. The controller shut down the thrusters while the lander was still 50 feet above the ground, causing the spacecraft to crash into the surface [Albee, 2000; Leveson, 2001; Leveson, 2004; Leveson, 2004].

Perrow defines system accidents as resulting from cascading failures [Perrow, 1999]. However, many accidents such as that of the Mars Polar Lander do not result from traditional failures. Digital systems and software introduce new types of accident causation where no component “fails” in the traditional sense of the word. In addition, while some level of automation is required to handle most complex systems, trouble often occurs at the interface between human operators and automated systems, when there are overlaps and/or conflicts in responsibilities and when human operators do not fully understand what the automation is doing. Accidents resulting from dysfunctional interactions between humans and system automation are often blamed on human error. However, automated systems are often poorly designed and the systems are so coupled and opaque that it is ludicrous to expect that operators will be able to remain in control when the situation turns sour.

The trend to blame accidents on human error is continuing. There are incentives for corporations to blame individuals in order to avoid or deflect possible suits and public anger. In fact, the less is known about the specific circumstances of an accident, the more likely it is to be attributed to human error [Johnson, 1980]. Perrow also argues that “human error” is a convenient classification for accidents whose real cause is uncertain, complex, or embarrassing to the organization [Perrow, 1983; Perrow, 1999]. In reality, major accidents are hardly ever caused by pure human error. All human activity takes place within a physical and social context that shapes behavior and it is almost always possible in hindsight to find deficiencies or deviations from prescribed behavior. After the fact, a more optimal string of decisions can be easily devised that would have allowed avoiding an accident. In their analysis of the Zeebrugge ferry accident, Rasmussen and Svedung [Rasmussen, 2000; Rasmussen, 2002] noted that those making decisions about vessel and harbor design, cargo and passenger management, scheduling and operation were unaware of the impact of their decisions on other decision-makers and on the system-level process leading to the accident. Rasmussen [Rasmussen, 1997] stresses that most decisions are sound given local judgment criteria (they are “locally” rational) and given the time and budget pressures and short-term contextual incentives that shape behavior. Each individual decision may appear safe and rational within the context of the individual work environments and local pressures, but when taken into the context of the entire system operation, these decisions and actions may interact in unexpected ways to produce an accident. Accidents in complex socio-technical systems

are often caused not by unknown variables or failure modes, but by flawed decision-making that may be the result of excessive performance pressure or poor safety culture. As such, to perform effective risk management in complex systems, it is necessary to use a more inclusive approach that encompasses the technical aspects of a system, as well as the managerial, organizational, social, and political aspects of the system and its environment.

As mentioned previously, systems evolve in order to accomplish changing objectives and adapt to environmental pressures and disturbances. Often times, accidents in complex systems involve the migration of the system toward an unsafe or unstable state where small deviations can cascade into catastrophes [Leveson, 1995; Rasmussen, 1997; Leveson, 2004]. Using a parallel to Turner's disaster incubation model [Turner, 1978], the foundation for an accident is often laid years before. Once the system has reached an unsafe state, a single event may trigger the loss, but if one particular event does not occur, another one eventually will [Leveson, 2006]. Preventing accidents in complex systems involves ensuring that risk remains at sustainable levels throughout the system lifecycle.

1.3.3 TRADITIONAL ACCIDENT MODELS

Accident models serve two major purposes: they are used to understand past accidents and to prevent future ones. Many different accident models have been used with various degrees of success. The common assumption underlying most traditional accident models is that accidents can be explained (or at least partially explained) as a sequence of events directly connected through a forward chain in time [Leveson, 2006]. Another related and necessary assumption behind such models is that there is one initial cause responsible for triggering the event chain. This section presents a short overview of variations on this chain-of-events model. The main limitations are discussed in order to understand the challenges that must be addressed to prevent accidents in complex engineering systems and to provide a stepping stone to the use of a more complete and comprehensive accident causation model. For further discussion of commonly used accident models, the reader is referred to [Leveson, 1995; Leveson, 2004; Leveson, 2006].

1.3.3.1 Event-Chain Model

An event-chain model describes accidents as the last event in a chain that includes multiple events or causal factors connected together through direct, sequential relationships. The assumption behind these models is that if the chain can be broken in any way, the accident can be prevented. This assumption implies that the relationships between events are direct and that a preceding event or condition is necessary for the following event to occur. Consequently, preventing accidents modeled using an event chain involves trying to break chains by eliminating some events or by intervening somewhere along a chain.

Chains of events need not be single strands. The chains may be branching and there may be multiple parallel chains synchronized in time or through common events. These branching chains may include logical conditions such as AND and OR that define relations between events. The selection of events included in the chain is rather arbitrary and heavily dependent on the domain knowledge of the analyst. The accident “cause” depends on where the chain is stopped, or what stopping rule is used, which may vary depending on industry standards and/or analyst preferences. The first event in a chain is called the root cause or initiating event. It is often convenient to start the chain using generic events such as “human error” or “software failure”, but the explanatory and prevention potential of these terms is very limited. In an analysis by Leveson [Leveson, 2001] of recent aerospace accidents, most of the reports stopped after assigning blame (usually to the operators) and never got to the reasons why the operators made the errors they did and how to prevent such errors in the future or why the software requirements error was made and why it was not detected and fixed before the software was used.

Many engineering techniques based on the event chain causation model have been created to help prevent accidents. Event chain analysis can be based on a temporal ordering of events or on the part-whole decomposition of hazardous states.

Temporal-based searches can proceed forward in time from an initiating event to a loss event or backwards in time from a loss event to initiating event(s). Forward chain methods such as Failure Modes and Effect Analyses (FMEAs) and event trees start from different failure modes of individual components and propagate them forward through the chain to see how

they could affect the operation of the system, whether the purpose is hazard analysis (event trees) or reliability analysis (FMEAs). As mentioned earlier, probabilities are often associated with individual component failure modes, which allows analysts to combine probabilities and obtain a probabilistic estimate of a system-level failure or unavailability mode. As discussed in section 1.1.4, there are many difficulties associated with these probability estimates.

Part-whole decomposition searches can proceed top-down (e.g. Fault Tree Analysis) where a hazardous state is decomposed from the loss event or associated system hazard at the top, and refine the chain, eventually leading to basic fault events. Analysis can also proceed bottom up, where arbitrarily chosen basic fault events combine to cause a loss event or associated hazardous condition. As discussed previously, quantitative analyses can also be performed using fault trees, given that probabilities are known for the occurrence of each basic event.

1.3.3.2 Other Variations of the Event-Chain Model

Many other variations exist that are still based on the general event-chain causation model. The Reason model uses a Swiss cheese metaphor to explain accident causation. Reason states that it is “a general model that traces the root causes of different accidents to organizational errors (latent failures) arising in the upper levels of any organization” [Reason, 1995]. The model is rather simple and easy to apply and it has been rapidly adopted in the aviation industry. The model has been used extensively by the Federal Aviation Administration (FAA) to investigate the role of management policies and procedures in aircraft accidents and incidents. According to this model, accidents occur when “holes in multiple layers of Swiss cheese align”. In reality, if the “holes” are viewed as events, the Swiss cheese model really is simply a way to visualize a chain of event. The Swiss cheese metaphor is interesting as a descriptive method, but is overly simplistic by assuming a single cause per “cheese slice” and does not provide much prescriptive power to improve safety management practices beyond ad hoc policies based on observations in different systems [Pidgeon, 1998].

Other variations include the Cause-Consequence Analysis (CCA), initially developed to perform quantitative risk analysis for the Danish Atomic Energy Commission [Nielsen, 1971]. Cause-Consequence Analyses combine deductive and inductive analysis; top-down searches (in the form of a fault tree) are used to perform cause analysis and are combined with forward

searches (in the form of event chain) to perform consequence analysis. Another event-based method is the Management Oversight Risk Tree (MORT) that was also developed in the early 1970s by Johnson [Johnson, 1973; Johnson, 1980] for the U.S. Atomic Energy Commission. The MORT analysis is a fault tree-type analysis, but the MORT framework provides a checklist-type repository of over 1500 basic events combined into 100 generic events coming from the fields of workplace accident prevention, management functions, human behavior and environmental factors [Suokas, 1993; Leveson, 1995].

1.3.3.3 Limitations of the Traditional Even-Chain Model

Accident prevention techniques based on the even-chain model such as FMEAs and FTAs have been around for more than 40 years. They were first invented for mechanical systems where component failures usually occur randomly and where redundancy is very effective at reducing the likelihood of certain types of failure modes. Later, these techniques were extended to electro-mechanical systems that included simple analog devices such as relays, motors and electrical hardware that also had a propensity to fail in a random way. With the advent of digital systems and software, the complexity of the systems being built increased dramatically, and it became difficult to model the relations between system components in a direct way. Nevertheless, since the component failure based methods were previously successful, they continued to be used widely and either ignored the indirect relationships and the impact of software on system safety or assumed all relationships were direct and assigned a probability distribution for software failure, even though software does not fail in a random way. In fact, software does not “fail” per se. It just does what it was programmed to do, which in some cases may contribute to an accident if the requirements were wrong in the first place or if there were implementation errors. The “software failure” box still found in many fault trees is a sign that traditional methods have reached their efficacy limits and that a new paradigm is needed.

In addition to difficulties in handling software-intensive systems, traditional methods are also limited in understanding the contribution of “softer” organizational factors affecting system safety such as management pressures, limited resources and independence of safety decision-making. Event-based methods are sometimes extended to include organizational factors. For example, researchers in the PRA field [Paté-Cornell, 1990; Paté-Cornell, 1996] have included

some human and management factors into their methods for determining system failure probabilities. However, these attempts, apart from being based on difficult to validate assumptions about the mapping from factor to event and omitting indirect relationships, are very static in nature, and thus will capture neither the time-dependent nature of risk mitigation nor the possible slow migration of systems toward a state of increasing risk.

1.4 ORGANIZATIONAL RISK THEORIES

It is increasingly well accepted that organizational factors play a role in almost all accidents and are a critical part of understanding and preventing them. Two prominent sociological schools of thought have focused their attention on the organizational aspects of safety: Normal Accident Theory (NAT) [Sagan, 1993; Perrow, 1999] and High Reliability Organizations (HRO) [Rochlin, 1987; Weick, 1987; Roberts, 1990; Roberts, 1990; La Porte, 1991; Rochlin, 1991; Weick, 1993; La Porte, 1996; Weick, 1999]. There has been an ongoing debate as to which theory or school of thought dominates. The purpose of this section is not to take a side or to provide an exhaustive review of the topic but rather to summarize the two approaches and discuss strengths and limitations as a stepping stone toward a more holistic system theoretic approach to accident modeling. Appendix C provides a summary of the strengths and limitations for each, as well as a discussion on the ongoing debate between the two approaches. For an even more extensive comparison and analysis, the readers are referred to [Sagan, 1993; Rijpma, 1997; Rijpma, 2003; Marais, 2004; Marais, 2005].

1.4.1 NORMAL ACCIDENT THEORY (NAT)

Charles Perrow's initial formulation of what has come to be known as Normal Accident Theory (NAT) was developed in the aftermath of the accident at the Three Mile Island nuclear power plant in 1979 [Perrow, 1982]. Perrow introduced the idea that in some technological systems, accidents are inevitable or "normal" [Perrow, 1999]. He defines two related dimensions: interactive complexity and tight coupling, which determine a system's susceptibility to accidents.

Interactive complexity refers to the presence of unfamiliar or unplanned and unexpected sequences of events in a system that are either not visible or not immediately comprehensible.

A tightly coupled system is one that is highly interdependent: Each part of the system is tightly linked to many other parts and therefore a change in one part can rapidly affect the status of other parts. Tightly coupled systems respond quickly to perturbations, but the response may be disastrous. Loosely coupled or decoupled systems have fewer or less tight links between parts and therefore have more capacity to absorb failures or unplanned behavior without major destabilization.

According to NAT, systems that are interactively complex and tightly coupled will experience accidents that cannot be foreseen or prevented. Perrow calls these system accidents. When the system is interactively complex, independent failure events can interact in ways that cannot be predicted by the designers and operators of the system. If the system is also tightly coupled, the cascading of effects can quickly spiral out of control before operators are able to understand the situation and perform appropriate corrective actions. In such systems, apparently trivial incidents can cascade in unpredictable ways and with possibly severe consequences.

1.4.2 HIGH RELIABILITY ORGANIZATIONS (HROs)

High Reliability HROs are defined by Roberts [Roberts, 1990] as the subset of hazardous organizations that enjoy a record of high safety over long periods of time:

“One can identify this subset by answering the question, ‘how many times could this organization have failed resulting in catastrophic consequences that it did not?’ If the answer is on the order of tens of thousands of times, the organization is ‘high’ reliability. [Roberts, 1990].”

The field of High Reliability Organizations research is based on observations made during the study of two aircraft carriers, U.S. air traffic control, utility grid management, a nuclear power plant and fire fighting teams [La Porte, 1991]. These observations seem to counter Perrow’s hypothesis by suggesting that some interactively complex and tightly coupled systems operate for long periods of time with very few accidents.

The literature associated with the HRO field is large and growing. Nevertheless, most HRO researchers agree on four primary organizational characteristics that they claim substantially limit accidents and simultaneously result in high levels of performance: (1) prioritization of

both safety and performance and consensus about the goals across the organization [La Porte, 1991]; (2) promotion of a “culture of reliability” in simultaneously decentralized and centralized operations [Weick, 1987]; (3) use of organizational learning that maximizes learning from accidents, incidents, and near misses [La Porte, 1991]; and (4) extensive use of redundancy [Rochlin, 1987]. Much of the recent HRO research focuses on applying these principles to various systems and/or attempting to correlate the application of these principles with organizational performance characteristics such as reliability [Roberts, 2005]

1.4.3 LIMITATIONS OF ORGANIZATIONAL RISK THEORIES

Organizational sociologists in general have made an important contribution to system safety by emphasizing the organizational aspects of accidents. At the same time, they have underemphasized or oversimplified the engineering parts, for example, focusing only on simple redundancy, not considering accidents where component failure is not the cause, or studying only systems that are relatively simple and loosely coupled and then drawing conclusions from them to apply to all systems. Complex, socio-technical systems need more sophisticated approaches to increasing reliability and safety for the non-random, technical, and organizational factors involved in accidents.

1.5 A SYSTEMS THEORETIC APPROACH TO SAFETY AND ACCIDENT MODELING

A group of researchers, including Rasmussen [Rasmussen, 1997], Hollnagel [Hollnagel, 2002], Woods [Woods, 2002], and Leveson [Leveson, 2004], most of whom come from a systems engineering and human factors background, have been advocating an alternative, systems approach to safety. The primary differences between a systems approach and the HRO and standard engineering approaches are: (1) top-down systems thinking rather than a bottom-up, reliability engineering focus and (2) a focus on the integrated socio-technical system as a whole and the relationships between the technical, organizational, and social aspects.

It is critical to recognize the difference between reliability and safety. HRO researchers talk about a “culture of reliability” where it is assumed that if each person and component in the

system operates reliably, there will be no accidents. Even Perrow seems to assume that accidents require failures of components. This assumption is not accurate. In complex systems, accidents often result from interaction among perfectly functioning components. The loss of the Mars Polar Lander was attributed to noise (spurious signals) generated when the landing legs were deployed during descent [Albee, 2000]. The onboard software interpreted these signals as an indication that landing occurred and shut the engines down prematurely, causing the spacecraft to crash into the Mars surface. The landing legs and the software performed correctly, but the accident occurred because designers failed to account for all interactions between the leg deployment and the software descent engine control software.

Highly reliable systems are not necessarily safe and highly safe systems are not necessarily reliable. Reliability and safety are different qualities and should not be confused. In fact, these two qualities often conflict. Increasing reliability may decrease safety and increasing safety may decrease reliability. One of the challenges of engineering is to find ways to increase safety without decreasing reliability. For example, some ways to reduce the accident rate on aircraft carriers would be to slow down the landing rates, only allow landing in the most perfect weather and the most ideal conditions, and only allow the most experienced pilots to make the landings. These operational conditions would most likely conflict with the achievement of other goals, such as training for combat.

Reliability in engineering is defined as the probability that a component satisfies its specified behavioral requirements over time and under given conditions. If a human operator does not follow the specified procedures, then they are not operating reliably. In some cases that can lead to an accident. In other cases, it may prevent an accident when the specified procedures turn out to be unsafe under the particular circumstances. Examples abound of operators ignoring prescribed procedures in order to prevent an accident [Leveson, 1995]. At the same time, accidents have resulted precisely because the operators did follow the predetermined instructions provided to them in their training. When the results of deviating from procedures are positive, operators are lauded but when the results are negative, they are punished for being unreliable. HRO researchers [Weick, 1987; Roberts, 1990; Roberts, 1990; La Porte, 1991; Schulman, 1993; Weick, 1993; Weick, 1999] correctly point out the need for operators to sometimes break the rules in order to prevent an accident, but incorrectly label their

behavior as reliable. Such behavior is in fact not reliable with respect to following the specified rules or training; it is unreliable but safe. The distinction becomes extremely important when multiple, perhaps conflicting, goals are involved. If the goal is to increase safety, then we should be talking about enhancing the safety culture, not the reliability culture. The safety culture is that part of organizational culture that reflects the general attitude and approaches to safety and risk management. Aircraft carriers do have a very strong safety culture and many of the aspects of this culture observed by the HRO researchers can and should be copied by other organizations, but labeling these characteristics as “reliability” is misleading and can lead to misunderstanding what is needed to increase safety in complex, tightly coupled systems.

Leveson defines an alternative, engineering-centric approach to organizational safety theories in [Leveson, 1995]:

“A systems approach to safety recognizes that safety is a property of the system as a whole, not a property of individual system components: The socio-technical system must be treated as an integrated whole using a top-down rather than a bottom-up perspective. This fact, in turn, implies that effectively tackling safety problems will require researchers and practitioners to step outside their traditional boundaries and take a broad view of the problems. Systems theory dates from the thirties and forties and was a response to the limitations of classic analysis techniques in coping with the increasingly complex systems being built [Checkland, 1981]. The systems approach assumes that some properties of systems can only be treated adequately in their entirety, taking into account all facets and relating the social to the technical aspects [Ramo, 1973]. These system properties derive from the relationships between the parts of systems: how the parts interact and fit together [Ackoff, 1971]. Thus, the systems approach concentrates on the analysis and design of the whole as distinct from the parts. The use of a systems approach creates the possibility of modeling and engineering the safety culture and organizational aspects of safety, including the entire socio-technical system.”

Some modeling techniques have been proposed as the foundation upon which dynamic socio-technical risk modeling can be performed. Those models and techniques will be discussed in the next chapter.

1.6 DEFINITIONS AND TERMINOLOGY

Many different terms will be used throughout this thesis to address risk and safety. Usually, detailed definitions will be provided along the way. But as a starting point for further discussion and analysis, this section provides baseline definitions of current terms and concepts used for safety and risk management in complex systems.

1.6.1 ACCIDENT

An accident is *an undesired and unplanned (but not necessarily unexpected) event that results in a specified level of loss [Leveson, 1995].*

1.6.2 INCIDENT

A near miss or incident is *an event that involves no loss (or only minor loss) but with the potential for loss under different circumstances [Leveson, 1995].*

1.6.3 SAFETY

Safety is defined as *freedom from accidents (losses) [Leveson, 1995].* Safety is defined absolutely as a quality that may not be entirely achievable, but that can still be defined in absolute terms as a desirable quality that can be improved.

1.6.4 RELIABILITY

Reliability is *the probability that a piece of equipment or component will perform its intended function satisfactorily for a prescribed time and under stipulated environmental conditions [Leveson, 1995].* Unreliability is the probability of failure.

1.6.5 FAILURE

Failure is *the nonperformance or inability of the system or component to perform its intended function for a specified time under specified environmental conditions [Leveson, 1995].*

1.6.6 HAZARD

A hazard is a state or set of conditions of a system (or an object) that, together with other conditions in the environment of the system (or object), will lead inevitably to an accident (loss event) [Leveson, 1995].

Hazards are a key concept in system safety engineering, and have been defined in many different ways. Some define hazards as an event, say an explosion, but for various reasons, it is customary in system safety practice is to define hazards as system states, and always in relation to the environment of the system or component. Moreover, what constitutes a hazard is dependent upon where the system boundaries are drawn. As with solving other engineering problems, the location of the system boundaries is arbitrary, but has a large impact on the problem resolution process. As with other problems (and problem-solving methodologies), the most important thing is consistency, but wisely choosing boundaries will facilitate the problem resolution and hazard analysis. As a rule of thumb, for hazard analysis, the system boundaries should be large enough to include the conditions related to an accident over which the system designer has some control.

In the system safety field, a hazard has two important characteristics: (1) *Severity*, and (2) *Likelihood* of occurrence. Hazard severity is defined as the worst possible damage that could result from the hazard given the most unfavorable environment conditions. Hazard likelihood of occurrence can be defined quantitatively or qualitatively. In some instances, usually for hazards associated with standard systems with abundant historical data, it may be possible to define likelihood using a quantitative probability of occurrence. However, for most complex socio-technical systems with little or no historical experience, which is mostly what this thesis is about, a qualitative assessment of likelihood is the best that can be achieved. The combination of hazard severity and likelihood is often called the hazard *level*.

1.6.7 HAZARD ANALYSIS

The identification of hazards and the assessment of hazard levels [Leveson, 1995].

1.6.8 HAZARD MITIGATION

The reduction of hazard levels through the use of various system design or operational mitigation strategies.

1.6.9 RISK

The definition used in this thesis is a system safety definition of risk:

Risk is the hazard level combined with the likelihood of the hazard leading to an accident (sometimes called danger) and hazard exposure or duration (sometimes called latency) [Leveson, 1995].

Exposure or duration of hazard is a component of risk. A hazard is only one of the conditions necessary for an accident to occur, but the longer the hazardous state exists, the greater the chance of the other required conditions occurring during the hazard exposure time. Figure 3 provides a summary of the many components of risk.

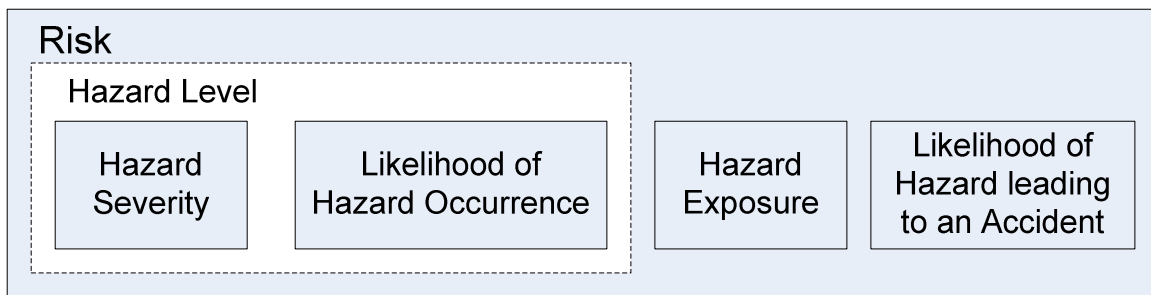


Figure 3: Components of risk (Adapted from Leveson)

1.6.10 RISK ANALYSIS

The identification and assessment of risks.

1.6.11 RISK MANAGEMENT

The process of maintaining risk at an acceptable level throughout the lifecycle of a system.

1.7 THESIS SCOPE AND BOUNDARIES

The focus of this thesis is on risk management in complex socio-technical systems. The objective is to create a framework to identify and model the factors that contribute to risk in the development and operation of complex engineering systems.

Just as solving an engineering or system safety problem requires the definition of system boundaries, writing a dissertation requires the definition of the problem scope, as well as the boundaries of the systems and factors to be included in the tentative problem solution. These boundaries can also be arbitrary, but must be large enough to include the factors that contribute to the problem at hand, without being so large as to waste resources on unimportant factors.

Most of the techniques upon which this work is based are derived from system safety engineering, system theory, control theory, and system dynamics. The definition of safety used throughout this thesis includes not only risks associated with human life, but also risks associated with mission failure, equipment loss and environmental damage. In fact, for the development of complex systems, safety and mission assurance cannot be decoupled from programmatic risks that include budget, schedule and performance risks, so these concerns are a critical part of modeling risk in the development of new systems.

The terms complex socio-technical systems and complex engineering systems will be used interchangeably in this thesis. However, there is a semantic difference between the terms. A thorough philosophical discussion of the characteristics of different systems is beyond the scope of this thesis. However, it should be noted that complex socio-technical systems must fulfill three main conditions: (1) They must have a large human-designed component, thus the technical term, (2) they must exhibit technical complexity, and (3) they must exhibit human/social/organizational complexity. Technical and social complexity is not a discrete characteristic, but can be defined along a continuum which ranges from very simple to extremely complex. Moreover, complexity is relative and a function of current intellectual manageability, which is evolving as new tools and techniques are developed [Leveson, 2000]. Consequently, it is extremely challenging to measure the level of complexity of different systems. However, it should be obvious for the readers that the examples used in this thesis

belong to the set of complex socio-technical systems. The term complex engineering system would not necessarily include high social complexity, but the working definition used in this thesis can be found among a larger complex systems taxonomy in [Magee, 2004]:

Engineering System: A system designed by humans having some purpose; large scale and complex engineering systems will have a management or social dimension as well as a technical one.

Complex System: A system with numerous components and interconnections, interactions or interdependencies that are difficult to describe, understand, predict, manage, design, and/or change.

In the tradition of systems theory research, in addition to defining the problem and system boundaries, it is necessary to define the level of abstraction at which most of the work will be performed. The risk modeling and management framework introduced in this thesis is based on the STAMP accident model. The models created include any component of the socio-technical system that contributes to risk management and/or safety-related decision-making. The components range from high-level congressional dynamics that influence the system, to the detailed technical design decisions that affect the physical characteristics of the system [Leveson, 2003; Dulac, 2004; Leveson, 2004; Leveson, 2004; Leveson, 2006]. However, for this thesis, the main focus is not on the technical details of the system, or even on individual human behavior, but rather on the organization-level interactions, processes and factors that influence safety and risk decision-making in organizations. The technical risk associated with individual system hazards is critical for a complete risk analysis; however, the focus of this thesis is on the organizational decision-making influenced by systemic factors that impact system safety.

1.7.1 THESIS OBJECTIVE

The goal of this dissertation is to lay down the foundations for dynamic modeling of risk in the development and operation of complex engineering systems.

More specifically, the main objective is to *develop and operate safer engineering systems by providing more powerful risk management techniques through the extension of the STAMP*

accident model with a dynamic modeling framework that supports safety-related decision-making. The new framework will assist with the design and testing of non-intuitive policies and processes to better mitigate risks and prevent time-dependent risk increase. Additionally, the framework enables the identification of technical and organizational factors to detect and monitor states of increasing risk before an accident occurs.

The objective is achieved by: 1) Explaining how the STAMP model and System Dynamics can be combined to create a more comprehensive risk management modeling framework, 2) Introducing a model-building methodology to create and validate custom dynamic risk management models, 3) Providing components (or templates) that can be assembled and customized to facilitate model-building, 4) Providing a toolset of analysis methods to assist in decision-making, analyze risks, test and create risk-mitigation policies and control actions, and monitor the system for potential risk increase, and finally 5) Demonstrating the usefulness of the entire methodology through real-life projects aimed at developing and operating safer complex systems.

1.7.1.1 Development Modeling vs. Operations Modeling

The safety and risk dynamics involved in developing a new system are very different than that involved in the operation of an existing system, even though there may be overlap in the time spans of the two activities when systems are partially deployed, or as systems in operation go through refurbishment or upgrades and evolution cycles. Nevertheless, for the purpose of this thesis, the development and operation phases, while influencing each other, will be considered as two different problems.

1.7.2 THESIS OUTLINE

The dissertation goes through a natural progression, from background to high-level dynamic patterns in complex systems development and operation, heavily rooted in past accident experience and current literature in system development and safety/risk management. The STAMP-based risk management process is reviewed, and the last steps of the process that include dynamic risk modeling and analysis are discussed including a detailed example of the model-building methodology and analysis using a real system.

More specifically, Chapter 1 provided background and literature review on safety and risk management in complex socio-technical systems. Chapter 2 follows by presenting two main theoretical foundations upon which this work builds, namely, the STAMP (System Theoretic Accident Model and Process) accident model and system dynamics, both heavily influenced by system and control theory. Chapter 2 continues by introducing and analyzing feedback loops and resulting dynamic patterns that impact the development and operation of complex socio-technical systems. Chapter 3 introduces guidelines and criteria for the creation, analysis and monitoring of static STAMP socio-technical safety/risk control structures. Chapter 4 defines the core of a methodology for creating dynamic risk management models using generic customizable dynamic components. Chapter 5 discusses the use of the newly introduced dynamic risk models for risk analysis and management in complex systems. Chapter 6 presents a case study of the methodology in action, i.e., how the risk models were created and used to inform the risk management decision-making at NASA's Exploration Systems Mission Directorate (ESMD) during the development of NASA's new space exploration system. Chapter 7 wraps up with a summary of conclusions, contributions, challenges, and future work.

CHAPTER 2: LIFECYCLE RISK MANAGEMENT MODELING FOR COMPLEX SOCIO-TECHNICAL SYSTEMS

This chapter defines some dynamic foundation of system safety during system development and operation. The chapter is divided into four major parts. The first two parts provide a review of the two major theoretical foundations upon which this work builds, namely: 1) the STAMP (System Theoretic Accident Model and Processes) accident model developed by Leveson [Leveson, 2004; Leveson, 2006] and 2) System Dynamics, which has a long history as a tool for addressing various problems of a dynamic, time-dependent nature [Forrester, 1961; Meadows, 1972; Sterman, 2000].

The third part will focus on the dynamic structures responsible for the integration of safety into the design and implementation of a new system, that is, “Safety-Centric System Development”. While safety is the main focus of this exercise, it cannot be separated from traditional risk management areas such as system development cost, schedule, and system performance. Thus the goal of the program and project manager is to develop a system that will meet the performance and safety requirements, while ensuring that the product will be developed on time and on budget.

The fourth part of this chapter focuses on safety dynamics during the operation of an existing system. As mentioned previously, accidents often occur following a slow migration of the entire system toward a state of high-risk. In other words, slow changes and erosion in the safety control structure and processes result in a failure to ensure safe system operation during the entire system lifecycle. In the last part of this chapter, we are attempting to capture the main dynamic feedback mechanisms responsible for this migration of the control structure toward states of higher risk where violations of safety constraints are likely to occur. Additionally, Appendix D is associated with this chapter and provides an example of the impact of balancing and reinforcing feedback loops on the dynamics of a real system, the NASA Independent Technical Authority.

The purpose of chapter 2 is to introduce and discuss some of the critical feedback loops responsible for creating dynamic patterns likely to occur during the development and operation of complex socio-technical systems. The emphasis of this discussion is on the feedback mechanisms themselves and the resulting behavioral patterns they create. The specific actors and system components responsible for the creation of these dynamic feedback effects will be discussed when we introduce the component-based model development methodology in later chapters.

2.1 THE STAMP MODEL OF ACCIDENT CAUSATION

Traditionally in engineering, accidents have been viewed as resulting from a chain of failure events, each directly related to its “causal” event or events. The event(s) at the beginning of the chain is labeled the *root cause*. Almost all hazard analysis or risk assessment techniques, such as failure modes and effect analysis (FMEA), fault tree analysis (FTA), and probabilistic risk analysis (PRA), use this chain-of-events paradigm.

Event-based techniques were created in an era of mechanical systems and then adapted for electro-mechanical systems. The assumptions underlying these systems do not fit the complex, software-intensive systems built today, which often involve complex human-machine interactions, systems-of-systems with distributed decision-making that cut across both physical and organizational boundaries. In these new, more complex systems, a new type of accident, *system accidents*, start to appear where the components function as designed (i.e., do not fail) but problems and accidents still arise because of dynamic interactions among components and systems.

Traditional event-based techniques do not support the complex human decision-making required to develop and operate modern automated large-scale engineering systems. Ensuring safety in these systems involve understanding the technical aspects of the system, but also the organizational and social aspects of safety and safety culture. Technical engineering decisions and organizational decisions are intimately related; good engineering decisions can be invalidated by poor management decisions.

Traditional event-based models using direct, linear causality must be augmented to handle the complex, indirect, and non-linear interactions of complex systems-of-systems. Leveson [Leveson, 2004; Leveson, 2006] developed a new accident model called STAMP (Systems-Theoretic Accident Modeling and Processes) that can handle these complex interactions in addition to the traditional direct causality interactions. STAMP is not based on chain-of-events. Instead, it uses a general notion that accidents result from inadequate enforcement of safety constraints in design, development, and operation [Leveson, 2004]. STAMP includes traditional failure-based models as a subset, but goes beyond physical failures to include causal factors involving dysfunctional interactions among non-failing components, software and logic design errors, errors in complex human decision-making, and various organizational characteristics such as workforce, safety processes and standards, contracting and procurement, and flaws in the safety culture.

STAMP is based on systems theory and draws on basic concepts from engineering, mathematics, cognitive and social psychology, organizational theory, political science, and economics. Leveson explains the systems theory foundation of STAMP in [Leveson, 2006]:

“Systems theory was developed after World War II to cope with the vastly increased complexity of the systems, particularly military systems, starting to be built at that time [Ashby, 1956; Von Bertalanffy, 1968; Checkland, 1981]. In systems theory, systems are viewed as interrelated components that are kept in a state of dynamic equilibrium by feedback loops of information and control. Systems are not treated as a static design but as dynamic processes that are continually adapting to achieve their ends and to react to changes in themselves and their environment. To be safe, the original design must not only enforce appropriate constraints on behavior to ensure safe operation (i.e., to enforce the system safety constraints), but the system must continue to operate safely as changes and adaptations occur over time to meet a complex set of goals and values under changing social and technical conditions.”

When using the new STAMP accident model, safety is treated as a control problem. Accidents occur when component failures, external disturbances, and/or dysfunctional interactions among system components are not adequately handled by the controller and result in the system heading to a hazard state. In other words, accidents result from inadequate control or enforcement of safety-related constraints during system development and operation.

Accidents arise because risk is not adequately managed by the socio-technical system (including the technical, social, managerial, organizational, and political system components) during design, testing, manufacturing and operation. For example, the Columbia accident involved inadequate controls on the physical process, as foam shedding commonly caused damage to the orbiter, in clear violation of original system requirements. In addition, controls were inadequate on the management launch decision processes and orbiter integrity monitoring, both during launch and on orbit, as many attempts at monitoring orbiter status were hindered, and even actively cancelled and prevented by mission management [Gehman, 2003].

Events and event chains are the result of inadequate control. These are the visible symptoms of dysfunctional interactions and inadequate enforcement of safety constraints. Conversely, inadequate control is only indirectly reflected by the events. Consequently, the socio-technical control structure that enforces safety constraints must be carefully designed, evaluated, and monitored to ensure that the controls are adequate to maintain the constraints on behavior necessary to ensure safety. This definition of safety and risk management is broader than definitions that specify particular activities or tools [Leveson, 2006]. STAMP, being based on systems and control theory, provides the theoretical foundation to develop the techniques and tools, including dynamic modeling tools, to assist analysts, designers and managers in managing safety risk in this broader context.

Control does not imply a strict military-type command and control structure. Control can occur through direct authoritarian intervention, but also indirectly through policies, procedures, standards, oversight, norms, shared values and other aspects of the organizational (and safety) culture. Behavior is influenced and at least partially “controlled” by the social and organizational context and norms in which the behavior occurs. Engineering this context can be an effective way to create and change a safety culture [Leveson, 2006]. STAMP has three basic concepts: safety constraints, hierarchical safety control structures, and process models [Leveson, 2006].

2.1.1 SAFETY CONSTRAINTS

The most basic concept in STAMP is not an event but a constraint [Leveson, 2006]. In systems theory, systems are viewed as hierarchical control structures where each level imposes constraints on the activity of the level beneath it [Checkland, 1981]. Constraints or lack of constraints at a higher level allow or control lower-level behavior. Safety-related constraints specify those relationships among system variables that prevent the system from reaching a hazardous state.

Instead of viewing accidents as the result of an initiating (root cause) event that cascades into a series of events leading to a loss, accidents are viewed as the result of interactions between components that result in a violation of safety-related constraints. The control processes (both social and technical) that enforce these constraints must limit system behavior to the safe changes and adaptations implied by the constraints [Checkland, 1981]. Preventing accidents requires designing a socio-technical control structure that will enforce the necessary constraints on development and operation.

2.1.2 HIERARCHICAL SAFETY CONTROL STRUCTURES

Figure 4 shows a generic hierarchical safety control model [Leveson, 2004]. Control structures have to be tailored to the specific organization and system. Accidents arise when components of the socio-technical system do not adequately enforce necessary constraints on behavior (e.g., the physical system, engineering design, management, and regulatory behavior). The model in Figure 4 has two main hierarchical control structures: one for system development (on the left) and one for system operation (on the right), with interactions between them. An aircraft manufacturer might only have system development under its immediate control, but safety requires control over both aircraft development and operations and neither can be achieved in isolation. Safety must be designed into the system as a whole, and safety during operation depends partly on the original design and partly on effective control over operations and the changes and adaptations in the system over time. Aircraft manufacturers must communicate with the airlines and pilots to share their assumptions about the operational environment upon which the safety analysis was based, as well as information

about safe operating procedures. The operational environment, in turn, provides feedback to the manufacturer about the performance of the system during operations.

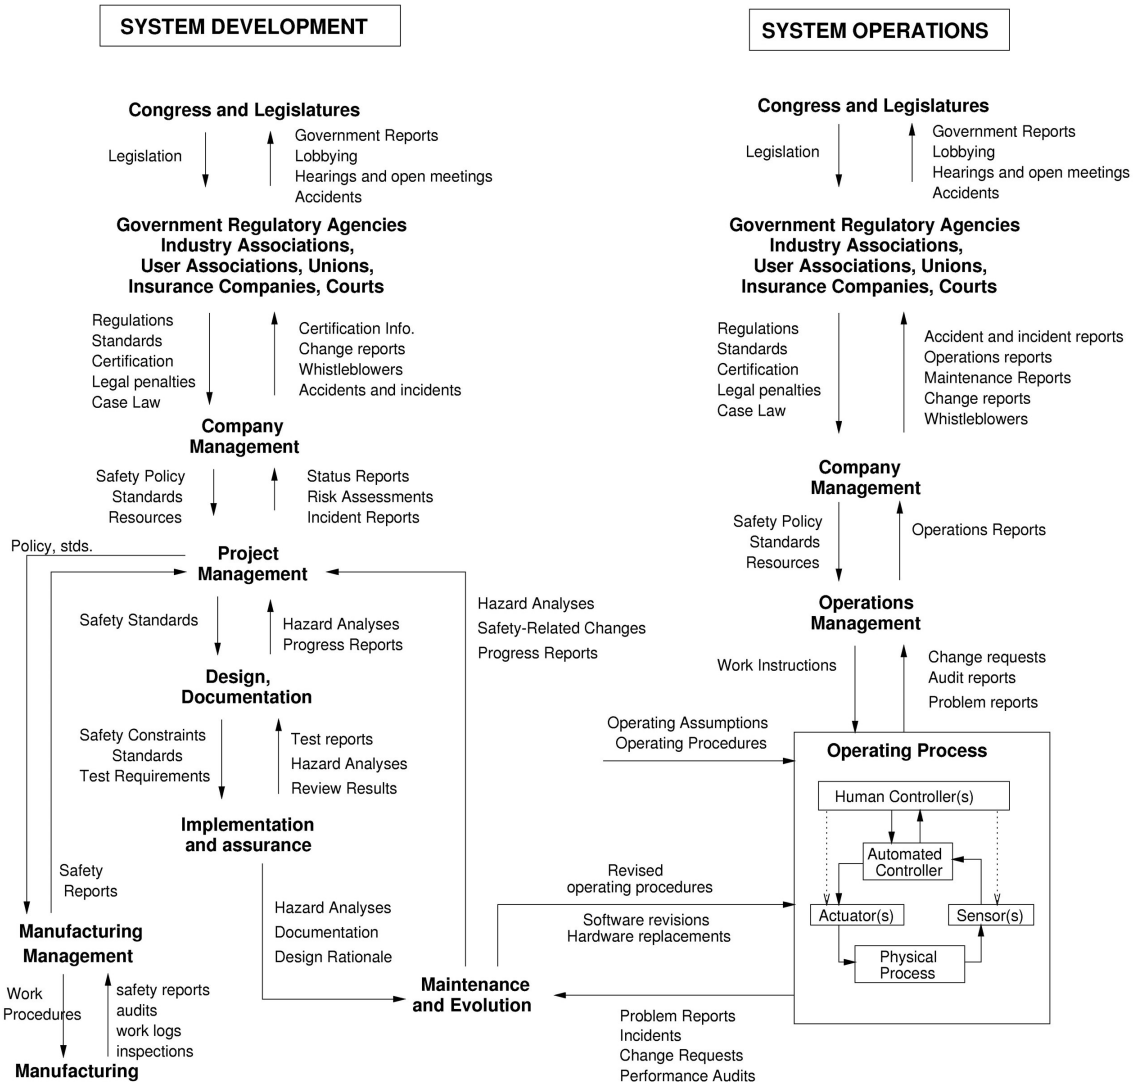


Figure 4: Generic Safety Control Structure (from [Leveson, 2004])

2.1.3 CONTROL LOOPS AND PROCESS MODELS

One or more control loops operate between the hierarchical levels of each control structure with a downward channel providing the information or commands necessary to impose constraints on the level below and a measuring channel to provide feedback about how effectively the constraints were enforced. For example, company management in the

development process structure may provide a safety policy, standards, and resources to project management and in return receive status reports, risk assessment, and incident reports as feedback about the status of the project with respect to the safety constraints.

At each level of the control structure, inadequate control may result from missing safety constraints, inadequately communicated constraints, or from constraints that are not enforced correctly at a lower level. Feedback during operations is critical. For example, the safety analysis process that generates constraints always involves some basic assumptions about the operating environment of the process. When the environment changes such that those assumptions are no longer true, the controls in place may become inadequate. This can happen when systems or system components evolve asynchronously [Leveson, 2004].

Figure 5 shows a typical control loop operating between levels. Leveson describes the functioning and requirements of a generic control loop in [Leveson, 2004]:

“Any controller must have a model (for human controllers this is a mental model) of (1) the current state of the system being controlled, (2) the required relationship between system variables, and (3) the ways the process can change state. Accidents, particularly system accidents, frequently result from inconsistencies between the model of the process used by the controllers and the actual process state; for example, the lander software thinks the lander has reached the planet surface and shuts down the descent engine or the mission manager believes that foam shedding is a maintenance or turnaround issue only. Part of STAMP-based hazard analysis efforts involve identifying the process models required for safe operation, examining the ways they can become inconsistent with the actual state (such as missing or incorrect feedback), and determining what feedback loops, redundancy, or other design features are necessary to maintain the safety constraints.

When there are multiple controllers and decision makers, i.e., distributed control and decision making, system accidents may result from inadequate coordination among several controllers and decision makers, including side effects and conflicts between independently made decisions and control actions. While decision makers usually make decisions that are “locally” rational, when taken into the context of the larger system design and operation, these decisions and actions may interact in unexpected ways to produce an accident. Accidents are most likely to occur in boundary areas between system components or areas of overlapping control. Such coordination flaws are often the result of inconsistent process models. For example, two controllers may both think the other is making the required control action resulting in neither doing it, or they make control actions that conflict with each other. Communication plays an

important role here, and one use for STAMP models is in the design of communication channels and the information each actor needs in a distributed control or decision-making environment.”

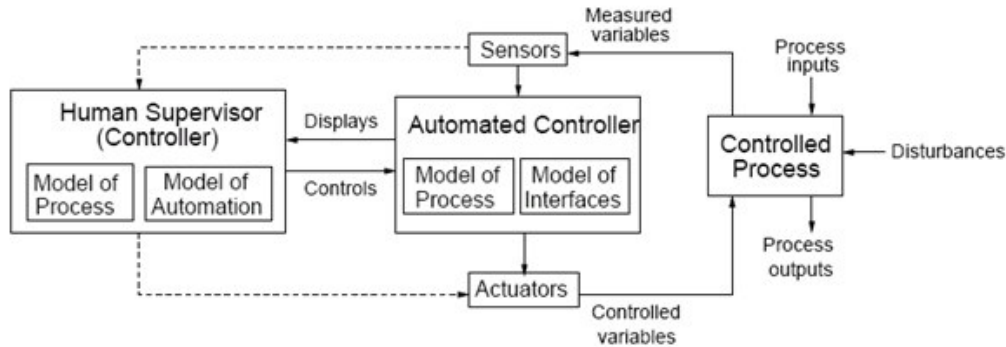


Figure 5: A Generic Control Loop (Adapted from [Leveson, 2004])

Leveson also describes the dynamic nature of safety control structures in [Leveson, 2004]:

“The safety control structure often changes over time, which accounts for the observation that accidents in complex systems frequently involve a migration of the system toward a state of heightened risk where a small deviation (in the physical system or in human behavior) can lead to a catastrophe. The foundation for an accident is often laid years before the loss actually occurs. One event may trigger the loss but if that event had not happened, another one would have. The Bhopal MIC (methyl isocyanate) release, which is among the worst industrial accidents in history, was blamed by Union Carbide and the Indian government on human error, namely the improper cleaning of a pipe at the chemical plant. However, this event was only a proximate factor in the loss. Degradation of the safety margin at the Union Carbide Bhopal plant had occurred over many years, without any particular single decision to do so, but simply as a series of decisions that moved the plant slowly toward a situation where any slight error would lead to a major accident. An argument can be made that both the Challenger and Columbia losses involved this type of long term degradation of safety margins and increasing system risk [Leveson, 2004].”

This dissertation builds upon the dynamic nature of safety control structures by providing methods and tools to analyze potential scenarios where the structure migrates toward a state of higher risk where it can no longer enforce safety constraints, resulting in a hazard state.

Figure 6 shows a classification of control errors that can lead to accidents. The factors are derived from the basic properties of control loops. The classification forms the basis for a new type of hazard analysis called STPA (STamP Analysis) [Leveson, 2003].

- **Inadequate control actions (enforcement of constraints)**
 - Unidentified hazards
 - Inappropriate, ineffective, or missing control actions for identified hazards
 - Design of control algorithm (process) does not enforce constraints
 - Process models inconsistent, incomplete, or incorrect (lack of linkup)
 - Flaw(s) in creation process
 - Flaws(s) in updating process (asynchronous evolution)
 - Time lags and measurement inaccuracies not accounted for
 - Inadequate coordination among controllers and decision-makers (boundary and overlap areas)
- **Inadequate Execution of Control Action**
 - Communication flaw
 - Inadequate actuator operation
 - Time lag
- **Inadequate or missing feedback**
 - Not provided in system design
 - Communication flaw
 - Time lag
 - Inadequate sensor operation (incorrect or no information provided)

Figure 6: A Taxonomy of Generic Control Flaws leading to Hazards (from [Leveson, 2005])

So far, the models presented have been static models of the safety control structure. But models are also needed to understand why the safety control structure changes over time in order to build in protection and monitoring features to prevent and/or correct unsafe changes. For this goal, system dynamics models are used, which are well suited to capture the dynamic processes responsible for changes in the safety control structure in the development and operation of complex socio-technical systems. The next section provides a review of system dynamics concepts and applications.

2.2 SYSTEM DYNAMICS CONCEPTS AND APPLICATIONS

The field of system dynamics was created at MIT in the 1950s by Jay Forrester [Forrester, 1961]. It is designed to help decision makers learn about the structure and dynamics of complex systems, to identify high leverage policies for sustained improvement, and to catalyze successful implementation and change. System dynamics provides a framework for dealing with dynamic complexity, where cause and effect are not obviously related. Similarly to STAMP, it is grounded in the theory of non-linear dynamics and feedback control, but also draws on cognitive and social psychology, organization theory, economics, and other social sciences.

System behavior in system dynamics is modeled by using feedback (causal) loops, stocks and flows (levels and rates) [Sterman, 2000]. In this view of the world, behavior over time (the dynamics of the system) can be explained by the interaction of positive and negative feedback loops [Sterman, 2000] through a stock and flow dynamic structure. The models are constructed from three basic building blocks: positive feedback or reinforcing loops, negative feedback or balancing loops, and delays. Positive feedback loops (reinforcing loops) are self-reinforcing while negative feedback loops (balancing loops) tend to counteract change and seek an equilibrium position. Another key component of system dynamics are delays, which can take various forms, and can cause overshoot and instability in the system.

Figure 7(a) shows a reinforcing loop, which is a structure that feeds on itself to produce growth or decline. Reinforcing loops correspond to positive feedback loops in control theory. An increase in variable 1 leads to an increase in variable 2 (as indicated by the “+” sign), which leads to an increase in variable 1 and so on. The “+” sign does not mean the values necessarily increase, only that variable 1 and variable 2 will change in the same direction (polarity). If variable 1 decreases, then variable 2 will decrease. In the absence of external influences, both variable 1 and variable 2 will clearly grow or decline exponentially. Reinforcing loops generate growth, amplify deviations, and reinforce change.

A balancing loop Figure 7(b) is a structure that changes the current value of a system variable or a desired or reference variable through some action. A “-” sign indicates that the values change in opposite directions. It corresponds to a negative feedback loop in control theory. The difference between the current value and the desired value is perceived as an error. An action proportional to the error is taken to decrease the error so that, over time, the current value approaches the desired value.

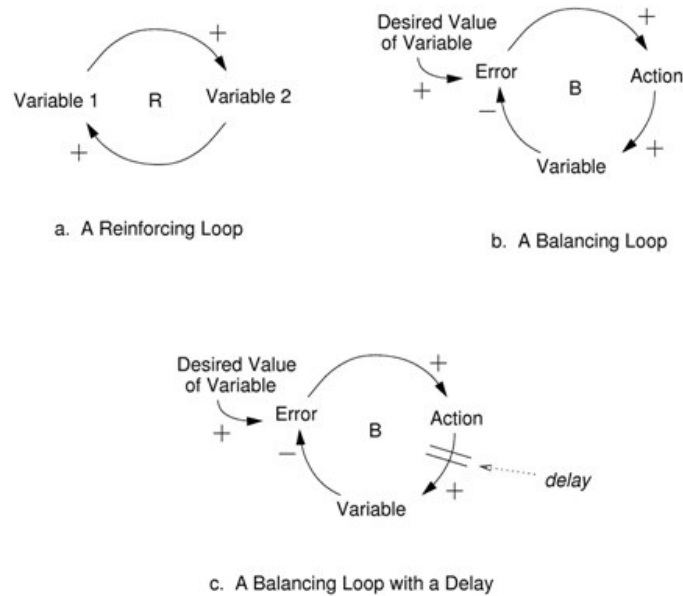


Figure 7: Three basic feedback structures

The third basic element is a delay, which is used to model the time that elapses between cause and effect. In some instances, a delay can be indicated with a double line perpendicular to the causal link, as shown in Figure 7(c). Delays can take all sorts of different forms, including pipeline delays (first in-first out), as well as single or multi-order material delays, and single or multi-order information delays. A detailed discussion of different types of delays can be found in [Sterman, 2000]. Using the right type of delay in modeling is critical. The type of delay used in the various models and components in this dissertation will be discussed as the models are introduced. Delays make it difficult to link cause and effect (dynamic complexity) and may result in unstable system behavior. For example, in steering a ship there is a delay between a change in the rudder position and a corresponding course change, often leading to over-correction and instability.

The dynamic behavior of system dynamics model is created by the stock and flow structure that provides the “integrators” necessary to obtain time-dependent differential equations. The analogy often used for an integrator is that of a bathtub with a faucet as an input and a drain as an output (see Figure 8). The water level in the bathtub is analogous to the level of a stock. In other words, the level of the bathtub at any point in time is the result of the time-integration

of the net flow into the bathtub (net flow = faucet flow – drain flow) from time $t=0$ to time t , given an initial value for the water level.




Figure 8: The bathtub analogy of stock and flows

A simple stock and flow structure of the Lotka-Volterra Predator-Prey model is shown in Figure 9. Two stocks are used in this model: the number of prey (H) in the population, and the number of predators (L) in the population. These two stocks correspond to the state variables of the model. Reinforcing and balancing loops are indicated. The a, b, c, d constants in the differential equations shown in Figure 9 correspond to the four exogenous parameters in the model. The behavior resulting from this stock and flow structure is shown in Figure 10. The amplitude and frequency of the succeeding predator and prey waves depend on the values of the a, b, c, d constants, but the stock and flow structure is inherently oscillatory. Oscillatory behavior will occur for every case situation except the trivial zero input and zero initial value case.

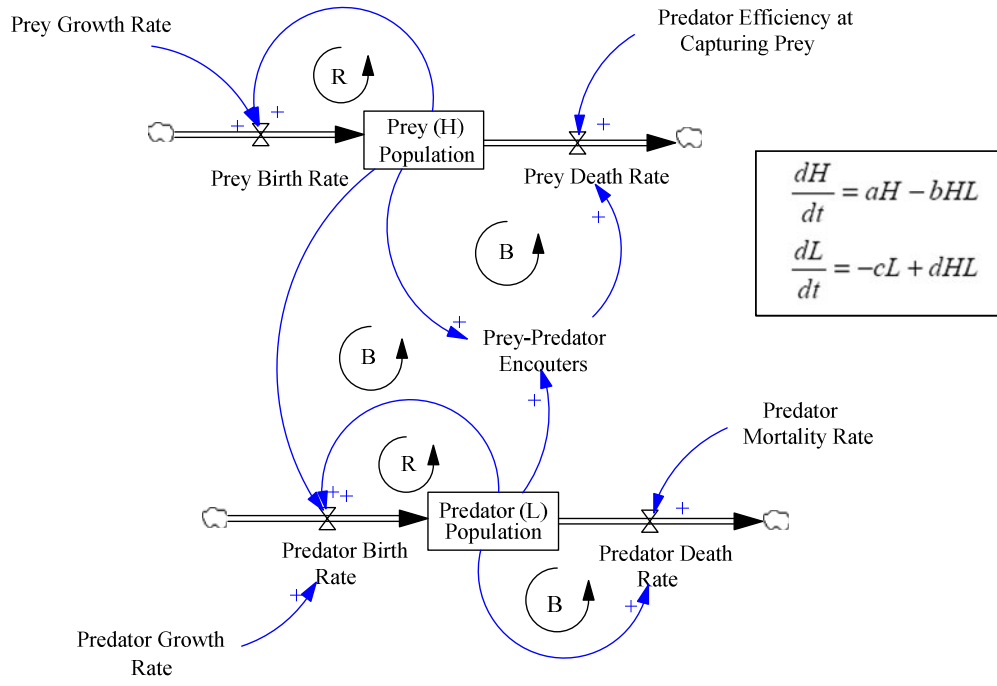


Figure 9: The stock and flow structure of a classic predator-prey dynamic model

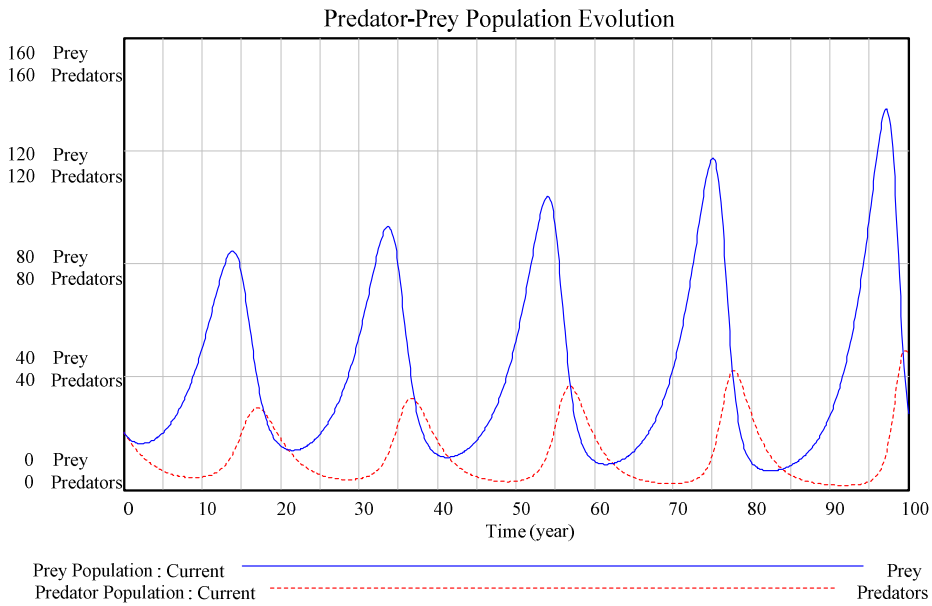


Figure 10: Dynamic behavior generated by the predator-prey structure

2.2.1 REVIEW OF COMMON SYSTEM DYNAMICS APPLICATIONS

System dynamics has been used throughout the years to find solutions and high-leverage policies to help solve the problems faced in complex dynamic systems. Complex dynamic systems are defined in the system dynamics field as systems that: (1) are extremely complex, consisting of multiple interdependent components; (2) are highly dynamic; (3) involve multiple feedback processes; (4) involve nonlinear relationships; and (5) involve both hard and soft data. One of the earliest application was the demonstration that some industrial production-distribution systems can endogenously create oscillatory behavior [Forrester, 1961]. The same concepts were further developed to include the impact of bounded rationality [Morecroft, 1983], feedback and delays on dynamic decision-making using a simple inventory control experiment that became the basis for the famous “beer game” played every year by thousands of company executives and MBA students from around the world [Sterman, 1989; Sterman, 1989].

Other classic applications include the study of urban dynamics using a formal system dynamics model to study urban growth and decay and show that many well intended policies for urban planning and development may lead to unintended negative side effects such as the creation of neighborhoods with high unemployment, cycles of low and high real estate occupancy, and gridlocks and long commute times [Forrester, 1969]. In another classical and controversial study, the limits to growth (WORLD3) model was used to argue that current policies are unsustainable and will lead to a grim future for humans including a highly polluted environment, high mortality and low standards of living [Forrester, 1972; Meadows, 1972; Meadows, 1992].

Many additional models and studies used system dynamics to understand the dynamics of cycles such as the economic long wave [Sterman, 1985; Sterman, 1986; Forrester, 1989] and to show how micro-behavior may create cycles and oscillations in the orders and delivery of products such as new aircraft [Lyneis, 2000]. System dynamics also has a long history of grappling with corporate strategy problems [Lyneis, 1980; Morecroft, 1984], as well as quality and process improvement including many case studies covering semiconductor companies, motorcycle manufacturers, and maintenance practices at Dupont and BP plants [Sterman, 1997; Repenning, 2001; Repenning, 2002].

The dynamic risk management work presented in this dissertation builds upon the tradition of system dynamics research mentioned above, but some of the research has a more direct impact, including the work related to complex project dynamics and management [Ford, 1995; Ford, 1998; Lyneis, 1999; Repenning, 2001], as well as decision-making [Morecroft, 1983] and learning [Morecroft, 1988; Senge, 1990]. Some important related work has also been done in the field of risk management, including the impact of variations in the quantity and timing of arising problems and interruptions on the onset of technological accidents [Rudolph, 2002], and the impact of management decisions and pressures on the safety of a coal production system, using the Westray Mine accident as a case study [Cooke, 2003].

The previous sections provided a short review of relevant work and topics in system-theoretic accident modeling, and system dynamics modeling. Further review of the literature associated with generic structures, model building, and model validation will be provided in later chapters. In the rest of this chapter, we will investigate complex systems behavior patterns according to lifecycle phase, divided into system development and system operation dynamics. The chapter will conclude with an example of the impact of reinforcing dynamic structures on the effect of the Independent Technical Authority (ITA) in NASA's space shuttle program.

2.3 SAFETY-CENTRIC SYSTEM DEVELOPMENT

Unlike other accident models, STAMP considers safety starting from the very beginning of system conceptual development, and continuing through system design, implementation and operation. The impact of the entire socio-technical structure on system safety is considered from the start. However, building safety into a new system and attempting to mitigate potential hazards during system development is very different from attempting to safely operate an existing system. In other words, the dynamic patterns that contribute to changes in the safety control structure of complex systems are very different whether we are trying to develop or operate a system. For most systems, development and operation happen on different time scales simply because a system must be developed and implemented before it is operated. There may be exceptions to this rule, such as systems developed entirely using a spiral process where the development and operation are performed concurrently by starting

with a very simple system and adding functions and capabilities while operating the system. However, for the type of complex systems we are concerned with, that is, large-scale complex socio-technical systems, we assume that system development and operations will be mostly performed on different timescales, while allowing for overlap periods during transition from development to operations as well as partial or timed system deployment and evolution.

The goal of sections 2.3 and 2.4 is to introduce (in action) the concept of causal feedback loops and stock-flow structures to analyze the dynamic parts of a system lifecycle. The focus of section 2.3 is on the dynamics occurring within the left column of the generic STAMP safety control structure (See Figure 11), that is, system development. Some of the most important feedback loops and their impacts will be discussed individually in the next subsections. These loops were derived from existing literature on the dynamics of project management and system development, as well as from dynamic safety archetypes and on the author's direct interactions and interviews with project management professionals at NASA's Exploration Systems Mission Directorate. Section 2.4 will describe some of the feedback loops and behavior modes occurring during system operation.

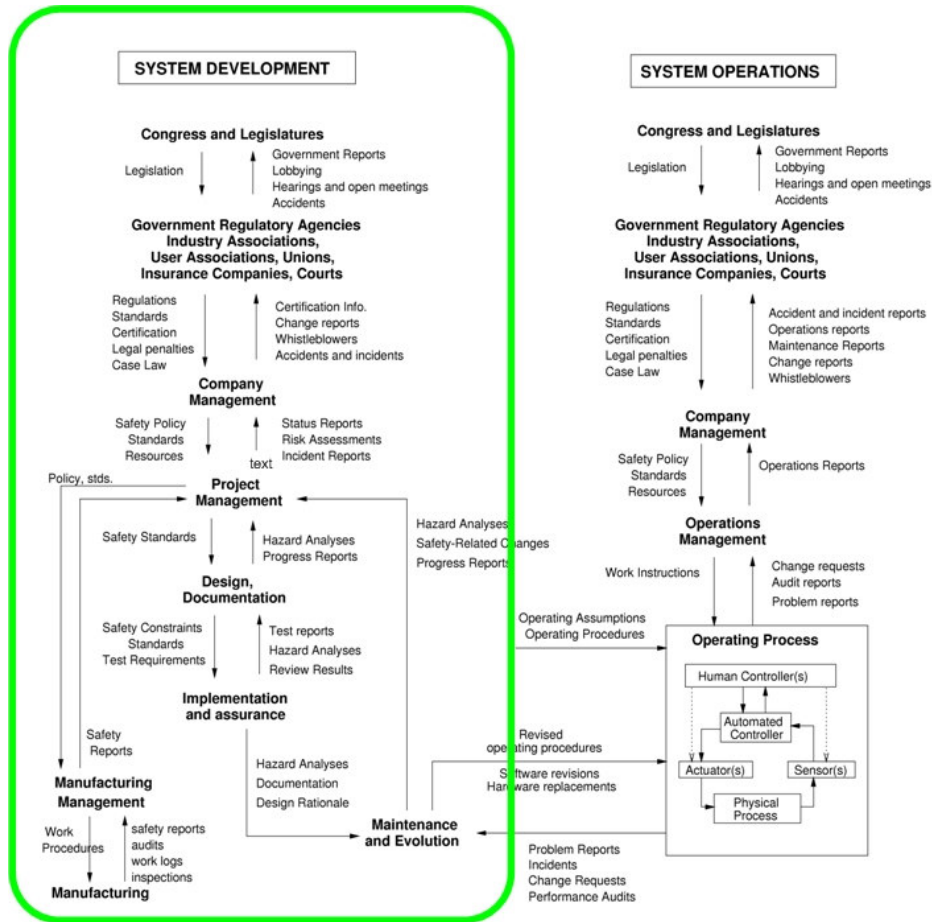


Figure 11: Generic System Development Structure (from [Leveson, 2004])

2.3.1 LOOP B1 - DELAYS CAUSE PRESSURE

The first and arguably most critical balancing loop is loop B1: "Delays Cause Pressure", or the schedule pressure balancing loop (See Figure 12).




Figure 12: Loop B1: Delays cause pressure

the desired completion fraction, schedule pressure will be applied to restore equilibrium. The amount of schedule pressure applied depends on the proportional gain (P gain) of the controller. For example, let's assume an externally applied decrease of 20% in workforce capacity at time=30 months. In the open-loop control case (P gain = 0), the completion fractions start to diverge at time t=30 months and the divergence increases until the project is finished (see Figure 14).




Figure 14: Impact of disturbance on completion fraction

In the closed-loop case (P gain > 0), the controller applies schedule pressure to reduce the schedule delays. The more schedule pressure is applied (the higher the P gain), the more schedule delays are reduced. However, it should be noted that once the equilibrium is disturbed, the proportional controller by itself will not be able to bring back the project exactly on schedule. In control theory term, this can be easily explained because the steady state error to a ramp input (which is similar to a perfectly planned linear project) does not converge to zero with proportional control alone (see Figure 15).




Figure 15: Completion fraction with disturbance and controller with $P>0$ and $I=0$

Adding integral control (I gain) is a slightly more sophisticated and realistic way of bringing a project back on schedule. The result is to effectively reduce the steady state error to zero by planning work ahead to compensate for the moving target effect of schedule completion. The result of adding integral control on completion fraction can be shown in Figure 16. This adjustment process is a common project management practice consisting of re-routing critical paths and schedules to compensate for inadvertent delays. In normal circumstances, and given sufficient system development capacity and/or schedule and resource reserves, it is a relatively straightforward management activity.




Figure 16: Completion fraction with disturbance and controller with $P>0$ and $I>0$

2.3.2 LOOP R2: THE REWORK/BURNOUT CYCLE

This very simple balancing loop is the main feedback mechanism responsible for keeping the project on schedule. However, other reinforcing mechanisms may reduce the strength of the “Delays cause Pressure” loop. One important loop is the “Burnout Cycle” loop that limits the impact of the “Delays cause Pressure” loop (see Figure 17).




Figure 17: Adding loop R2: Burnout Cycle

The burnout loop mitigates the impact of schedule pressure on completion rate because as more work is performed by the same number of employees, burnout starts to occur and productivity decreases over time as employees become tired and overwhelmed. Consequently, adding the burnout loop to the previous structure reduces the impact of the balancing loop, preventing the project from getting back on schedule, despite the pressures and coordination efforts of project managers (see Figure 18).




Figure 18: Completion fraction including the burnout cycle (R2)

2.3.3 LOOP R3: THE BURNOUT CYCLE

Another reinforcing loop affecting development completion is the basic rework cycle. The rework cycle is a standard component of development dynamics and has been discussed in great detail in the project dynamics literature [Ford, 1995; Lyneis, 1999; Sterman, 2000; Lyneis, 2001]. As burnout increases, people are overwhelmed and while burning the midnight oil to remain on schedule, they make subtle mistakes that create the need for more rework. A systems engineering manager at NASA we interviewed about schedule pressure told us:

“Shorter cycle makes you make the schedule, but you lose fidelity. But then, you can’t drag the cycle so long that you get better answers but you waste other people’s time. The fidelity of analysis may go down with shorter cycles, and the fatigue kicks in. People make subtle mistakes. You need some schedule pressure, but not so much that you burn out people or make mistakes.”

In addition, there is strong evidence that the relationship between “Pressure to deliver” and subtle flaws and mistakes is not linear, as a NASA employee explains:

“Schedule pressure is not a bad thing if it’s applied right. Schedule pressure is necessary, so it’s a positive thing too. People don’t produce as well without schedule pressure. It’s a matter of when the schedule pressure goes over the edge, and your fraction of tasks with flaws goes up too high. It’s almost like an exponential curve: for a long time, the effect is not too bad, but when schedule pressure is too high, people just give in, and they say: “Whatever you want, you got... You want that thing out the door? You got it!”. Productivity increases with schedule pressure, but flaws increase too.”

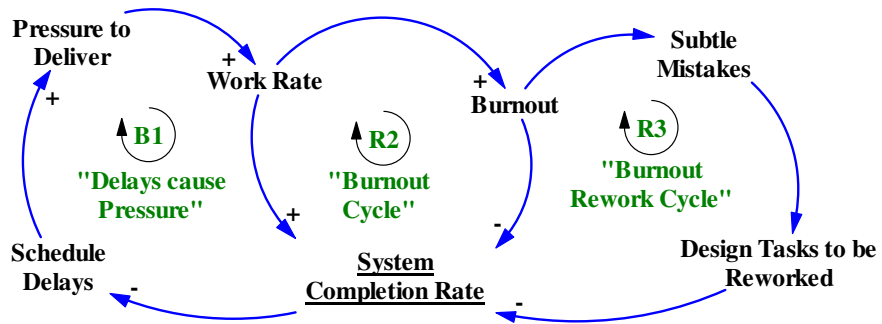


Figure 19: Adding loop R3: Burnout Rework Cycle

Figure 20 shows the impact of the Burnout Rework Cycle (R3) on the fraction of development tasks that need to be reworked because of subtle mistakes. As workload increases and burnout starts to accumulate, the fraction of tasks requiring rework increases until the project is completed. In addition to having an impact on schedule, it also increases overall cost because of avoidable scrapped work and rework.

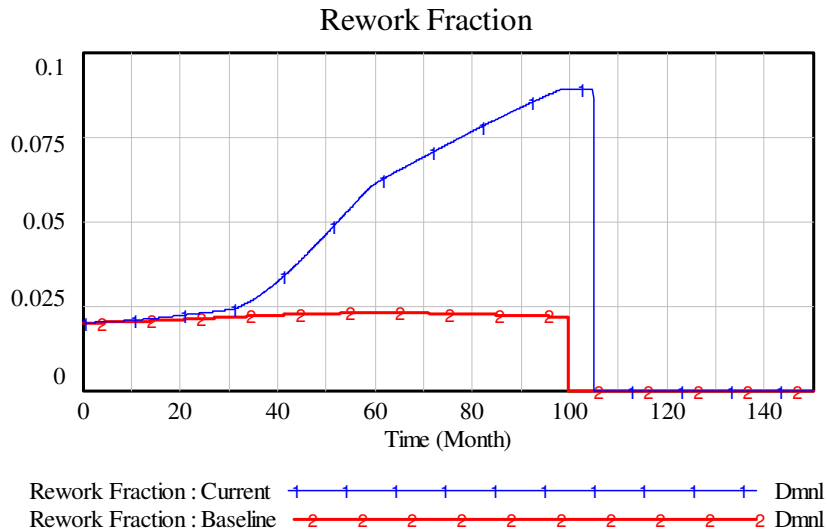


Figure 20: Impact of burnout cycle on rework fraction




Figure 21: Impact of burnout and rework on completion fraction

2.3.4 LOOPS R1A/B: SAFETY AND INTEGRATION

Other loops having an impact on system development are related to the “Quality and Timeliness of Safety and Integration Activities”. In the high-level model, those two types of activities are combined and have a similar impact on development. In the more detailed models to be presented in Chapter 4, safety and integration activities are shown to have different impacts on the system behavior and will be decoupled. However, as a first order approximation, as schedule pressure increases because of disturbances or development delays, the effective priority of the safety and integration activities increase. As schedule pressure increases because of development delays or overoptimistic planning, more effective priority is allocated toward getting the hardware built and delivered, at the expense of less visible activities such as safety and integration. There are multiple effects of this feedback loop that will be discussed in detail in chapter 4. At a high level, less effective priority toward safety and integration efforts reduces the impact, quality and timeliness of analyses through soft factors such as a loss of influence and power of the safety and integration efforts and people as all the effort and resources are allocated toward product delivery. In a similar way, the inevitable resource pressure coming from Congress or the NASA administration will be aimed primarily at activities not directly and immediately critical to system delivery. Those two reinforcing feedback loops are shown respectively as the R1 and R1b loops in Figure 22.

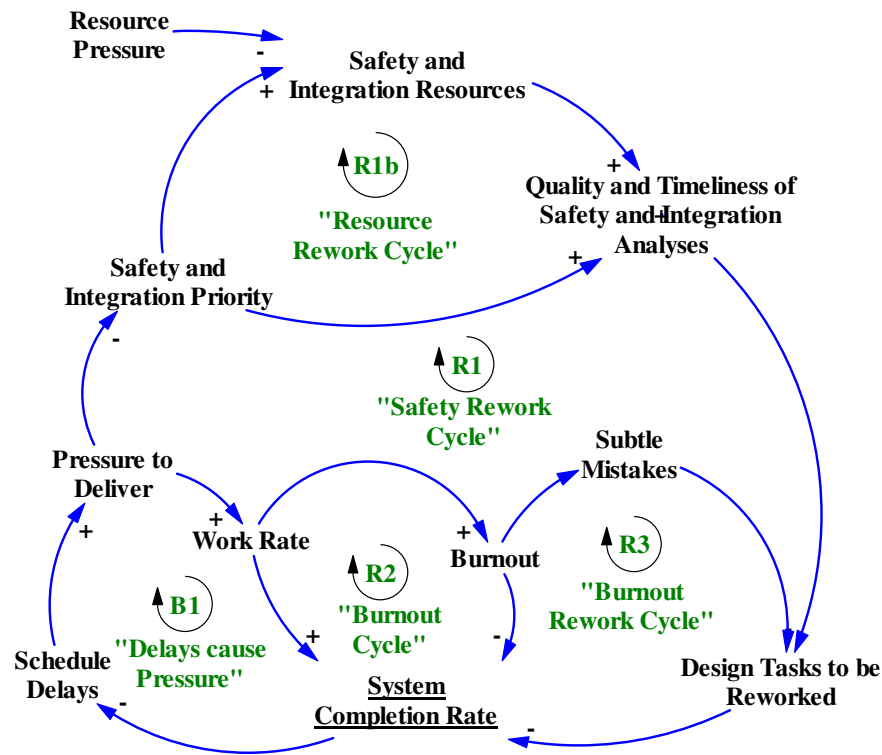


Figure 22: Adding loops R1 and R1b: The impact of safety and integration on rework cycle

Excellent comments were collected during the interviews to support this claim. This particular one on the topic, from a high-level NASA/ESMD manager is rather long, but enlightening:

“Well I think that the way to look at it is, if you just take this one side [pointing at the chart], incoming resource pressure from Congress. They always want to do more with less, [...] they want you to do more program or they want to have more content or they want to have more influence, whatever it may be. It’s a good thing and it’s a bad thing because on the one hand they want you to do it, but on the other they want you to do it for less than you believe that you can do it, so what’s the impact of that?”

The most obvious impact is, it’s going to result in budget pressures because they’re going to say, ‘Well we want you to do these 10 things, but we really are only going to give you enough to do 8.’ So what do you think we are going to do with our estimates when we come in, we’re going to low-ball the estimates, I mean not intentionally, but that’s the dynamic that’s been set in there, right, you know, do 10 things for the price of 8... so you go off and you figure out how to do that and you take risks and challenges in your budgets and everything else

and you say, 'Yes I can do it for this price IF—now everybody forgets the if—if everything goes right, if I get the money exactly when I'm supposed to get the money, if there are no continuing resolutions, if Congress doesn't change its intent, right—and that's all of the stuff that can happen from the outside—and if we don't run into any problems, and if the technology is at the maturity that we said it was going to be at that time, and if the contractors all do what they say they are going to do, and if the NASA workforce is where it's supposed to be, and if all those other activities that I'm doing now that I expect to be done when this starts are done when this starts.' So you've got all of those things that come in, but when you make that commitment you can say all of those ifs—I've been to this movie and I've said all of these ifs—and you know what?... they never remember those ifs, for whatever reason they don't remember all of those ifs.

So then what happens? Well, Congress gets upset because we've overrun our budget, we get upset because we say, 'Well we told you!,' (and this is true for anyone above the project including the AA) and so they say, 'Well it doesn't make any difference, (we've got) a different Administrator, different AA, different center director, whatever, different Congress, live with it.' So then what happens, well training is the easiest thing to impact. Now remember what I said about having good people that know what they're doing? Well guess what? They aren't going to be as good and they aren't going to know as much because I'm cutting back on training. So what does that mean for the quality and quantity of safety analysis: It's less. What does that mean as far as what I can do in terms of -- since you're talking about safety and you could put anything that you want in here -- it means that I'm going to put all of my resources to getting the job done. So getting the job done has a nice, clear path: build the hardware, build the software, assemble the hardware, assemble the software, test the hardware, test the software, and launch it. Anything else on the outside is going to be starved as much you can possibly starve it. So that's the problem and probably one of the most important things that you can say for a project or a program is it succeeds or fails the day you put your first plan up.

You know everybody uses (the example) of Apollo as a great success, but everybody forgets the IF there, when James Webb was asked how much was Apollo going to cost, he initially said \$12 billion or \$13 billion. When he was asked by—and I don't remember the story, it changes, I'll have to find out what the reality is—either by Congress or OMB, 'Oh come on you guys always low-ball it,' he said \$25 billion. Okay, so they had a good plan and (they said it was going to cost) roughly twice what they thought that it was going to cost. Now it cost less than \$25 billion, but it cost more than \$12 billion. Okay, what have we done since then, we've gone kind of the opposite. Every program—Shuttle, EOS that I worked on—the first estimate for it was unrealistic, it was like \$40 billion and everything knew it wasn't going to get \$40 billion so we knew we had to go back to the drawing boards on that. But when we came up with a realistic estimate, and we had what everybody agreed was the right content, we went in and said that it was \$17 billion (and they said that) it's too much (you need to)

cut it back. You want to do this program we know that you can do it for more, it doesn't make any difference who you're talking to, you know, you've got to be able to do it for less. We went from \$17 billion to \$8 billion. Now we had to take content out to get that done, but if you look at every program that we do, it's exactly the same thing. So the pressure from the outside is there. [sic]"

The impact of the R1a and R1b outer loops on the development dynamics is significant, perhaps even more than the reinforcing loops documented previously. While the detailed impact of development tasks sequencing, overwork, burnout, and mistakes during a project has been addressed extensively in the literature, the impact of the safety and integration activities on the strength of the rework cycle has not been the focus of attention. However, for the development of complex safety-critical systems, safety, integration and software development activities are likely to be the bottleneck. Consequently, the thoroughness, quality, and timing of safety and integration activities must be modeled as accurately as possible if we aim at improving complex system development processes. Figure 23 and Figure 24 show the cumulative impact of adding the effect of the R1a and R1b loops to the previous feedback structure. The cumulative impact is significant, creating an order of magnitude increase in the fraction of tasks requiring rework, which also causes large development delays, with associated cost overruns. The exact amount of the increase is difficult to measure and assess, and the high-level formulations used for this simple model could be refined, but in the end, the cumulative effect of the three reinforcing loops creates significant delays, cost overruns, and safety and quality problems. In these conditions, without the addition of significant additional development resources, the development cost, schedule, safety, and most likely system scope will suffer significantly from disturbances that affect initial program/project planning.

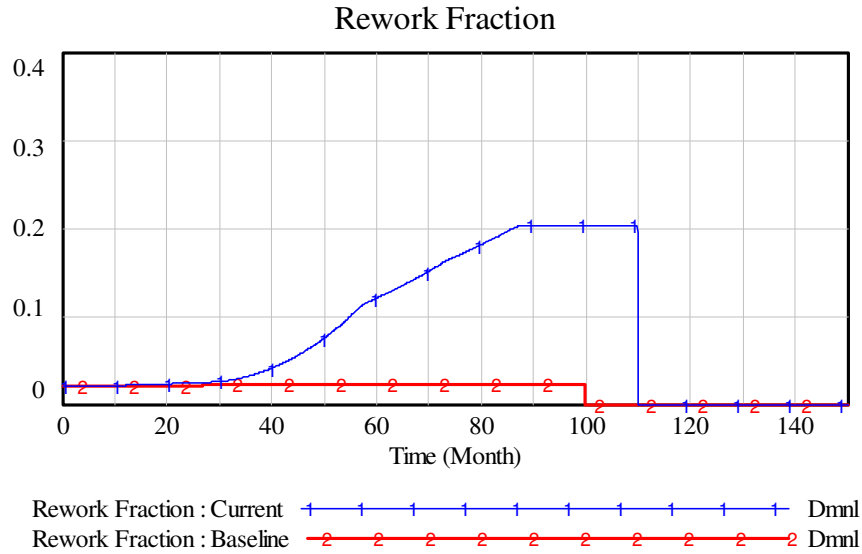


Figure 23: Impact of loops R1 and R1b on rework fraction




Figure 24: Impact of loops R1 and R1b on completion fraction

2.3.5 OUTER LOOPS: WAIVERS, COST AND RESOURCES

The main high-level loops that create the system development dynamics were described in the previous subsections. Outer loops that may not necessarily have an impact on the system during development, but that may impact operational characteristics, are shown in Figure 25 using dotted causal links. Lifecycle cost was a recurrent theme during interviews. Often,

lifecycle cost is informally defined by system designers and manufacturers as the costs incurred from the start of system development up to hardware delivery. In many instances, operating cost is not explicitly included in lifecycle cost. The rationale is that it is very difficult to estimate operating cost before the system is delivered and operated. For complex systems such as NASA's space exploration system, where the system will be operated for a long period of time, and with limited reuse between missions, it is more convenient and common practice to define lifecycle costs as the costs incurred before operation. The author believes it is a mistake to neglect the impact of built-in system characteristics on the system operation just because it is difficult to assess early in the lifecycle.

Cost estimation is very difficult and data intensive, especially for radically new, large-scale complex systems including much software and new technologies. The usual method for evaluating cost is to divide the work to be performed according to a standard work breakdown structure, evaluate cost for the development and/or acquisition of each component or subsystem, and assemble the cost of subsystems, while allowing for system engineering and integration costs, as well as management and overhead. At our level of analysis, the resolution of standard cost estimation methods is too high to be useful. Instead, we create a system lifecycle cost variable based on a few proxy variables. According to interview data, two main factors will influence the cost overruns for a project. The first factor is the amount of work done to correct mistakes and problems found at any stage of system development. The second factor is the project or program completion time relative to planned time. Those two factors are correlated. Unless enough management reserves are available to allow for additional development work, there are two possible options: either the problem will be accepted and requirements waived, or the program completion will be delayed in order to use the resources budgeted for the following fiscal year.

Cost overruns and schedule delays have an impact on the satisfaction of project sponsors or funding organizations. This impact is very important because it can switch the polarity of outer loops from reinforcing to balancing or vice versa. If the project sponsor reacts negatively to delays and overruns by adding more resource pressure, the polarity of the loop will be reinforcing, creating more delays and problems. On the other hand, if the sponsor reacts by adding resources to alleviate delays and problems, the outer loops are balancing and

pressures are diminished. The response of project sponsors and funding organizations usually depends on project context and criticality, as well as on the magnitude of delays and overruns. Consequently, the impact of outer loops (shown dotted in Figure 25) will be discussed further in more detailed models.

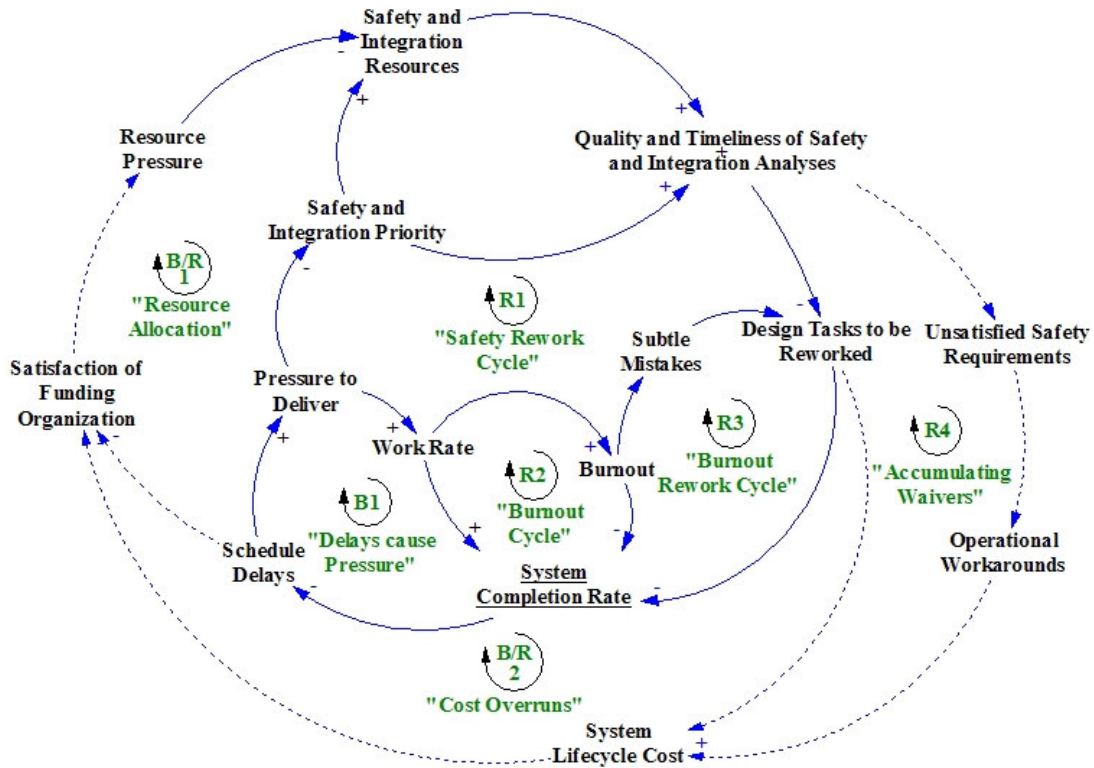


Figure 25: The Impact of Cost and Schedule on Funding and Pressure

2.3.6 IMPACT OF SYSTEM DEVELOPMENT ON OPERATIONS

The purpose of section 2.3 was to introduce the concept of using causal loops and stock-flow structures to analyze some safety aspects of system development. In a more complete model, dozens of additional variables would be used to track real system characteristics, such as the amount of resources (material and human) allocated to a particular project, as well as the number of tasks allocated, the number of tasks completed, the number of safety analyses completed and used in design. However, very simple models such as the high-level development model presented in this section can provide useful information on behavioral patterns that may create a system with undesirable characteristics such as high operations cost

and poor system safety. The design characteristics of a system have a direct impact on system operations, which is discussed of the next section.

2.4 SYSTEM OPERATIONS, SAFETY EROSION AND MIGRATION TOWARD HIGH-RISK

Figure 26 shows a simple model of safety dynamics initially created based on the Columbia accident and later generalized. A high-level model is useful in understanding some of the behavior patterns responsible for the migration of systems toward states of high risk. For example, this simple model proved useful in communicating with system safety experts, NASA managers and astronauts about some safety dynamics factors leading to the Columbia shuttle accident.

There are three main state variables in the model: Safety, Complacency, and Success in meeting expectations. The feedback loop in the lower left corner of Figure 26, labeled R1 or Pushing the Limit, shows how as external pressures increased, performance pressure increased, which led to increased launch rates and thus success in meeting the launch rate expectations, which in turn led to increased expectations and increasing performance pressures. This is an unstable reinforcing system that would create exponential growth or decay if left unbounded.

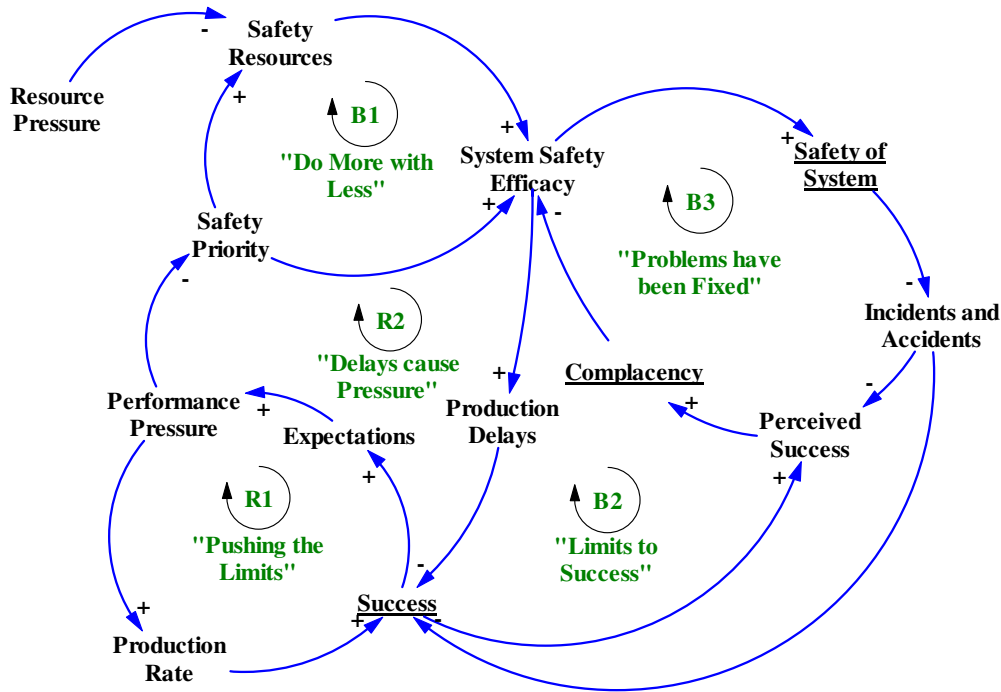


Figure 26: Simple generic model of safety dynamics during system operation

This unbounded reinforcing situation cannot be maintained indefinitely. Other balancing feedback loops such as loop B2, labeled *Limits to Success*, constrain the reinforcing dynamics. As success increases, perceived success follows, which causes an increase in complacency that eventually limits success through lower system safety efficacy. The upper left loop represents part of the safety program loop. The external influences of budget cuts and increasing performance pressures that reduce the priority of safety procedures lead to a decrease in system safety efforts. This decrease combined with loop B2 in which fixing problems leads to higher perceived success and increased complacency. Higher complacency also contributes to a reduction in the effectiveness of system safety efforts, eventually leading to a situation of (unrecognized) high risk.

In modeling the safety dynamics leading to the Columbia accident, another important factor had to be added to the model: increasing system safety efforts leads to additional launch delays, another reason for reducing the priority of safety efforts in the face of increasing

launch pressures. This delay in launch or production is shown in Figure 26, but it may be neither possible nor relevant in every system to reduce production based on safety concerns.

Delays are not explicitly shown in the diagram of Figure 26, but are embedded within the stock-flow structure and are critical to define the dynamics of the system. While reduction in safety efforts and lower prioritization of safety concerns may lead to accidents, systems can operate in a high-risk state for a period of time before an accident occurs. The result is that false confidence is created that the reductions in safety efforts do not have a significant impact on safety. As a result, pressures increase to reduce the efforts and priority even further as the external performance pressures mount. The result of these pressures is a succession of cycles of success, created in part by high commitment to safety, followed by a delayed increase in risk and complacency, causing a reduction in success until it is so low that safety is seen as highly critical again, creating another cycle of high success. Figure 27 shows the oscillatory behavior created by the stock and flow structure shown in Figure 26. This conceptual model is used to analyze the cyclic behavior of alternating success and high-risk phases. The magnitude and the frequency of the behavior cycles are highly dependent on model parameters, which is as expected since this model is highly generic and various systems have different characteristics. However, regardless of the value of these parameters, the oscillatory behavior remains as well as the phase between the peaks of alternating variables.

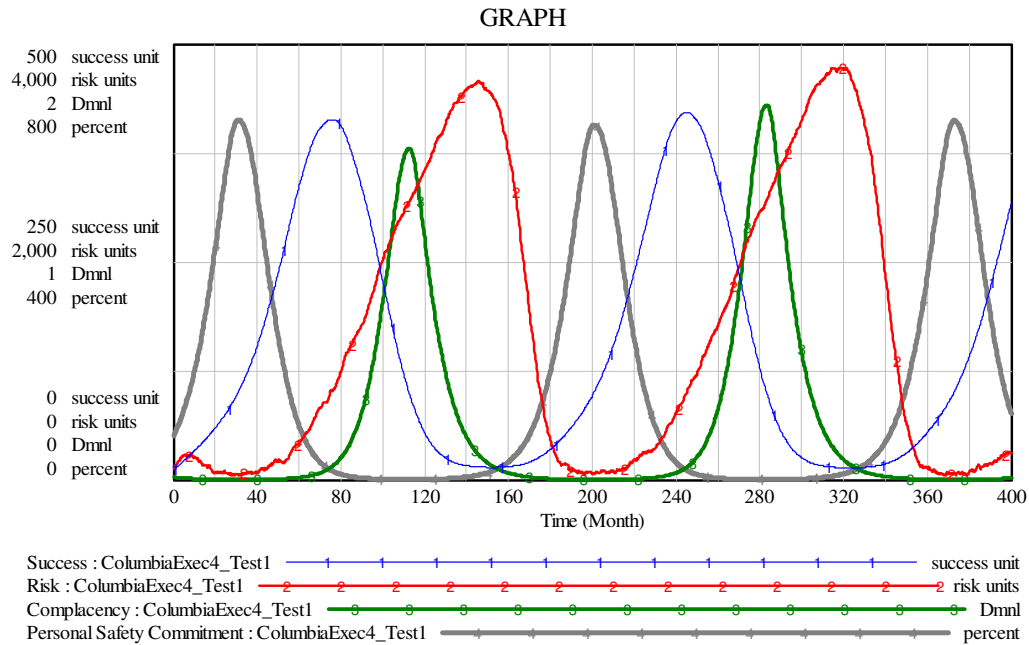


Figure 27: Structure-created oscillatory behavior of model state variables

The models introduced in this chapter can be used to devise and validate fixes for the problems and to design systems with inherently lower risk, or systems that will resist or detect and actively correct a potential migration toward a state of high risk. For example, one way to control the oscillatory behavior of the model in Figure 26 is to anchor the safety efforts by, perhaps, externally enforcing standards in order to prevent schedule and budget pressures from leading to reductions in the safety program. Other solutions are also possible, such as working on programs to ensure that even though some level of complacency occurs, which is almost unavoidable for highly successful systems without accidents or serious incidents, it is either limited or has a limited impact of the efficacy of system safety efforts. One example of a successful program to limit the impact of complacency on system safety has been implemented by the Navy submarine program. In this program, engineers are required to participate in a system safety training session each year. At that time, the tapes of the last moments of the crew during the Thresher loss are played, reminding engineers about the importance of staying alert, curious and inquiring, even in systems that have extraordinary safety records. The potential impact of limiting the effect of success on complacency on the system behavior is shown in Figure 28. Just as in the simulation of Figure 27, a stochastic component was added to the risk input in order to account for some randomness in the way

problems arise and are addressed in complex systems. The impact of effectively disconnecting success from complacency (in Figure 26) is significant. It has the potential to completely neutralize the oscillatory tendency of the system. See how the values for *complacency* and *commitment to safety* rapidly stabilize in Figure 28 while *risk* remains at a relatively low and stable level (notwithstanding the stochastic component), allowing *success* to continually increase.

Depending on the system, it may be very difficult to devise and implement policies to completely neutralize this tendency. However, the models can be used to find alternative solutions and policies, and to evaluate them for their potential effect on the system and impact on risk. The simple models presented in this chapter can only provide high-level conceptual insights into the system dynamics. However, the creation of more complete models customized to the system under analysis allows the user to perform much more thorough analyses based on real life detailed policies and metrics. This topic will be discussed in the following chapters.

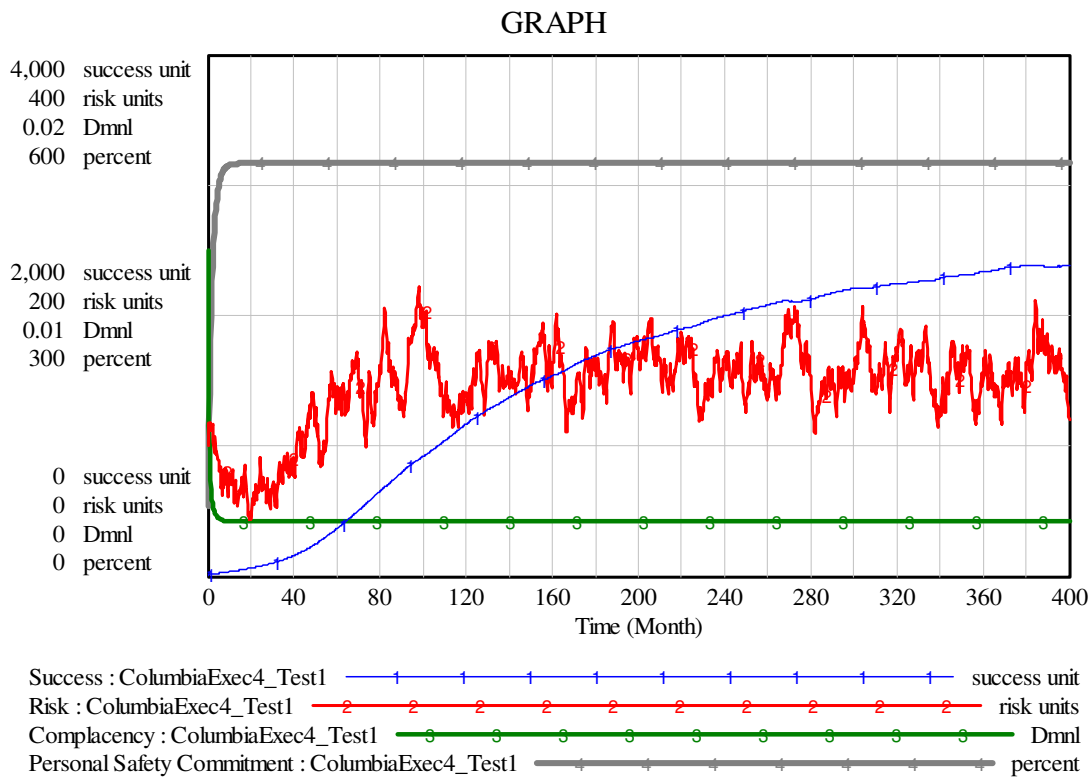


Figure 28: Anchoring safety efforts by limiting the impact of budget pressure and complacency

2.5 EXAMPLE AND SUMMARY

This chapter provided a review of some of the theoretical foundations upon which the rest of this thesis builds. High-level feedback loop structures were presented for complex systems development and operations. A real-life example of their application is provided in Appendix D. Until now, the feedback loop structure was presented by abstracting away the STAMP safety control structure and organizational components within which the high-level loops are embedded. The next chapter will present guidelines for the creation and analysis of STAMP control structures as a preamble to the component-based modeling methodology and analysis presented in chapters 4 and 5.

CHAPTER 3: GUIDELINES AND CRITERIA FOR THE CREATION, ANALYSIS AND MONITORING OF STATIC SAFETY CONTROL STRUCTURES

The purpose of this chapter is to provide guidelines for the creation, analysis and monitoring of static safety control structures that are at the core of the STAMP accident model [Leveson, 2004; Leveson, 2006]. The component-based dynamic model creation and analysis methodology introduced in the two following chapters builds on the concept of STAMP safety control structures. A complete overview of a STAMP-Based risk analysis process is shown in Figure 1. The steps were thoroughly documented using the ITA Analysis in [Leveson, 2005]. They include: 1) A standard Preliminary Hazard Analysis (PHA) where hazards and associated safety requirements and constraints are identified, 2) The modeling of the safety control structure of the system including the safety-related roles and responsibilities of each component, as well as the feedback and control channels across components, 3) A gap analysis where safety requirements (from step 1) are compared with the safety-related responsibilities of components (from step 2), 4) A detailed hazard analysis where detailed risks and inadequate control actions are identified, 5) An analysis and categorizing of risks identified, 6) the creation of a STAMP-based dynamic model and its use to further analyze risks identified in step 4, and 7) The testing and formulation of policies and structure changes to mitigate risks and the identification of early indicators of risk increase. A short example of steps 3-5 is provided in section 5.2.1.1.

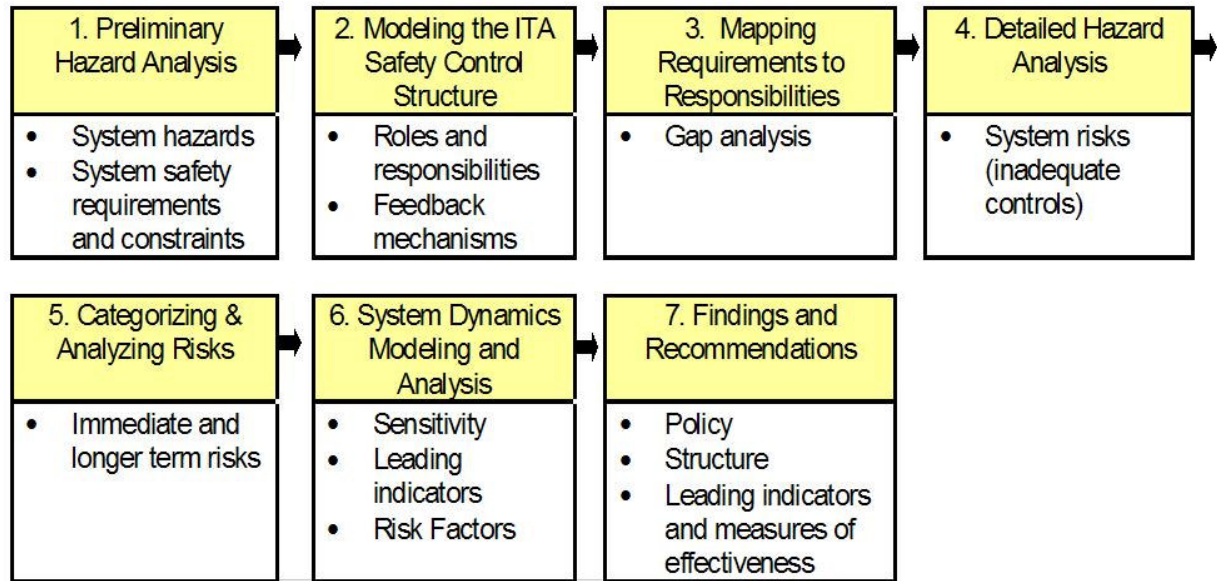


Figure 29: The STAMP-Based risk analysis process (adapted from [Leveson, 2005])

The key contributions of this dissertation are steps 6 and 7 of the STAMP risk analysis process, where STAMP-based dynamic modeling and analysis are introduced and demonstrated in a real system (in chapters 4-6). However, before dynamic models can be created, an accurate static model of the static safety control structures must be defined. Safety control structures are at the center of a STAMP-based risk analysis. The present chapter focuses on the Step 2 of the STAMP process: Modeling the Static Safety Control Structure.

In this chapter, additional guidelines are provided to facilitate the creation of effective and accurate safety control structures for risk analysis and management. In the earlier part of the chapter, guidelines are provided to select the components that will form the base of the control structure. Next, the concept of a generic control structure connector is introduced and a taxonomy of generic connectors is proposed to assemble components into a complete control structure. Later in the chapter, criteria are introduced to perform a first order evaluation of static control structure models based on the generic connector concept.

3.1 SELECTING CONTROL STRUCTURE COMPONENTS

The concept of a system safety control structure was introduced by Leveson [Leveson, 2004]. A sample generic complex system safety control structure was shown in Figure 4. Each level or component of the control structure has roles and responsibilities in order to ensure that system safety constraints will be enforced throughout the lifecycle of a system. The second step of the STAMP-based risk analysis process involves modeling the safety control structure of a specific system. This section provides guidelines and templates to create customized safety control structures for the operation and/or development of a system. These guidelines are not meant to be universal rules, but rather a starting point to create useful safety control structures to be used as the basis for a complete risk management process.

3.1.1 LISTING RELEVANT ORGANIZATIONAL COMPONENTS

The first step in the creation of a system safety control structure is to decide what to include and what to leave out. This decision is analogous to defining the boundaries of the problem. It is critical to include in the analysis the actors who will impact safety throughout the system lifecycle. Choosing an analysis boundary involves an important tradeoff. A large boundary may improve the likelihood of discovering causal factors that cannot be found otherwise, at the cost of more analysis effort and resources. This section provides guidelines in listing relevant socio-technical system components to be included in the analysis. This is not a one-time step. Systems change over time, and the list of components included in the analysis has to be reviewed periodically - structural changes may have a large impact on the dynamics of the system. Structure impacts behavior, which itself causes structural changes, so the model structure and formulations must be revisited in a Plan-Do-Check-Act or PDCA (Deming) cycle fashion in order to ensure adequate risk identification and analysis results [Shewhart, 1939]. Both the model and the analysis must go through revision cycles, following a standard system safety process where hazards and analyses must be revised, updated and refined throughout the system lifecycle.




Figure 30: Deming's Plan-Do-Check-Act (PDCA) cycle

3.1.1.1 Using Org Charts

Organization charts are a good starting point for a list of components to include in the model. In every analysis we performed, organization structure documents and charts were an invaluable source of information to decide whether to include or leave out organization components.

For the development part of the system lifecycle, every component on the chart that has a direct impact and role to play in system design and integration, safety analyses, technology development, testing and manufacturing must be included in the model. For the operations part of the system lifecycle, every component on the chart that is involved in the funding, procurement and supply chain, production, throughput and capacity decision-making, safety decision-making, safety analyses and tracking, standards and processes definition as well as safety and quality assurance functions must be included in the model.

As an example, organization structures based on authority flows, such as that shown in Figure 31, should be used as the basis for the creation of the safety control structure. Typically, charts of the organization ultimately responsible for developing or operating the system should be used as a starting point. Organization charts are generally limited to components within a single organization. A complete safety control structure must include the technical and organizational components that impact system safety, as well as outside influences in the development or operating components that can impact safety. Consequently, organization charts must often be supplemented has to be extended to include outside components such as suppliers and contractors, regulatory and government agencies.

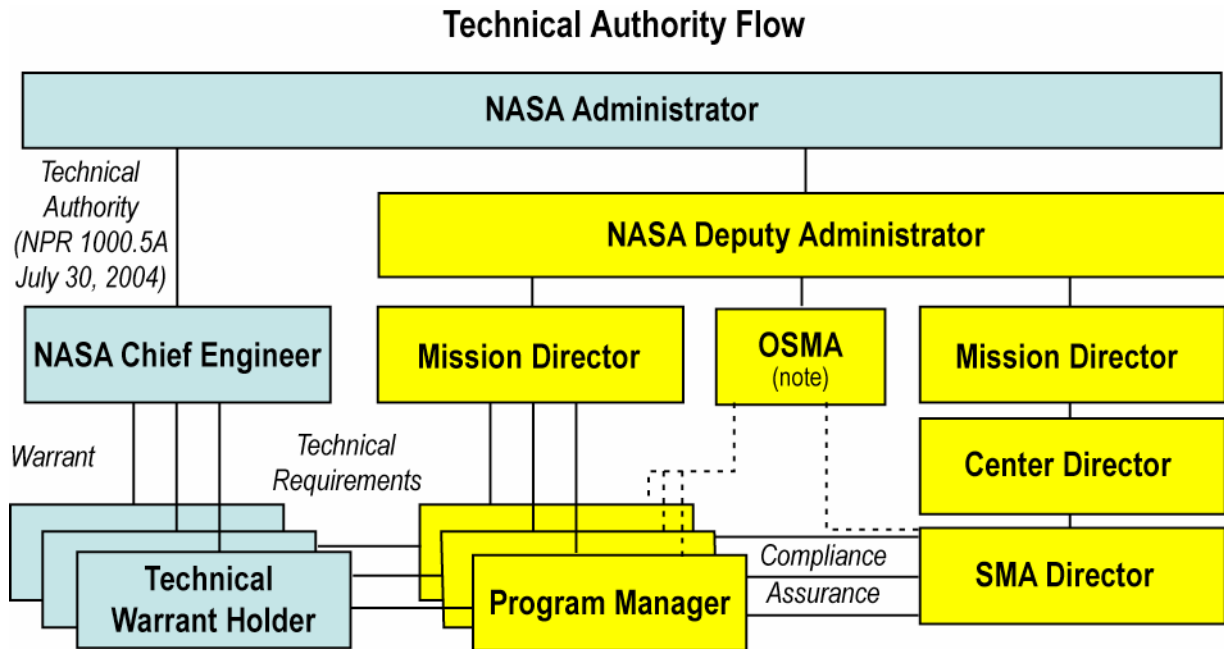


Figure 31: Technical Authority Flow (Adapted from [NASA, 2005])

3.1.1.2 Using generic STAMP structures

The generic control structure shown in Figure 4 can be used as an effective checklist to verify that important components have not been mistakenly left out of the model. For many systems, it may not be necessary to follow the hierarchical structure all the way up to the Congress and Executive components. As a rule of thumb, components should be included at least up to the level where funding originates for the system development and/or operation. For example, in NASA projects, it is necessary to include components all the way up to Congress and Executive, because this level defines the high-level mission objectives and provides Agency funding.

In addition to using generic structures as a checklist, it is possible to define generic components that inherit some characteristics and connectors of generic components. Customizing structures based on generic components and connectors is the first step toward the definition of a semi-formal process to facilitate the creation, analysis and monitoring of static safety control structures, as well as pave the way for the creation of dynamic risk management models. The list of generic components identified thus far is shown in Table 1. This list is not exhaustive and further generic components will be added as more analyses are

performed and more component types are identified. The purpose of this list is to facilitate the customization and characterization of real-life system components to create a custom static safety control structure and enable the creation of dynamic system components (see Chapter 4).

System Development	System Operation
Congress and Executive	Congress and Executive
Government Regulatory Agency	Government Regulatory Agency
Industry Association	Industry Association
User Association	User Association
Insurance Company	Insurance Company
Court	Court
Company Management	Company Management
Program and Project Management	Operations Management
Development Engineering	Operations Engineering
Development Safety Analyst	Operations Safety Analysts
System Evolution	System Maintenance
Development Assurance (Safety, Quality...)	Operations Assurance (Safety, Quality...)
Manufacturing	Supplier
Development Contractor	Operations Contractor

Table 1: List of Development and Operation Generic Components

In some cases, it will be trivial to associate a real-life system component to one of the generic components shown in Table 1. For example, the Federal Aviation Administration (FAA) is clearly a government regulatory agency that is involved in both the development and operation of air transportation systems. In other situations, the generic component type will be defined through the function of the component and its connections and interactions with other components. In still other situations, the component will be unique and may become the model upon which another generic component can be generated and added to the list in Table 1.

3.1.1.3 Using Interview Data

In addition to using organization charts and generic structures, another source of input for the generation of a list of structure components is from individuals within the structure itself.

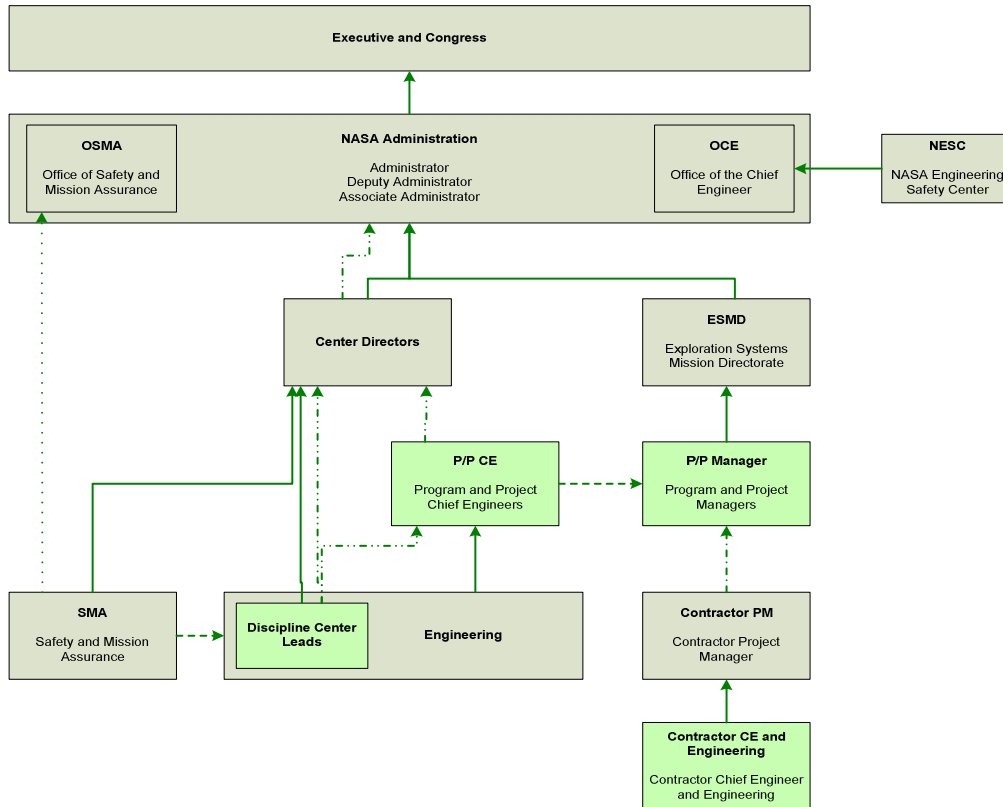


Figure 32: Structure of NASA Exploration System Mission Directorate (ESMD)

In an interview protocol (the complete interview protocol is provided in Appendix F) we developed to perform a risk analysis of the NASA Exploration Systems Mission Directorate (ESMD), we presented interviewees with a simplified structure of NASA ESMD (see Figure 32) and asked the following questions:

1. *Where does your position fit in this structure and what kind of expertise do you bring?*
2. *Focusing on your area of expertise, how would you describe the flow of resources and information across participants (in boxes) in the system (examples)?*
3. *How do you describe the role you play in safety during the development of the space exploration system?*

In refining the control structure, it is important to interview system participants within every box in the chart and to ask interviewees to explain the structure in their own terms. The emphasis should be on reviewing, improving and refining the structure, eliciting informal structural connections that are not represented in the official “party-line” organization chart and extending the boundary to components outside the organization, as necessary.

3.1.1.4 Inclusion Criteria

One of the problems of organization charts and interview data is that it will often generate a large list of components that could be included in the model. However, it will not, by itself, provide the information necessary to decide which components should be left out to keep the model size and complexity within practical bounds while still being able to perform detailed and useful analyses. Interview data, while extremely helpful, is often misleading because every interviewee wants his/her function to play a critical role in the model. In the end, it is necessary to decide which components will not provide much insight in the modeling analysis and should be either left out or combined with a more critical component.

In order to decide whether an organization component should be included in the model, the following questions can be asked:

For system development:

- 1- Is the component responsible for defining high-level system requirements and/or mission objectives and/or development schedule objectives?
- 2- Is the component responsible for funding decisions for the project or program?
- 3- Is the component responsible for enforcing schedule, budgets, and/or system requirements (including safety requirements) during development?
- 4- Is the component responsible for defining development standards and processes (especially safety-related standards and processes)? If so, does it have enforcement power?
- 5- Is the component responsible for, or heavily involved in, initial system certification?
- 6- Does the component have the knowledge and authority to halt or slow down system development when problems arise?
- 7- Does the component include a significant number of people working on activities such as technology development, safety analyses, system design, system integration, testing, and/or quality and safety assurance?

- 8- Is the component an important contractor of the main development organization, providing a significant portion of the organization's product and/or technical personnel?
- 9- Will the component be responsible for, or heavily involved in system evolution and upgrades?

For system operation:

- 1- Is the component responsible for or involved in defining criteria and metrics for system performance, production requirements, and/or mission objectives?
- 2- Is the component responsible for funding decisions for the system operation?
- 3- Is the component responsible for enforcing schedule pressure, budgets, and/or requirements (especially safety requirements) during system operation?
- 4- Is the component responsible for defining operation standards and processes (especially safety-related standards and processes)? If so, does it have enforcement power?
- 5- Is the component responsible for, or heavily involved in, system certification renewal or review?
- 6- Does the component have the authority to delay or stop production when problems arise?
- 7- Does the component include a significant number of people working on activities such as safety analyses, system evolution and/or upgrades, system maintenance, system integration, testing, safety and quality assurance?
- 8- Is the component an important contractor of the main system operator, providing a significant portion of the system hardware and/or personnel?

3.1.1.5 Combination Criteria

For many systems, it will be possible and desirable to combine multiple real-life components into "functional" model components that will conserve the critical functions and responsibilities of individual components while limiting model complexity. There are no universal rules to decide whether to combine different components, but some guidelines can be used. For example, when going through the previously introduced list of questions, if the answer to one specific question or sub-question is positive for multiple components, then combining these components should be considered. In other words, components that are structurally independent, but functionally similar, should be combined unless they receive funding from completely different or competing sources, or if they have competing

incentives. For example, if two certification agencies have different funding sources and competing incentives, they should be kept separate.

3.2 CONNECTING CONTROL STRUCTURE COMPONENTS

Once a list of components has been established, the components must be connected to ensure that they have the feedback channels necessary to be able to observe the controlled process. In addition, the components must have the control channels necessary to be able to control the process. These are the typical control theory observability and controllability conditions that must be met to ensure the safety constraints can be enforced. The observability and controllability conditions are necessary, but not sufficient to ensure the enforcement of safety constraints.

Many different types of generic connections can be used to connect components. Having a taxonomy of connections or relations between components has multiple purposes. By using a framework where different generic components are connected through various generic connectors, it is possible to create a formal structure made of nodes (components) and edges (connectors) that can form the basis for a mathematically analyzable control structure. Moreover, having multiple types of generic connectors allows building safety control structures using a “layered” approach, where a structure can be created and analyzed with respect to a limited number of connection types at the time. This allows the analyst to easily identify gaps and inconsistencies in the control structure, a process that can even be formalized and ultimately automated, if tools are built to analyze formal control structures. In addition, connecting various components using a taxonomy of generic connectors creates a library of generic connectors that will greatly facilitate the component-based dynamic model creation process introduced in the next chapter.

Using generic connections will also facilitate and provide the basis for automating portions of the lifecycle control structure monitoring process. Once critical connections have been identified using either the static control structure analysis criteria presented in section 3.3 and/or the results of the dynamic model simulations based on the control structure components introduced in chapter 4, it is important to monitor the health of critical control

structure connections at regular intervals to ensure the safety control structure remains in a condition where it can enforce safety constraints. In many accidents such as the water contamination accident of Walkerton, Ontario [Leveson, 2004], the control structure was initially adequate to control the process and ensure the enforcement of safety constraints, but under various dynamic pressures, some critical feedback channels became ineffective and were eventually lost, leading to a loss of observability in the control of water contamination. This hazardous state combined with specific environment conditions, including heavy rains and local area manure practices, resulted in bacterial contamination of the water supply, ultimately leading to seven deaths and over 2000 illnesses in the small community. The following section provides an initial taxonomy of generic connectors to help in creating, analyzing and monitoring static control structures and to ensure that safety constraints can be enforced throughout the system lifecycle.

3.2.1 GENERIC CONNECTOR TYPES

The use of these generic connectors is demonstrated here using the control structure of NASA's independent technical authority (ITA) as described in the ITA implementation plan [NASA, 2005]. The NASA-ITA control structure includes many interdependent components that all have responsibilities in ensuring safe system operation (see the components in Figure 34). The Agency chief Engineer, located at NASA Headquarters is the head of the ITA, and technical warrants are distributed to carefully selected, highly competent individuals within NASA. System technical warrant holders (STWHs) hold warrants for specific systems (e.g. shuttle system, space station, etc...) and Discipline Warrant Holders (DTWHs) hold warrants for specific disciplines (e.g., aerodynamics, structures, etc...). STWHs and DTWHs allocate tasks and responsibilities to selected Trusted Agents (STrAs and DTrAs) who are the "eyes and ears" of the warrant holders in the field. A complete list of Acronyms is provided in Appendix A.

The list of generic connectors presented in this section should not be seen as an exhaustive set, but rather as a working list that may evolve as new applications and structures are created.

3.2.1.1 Direct Report (Authority)



Figure 33: Direct report (authority) link between components

Direct report is the typical connector type used in organization charts. It shows official authority links between components, for example, a company president may have a half-dozen vice-presidents as his/her direct reports. In most situations, a component should have a limited number of direct reports. One outgoing direct report connection is typical for a standard hierarchical structure; two or more outgoing direct report connections are typical for matrixed structures. Components with more than two outgoing direct report links should be examined carefully for errors or inconsistencies. Multiple direct reports may lead to mixed loyalties and confusion in authority and control. For example, using the ITA structure example (see Figure 34), the connectors show that System and Discipline Technical Warrant Holders (STWHs and DTWH) report to the Agency Chief Engineer (OCE).

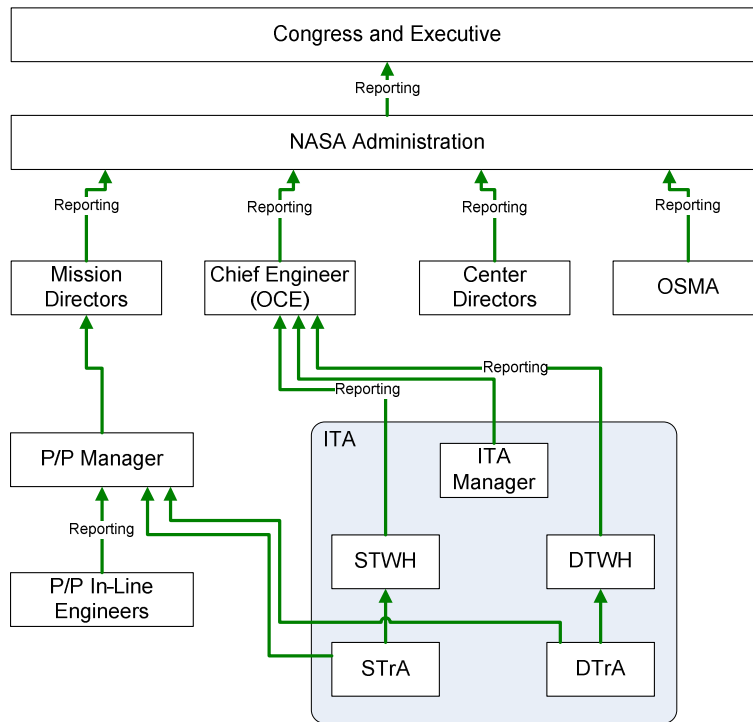


Figure 34: Direct report structure in the NASA ITA

3.2.1.2 Direct Oversight



Figure 35: Direct oversight link between components

Direct oversight is a generic connection where a component supervises the activities of another component. It does not necessarily involve authority of a component over another, but the oversight component should report to another component that has authority over the activities being performed, otherwise, even though the oversight component is responsible for overseeing another component, it does not have the authority to impose changes if necessary. A component can oversee multiple components. A component can also be overseen by multiple components, but this situation should be examined carefully as it may lead to ineffective oversight. For example, the Agency Chief Engineer (OCE) provides direct oversight to the Technical Warrant Holders (see Figure 36).

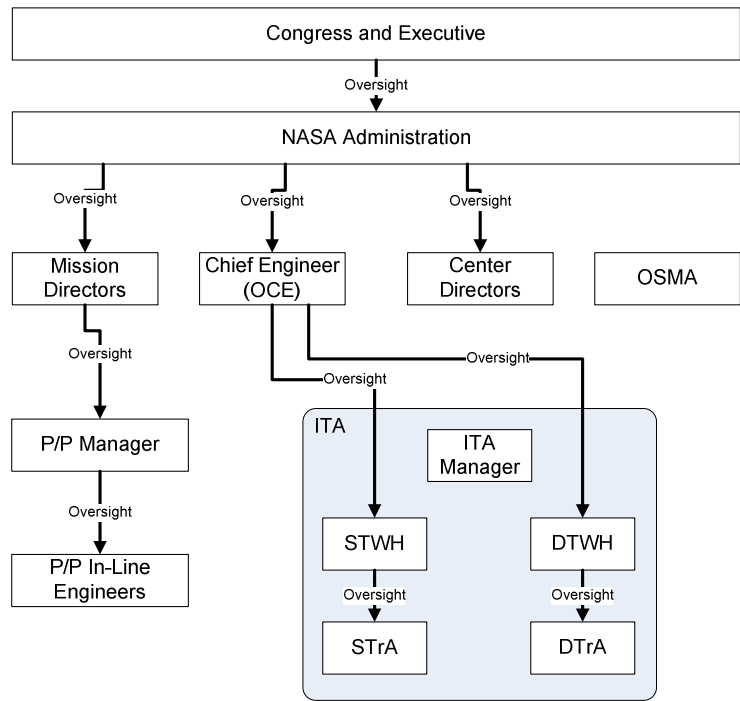


Figure 36: Direct oversight structure in the NASA ITA

3.2.1.3 Progress Report




Figure 37: Progress report link between components

The information report connector is at the basis of the typical feedback connection necessary for a component to observe the state of another component. The information report process can take the form of status and operations reports, informal reporting and information sharing, change reports, test reports, whistle blowing, incident reports, and reporting of analysis results [Leveson, 2004]. For example, System Trusted Agents (STrAs) provide progress reports to System Technical Warrant Holders (STWHs) and Program/Project Managers (see Figure 38).

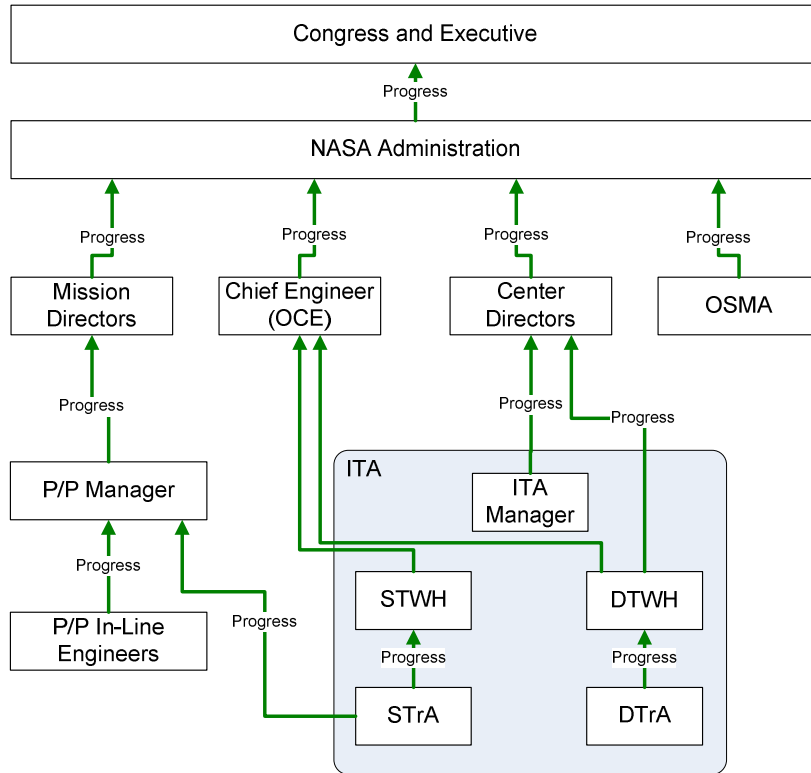


Figure 38: Progress report structure in the NASA ITA

3.2.1.4 Performance Appraisals




Figure 39: Performance appraisal link between components

The performance appraisal connector indicates the flow of performance appraisals between components. This connector type is important because performance appraisals provide incentives when they are linked to performance-based awards. When properly aligned, performance appraisal connections can have a positive impact on system qualities such as performance and safety, but when misaligned, they can lead to conflicts of interest, internal battles, and have a negative impact on desired system qualities. A component can receive performance appraisal from multiple other components. Similarly, a component can participate in the performance appraisal of multiple components. It is allowable, though not

always desirable, to have reciprocal loops in performance appraisals, when lower level components are asked to evaluate their superiors. This can also have an impact on incentives. Using the ITA example, as Trusted Agents (STrAs and DTrAs) are chosen from the technical workforce, they receive performance from P/P managers for the project tasks they perform.

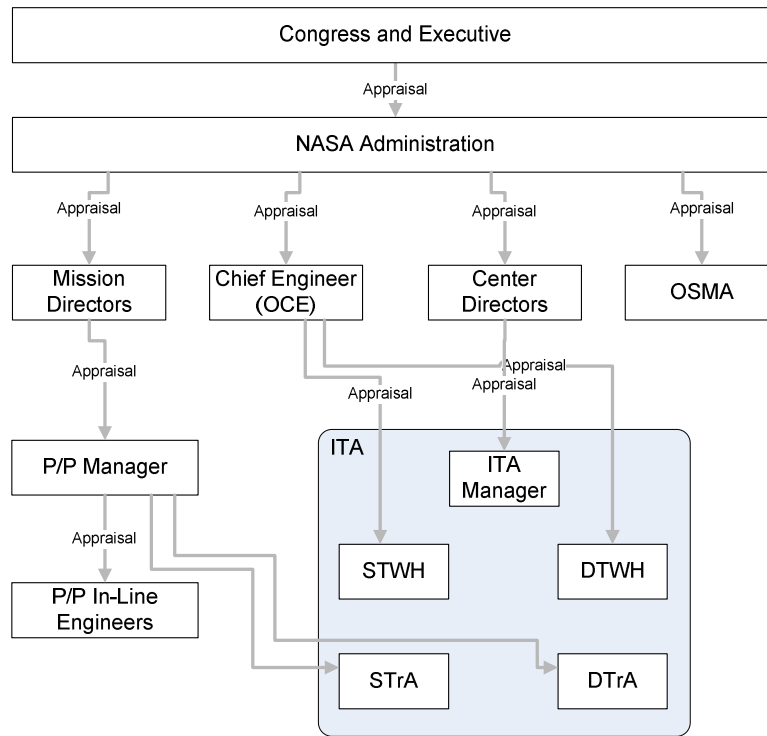


Figure 40: Performance appraisal structure in the NASA ITA

3.2.1.5 Resource Allocation



Figure 41: Direct resource allocation link between components

The resource allocation connector indicates the resource flow between components. A component can either provide funding, material, or human resources to another component. A component can receive partial resources from multiple other components. This is an important connector that should be followed from the main funding source(s) of the system all the way to (or through) every other component in the system control structure. Gaps in

resource allocation indicate an incomplete structure and should be examined carefully. For example, Trusted Agents receive funding from Technical Warrant Holders for the ITA tasks they perform and from the P/P managers for the project tasks they perform.

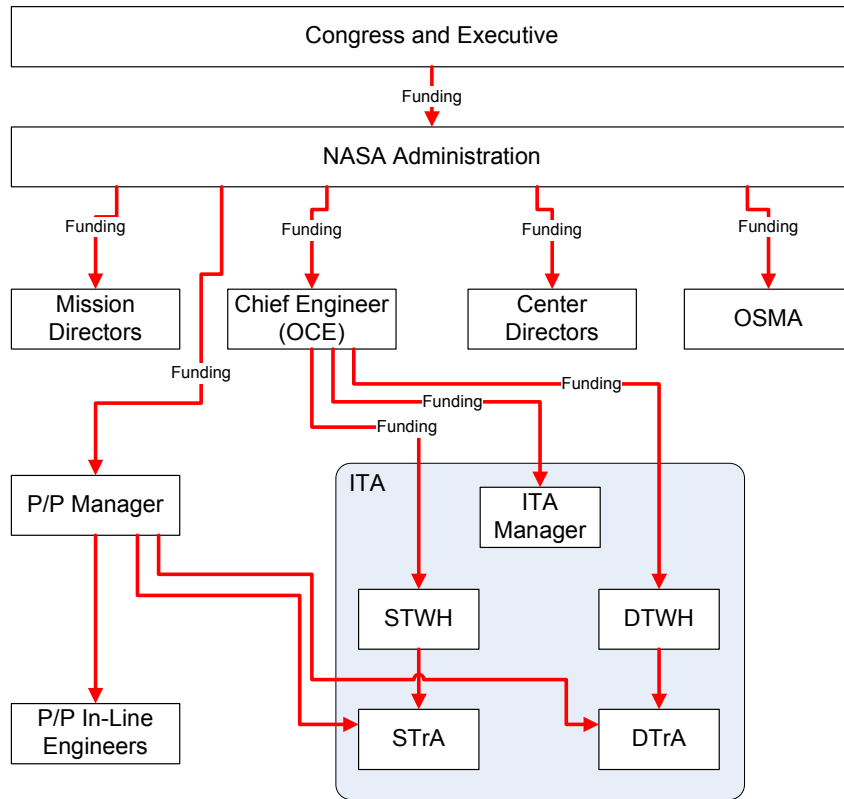


Figure 42: Direct resource allocation structure of the NASA ITA

3.2.1.6 Coordination and Technical Information Exchange




Figure 43: Coordination and technical information exchange link between components

The coordination connector is a non-directional connector that reflects the need for coordination between various components. Coordination and technical information should flow through this connector. Coordination is not a precise term, and the information that flows through this connector can overlap with that transferred through other connectors. The

importance of this connector comes from the need for a large amount and high quality of information sharing between various components of the system. Monitoring the strength of the coordination connection between critical components can be an effective way to measure and manage the level of information sharing and interaction within the system. Monitoring these parameters may be even more important for components that are separated physically and/or geographically.

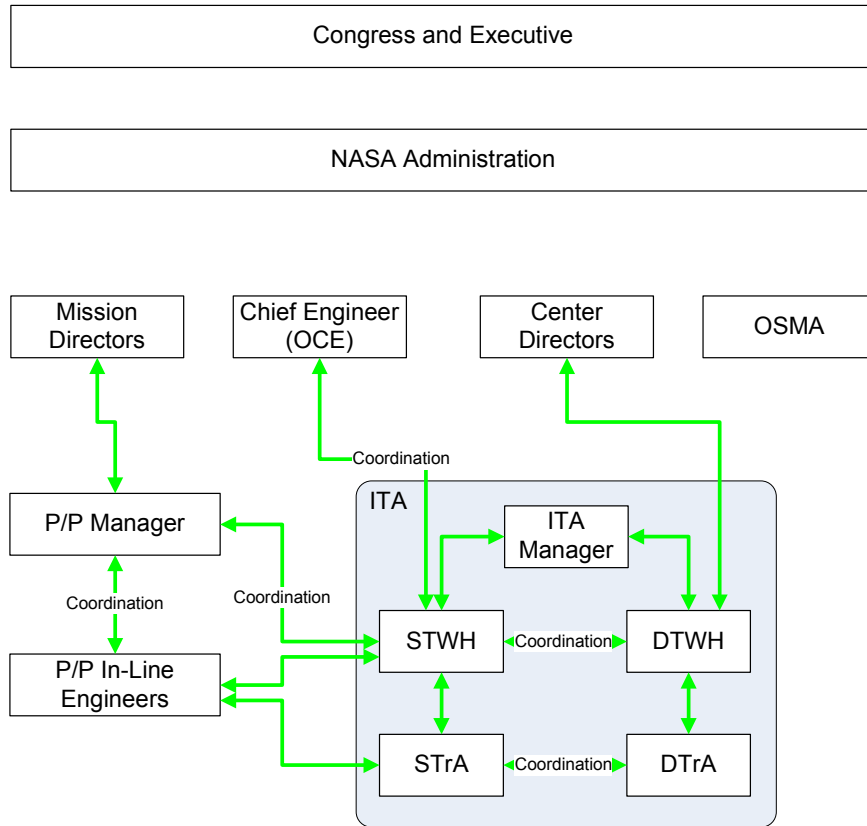


Figure 44: Coordination and technical information exchange structure of the NASA ITA

3.2.1.7 Physical Co-Location




Figure 45: Physical co-location link between components

The physical co-location connector is a non-directional connector that regroups components that have a high level of co-location. The exact amount of physical or geographical co-location indicated by the connector must be defined for each control structure, whether it means the components are located in the same office, in the same building, in the same city or else. One of the properties of this connection is that a transitive closure operator can be applied to the resulting graph to ensure completeness and verify consistency. In other words, if component B is co-located with component A and C, then component A is also co-located with component C (see Figure 46). This property check can be automated to ensure consistency of the structure. Co-location is an important property in some systems because co-located components tend to have more informal and richer interactions. Consequently, components that are not co-located may necessitate more monitoring and intense coordination, feedback and communication channels. In some cases, for systems distributed across many physical locations, the transitivity condition may not apply. The co-location relation may be defined as a continuum and not be completely discrete. Depending on to the definition of co-location chosen for the analysis (same office, building, city, state, etc...) it is possible that the transitivity condition may not apply. For example, if the physical distance between component A and B is small, and that between B and C is small, but the total physical distance between A and C is larger (being the sum of the distance between A-B and B-C), then the transitivity condition may break down. In this case, the transitivity condition does not apply, but can be still used to uncover problems in co-location (and communication richness) assumptions.




Figure 46: Using the transitivity connection relation

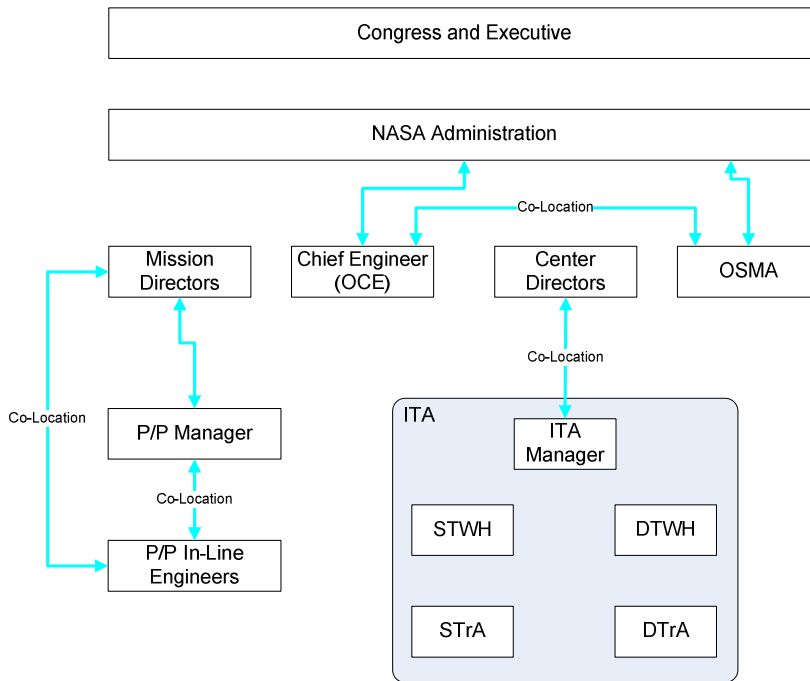


Figure 47: Physical co-location structure of the NASA ITA

3.2.1.8 Personnel Appointment



Figure 48: Personnel appointment link between components

The personnel appointment connector indicates that the management of component A appoints or participates in the appointment of the management of component B. In cases where the personnel appointment is a joint process between multiple components, a component can receive personnel appointment connections from multiple other components. It is not absolutely necessary to have a personnel appointment connector pointed to every single component, but the rationale should be made explicit when this is not the case.

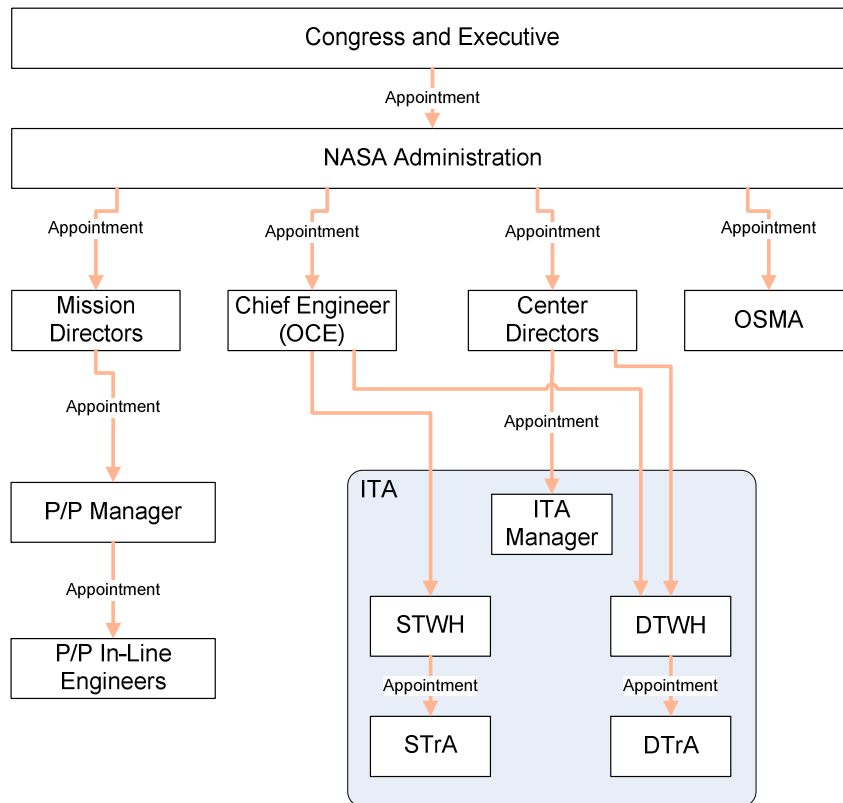


Figure 49: Personnel appointment structure of the NASA ITA

3.2.1.9 Procurement



Figure 50: Procurement link between components

The procurement connector indicates that an outside component that is not part of the main system development or operation organization provides material and/or human procurement services to another component in exchange for financial resources. For completeness, a component where a procurement connector originates should have a funding connector pointing to it. In the ITA analysis, it was decided not to include contractors explicitly within the analysis boundary. Consequently, no procurement connector was defined between components of the system. However, procurement services were used from support contractors embedded within the line engineering structure, and procurement hardware was

provided by outside contractors. In some cases, it may be useful to explicitly show the procurement connections coming from outside the analysis boundary.

3.2.2 DOCUMENTING ORGANIZATIONAL COMPONENTS

While generic connectors provide a precise and concise syntax that can be used as the starting point for building safety control structures, in most situations, the discrete syntax of generic connectors is not sufficient to understand the background, context and environment within which components and control structures operate. This section provides a list of documentation topics that may be required to allow a better understanding of the components, the control structure, and the risks associated with the development and operation of the system within its environment. Some examples are provided based on the documentation created for the System Technical Warrant Holders (STWH) component in the Independent Technical Authority (ITA) project.

3.2.2.1 Description

A high-level description of the component and its safety-related functions should be provided to quickly explain the purpose of the component and its role within the larger structure. If multiple smaller components or system members are aggregated in a larger component, the combination/aggregation rationale should be documented.

3.2.2.2 Overall Role and/or Mission Statement

This documentation topic provides a description of the role of the component (especially with respect to safety-related activities) as defined from the point of view of the organization within which the component operates. In many cases, a description of the role and responsibility of different business units in an organization will be readily available. A mission statement should also be recorded when available as it may provide insights into the function of a component from the point of view of management, and how this function may contribute to or hinder system safety. For the STWHs, the overall role was documented as: *“Primary ITA interface to one or more mission-related systems and representative of ITA for lower-level decision-making and actions. Ensure, as appropriate, the evaluation of technical issues and identification of risks through the use of existing engineering organizations.”*

3.2.2.3 Safety Responsibilities

Every component of the control structure has a role to play in ensuring that the constraints necessary to ensure safety are enforced throughout the system lifecycle. A complete STAMP analysis includes the allocation of safety responsibilities to each system component. In this section, the safety responsibilities for each component should be documented. In addition, a gap analysis should be performed to ensure that the official component responsibilities defined by the organization include as a subset every safety responsibility identified in the STAMP analysis [Leveson, 2005; Leveson, 2006]. A sample safety responsibility for the STWHs is to ensure the: *“production, quality, and use of Failure Modes and Effect Analyses, Critical Items List, trending analysis, hazard, and risk analyses”*.

3.2.2.4 Relevant Non-Safety Responsibilities

Non-safety responsibilities that may impact system development or operation cost, schedule and performance should be documented carefully as they indirectly impact system safety. For example, a non-safety responsibility for the STWHs is to: *“establish and maintain technical policy, technical standards, requirements, and processes for a particular system or systems”*.

3.2.2.5 Authority and Control

In addition to the generic connector indicating direct report, authority and control, it may be required to define more specifically the ways by which authority and control is realized, whether it is through requirements, regulation, standards, organization processes, norms, culture, as well as the information necessary to ensure compliance, that is, to be able to observe the state of other components in the system. For example, one of the ways the STWHs exert control on the system operation is through the: *“approval of hazard and risk analyses, signature on design reviews and CoFR, approval of hazard and risk analysis methodologies and definition.”*

3.2.2.6 Accountability

In addition to the documentation of a component responsibility and authority, it is necessary to “close” the responsibility-authority-accountability triangle by documenting the accountability of each component to ensure that its safety and non-safety responsibilities are fulfilled.

3.2.2.7 Resources (Funding, Material, Human)

In addition to the funding connector defined earlier, it is very useful to document the sources of available resources for each component. In addition, the conditions for resource allocation should be documented, as well as the uncertain factors that may affect the availability of resources over time for each component, if available. For example, funding for the STWHs and Trusted Agents is dependent on resources from the Office of the Chief Engineer (OCE). Changes in OCE responsibility or resources could affect STWH resources.

3.2.2.8 Hierarchical context

The hierarchical context within which the component operates must be documented as thoroughly as possible because it shapes decision-making within components. The hierarchical context should include the governing structure within the component, as well as the hierarchical path(s) within which the component operates. In addition to this, the appointment process for managers, and the promotion rules or process including tenure of managers should be documented.

3.2.2.9 Required Feedback

For each component, in addition to the feedback connector defined earlier, it is important to document the feedback requirements coming from each component, and going out to other components connected through a feedback link. This documentation is critical to evaluate the observability of various state variables. The feedback requirements are directly related to the safety control responsibilities. In general, every safety control responsibility (control action) should have a related feedback path to allow effective control to be exercised. Among other feedback, STWHs require program technical data, reviews, evaluations, analyses, tests, and process evaluations.

3.2.2.10 Mental Model Requirements [Leveson, 2004]

In this section, the mental model requirements of managers and employees within a component must be documented thoroughly. In order to ensure the enforcement of system-level safety constraints, each component needs to maintain and update a mental model of the functioning, purpose, and state of the system-as-a-whole, as well as a mental model of the functioning and state of each component with which direct interaction is required. The mental

model requirements documentation should include the feedback necessary for the component to create an accurate model of the system, its current state, and its functioning (i.e. the way it can change states and reacts to inputs and disturbances). For the STWHs, mental model requirements include the: *“Understanding of risk management process and expertise (including knowledge of uses and limitations) of tools such as FTA, PRA, FMEA/CIL and hazard analysis techniques.”*

3.2.2.11 Coordination and Communication Requirements [Leveson, 2004]

In this section, coordination and communication requirements between components should be documented, including the frequency, content of communications, typical communication media (direct interaction, email, phone, web-based, teleconference, videoconference...), meeting frequency, etc. These requirements should match the information identified in the previous section to create an accurate mental model of the component and system state and functioning.

3.2.2.12 Environment and Behavior-Shaping Factors (context) [Leveson, 2004]

In this section, information about the component environment should be collected, as well as information about how the environment and context, including the cultural context may shape behavior and decision-making within the component. For example, if the component operates within an environment and cultural context where speaking up and surfacing problems is discouraged through a tendency to “shoot-the-messenger” or to ignore problems, or if competition between divisions of a same company or agency is embedded within the organization culture and hinders communication between divisions, these factors must be documented as they will impact decision-making and may hinder the creation of accurate mental models of the system operation. For example, behavior shaping factors for the STWHs include performance pressures, career and performance appraisal factors, as well as the need to maintain credibility and influence to have an impact on system safety.

3.2.2.13 Potential Inadequate Control Actions [Leveson, 2004]

Steps 3, 4 and 5 of the STAMP risk analysis process provide some information on the safety responsibilities allocated to each component [Leveson, 2005], and help to identify and classify risks resulting from inadequate control actions. Potential inadequate control actions

for each component should be documented in this section, as well as requirements on the tracking and update of potential inadequate control actions. An example potential inadequate control action for the STWHs is the “*approval of inadequate safety and reliability engineering products (FMEA/CILs, hazard analyses, etc.)*”

3.2.2.14 Physical Location

The physical location of the component should be documented. If the operation of the component is distributed across multiple physical locations, then the distributed location should be documented and the need for additional coordination and communication should be reflected in the “Coordination and Communication” section.

3.2.2.15 Key member(s)

The key members (directors, president, managers, leaders, external relations, coordinators) of the component should be listed, as well as any special responsibilities they may have in addition to that necessary to fulfill component-level responsibilities.

3.2.2.16 Comments and Sources

Additional relevant information collected by analysts should be listed in this section, as well as documentation and other information sources upon which the analysis is based.

3.2.2.17 Monitoring and Tracking Requirements

In this section, requirements for tracking, monitoring, gap analysis, reviews and updates of the control structure model should be defined and the rationale for the requirements should be documented.

3.3 ANALYZING SAFETY CONTROL STRUCTURES

This section provides a summary of criteria that must in most cases be met by a complete and consistent control structure that will have the capability to enforce system safety constraints. These criteria can be used as guidelines for creating and analyzing safety control structures. Meeting the criteria is only one of the conditions necessary for an effective safety control structure. A control structure that does not meet every criterion is not necessarily “unsafe”,

but it should be examined carefully and the rationale for not meeting the criterion should be well documented and reviewed.

There are two types of criteria in this section. Completeness criteria support the analyst by providing a checklist of conditions related to various combinations of generic components and connectors. Completeness criteria are usually based on a single type of generic connection. Consistency criteria provide a repository of “best-practice” based on domain expertise and accumulated experience. This repository should be customized to the specific application domain of the system under analysis. This section provides a limited number of these criteria. The purpose is to define the syntax and provide a framework for analysts and domain experts to build their own repository of domain-specific criteria.

It is not possible to have a completely defined organization where every participant’s job description and responsibilities are formally defined. In many cases, if a criterion is not met by the formal structure but is critical for the effective operation of a system, participants will “fill the formal gaps” by creating a network of informal connections to compensate for weaknesses or gaps in the formal structure. A the problems with this informal “gap-filling” is that it is highly dependent on the tacit knowledge and experience of some system participants. If the key tacit knowledge and experience is lost, the gaps may “re-open” and problems will surface. An effective set of structure analysis criteria can help uncover some of those gaps, and mitigation strategies can be created to address the potential problems before they arise.

3.3.1 COMPLETENESS CRITERIA

3.3.1.1 Resource Criterion

The resource criterion states that every component of the safety control structure must receive resources (human, material, funding) from at least one other component of the control structure. In the case where a component does not explicitly receive resources from another component, a resource “source” must be defined and documented at the boundary of the system for completeness. For example, in Figure 42, the Congress and Executive component is the only component for which the source of resources is not documented. For

completeness, a resource “source” should be documented to explain where the original resources come from (federal taxes in this case).

3.3.1.2 Reporting Criterion

The first condition of the reporting criterion states that a very limited number of components at the top of the hierarchical structure do not report to any other components. In most control structures, there will be a single component at the top. For systems where different components are ultimately responsible for the system development and operations, configurations with multiple top components are possible, but should be examined and documented carefully. Distributed responsibility is a potential source of dysfunctional interactions and coordination problems leading to unsafe system behavior. If the control structure analysis boundary includes conglomerate-type extended enterprises, the top components should be sought at the level of government bodies or regulatory agencies that are responsible to provide the legislation and standards that regulate the safety of the system.

The second condition of the reporting criterion states that every component should be connected to at least one top component through a reporting link. This ensures that an official path (or process) to surface safety concerns exists, whether it is used or not.

3.3.1.3 Appointment Criterion

The appointment criterion states that every component but the top component(s) should have an appointment link pointing to it, i.e. the appointment process and mechanics for managers and leaders in each component should be made explicit and documented.

3.3.1.4 Appraisal Criterion

The appraisal criterion states that every component but the top component(s) should have an appointment link pointing to it, i.e. the performance appraisal process for managers and leaders in each component should be made explicit and documented.

3.3.1.5 Physical Co-Location Criterion

The physical co-location criterion simply states that if component B is co-located with components A and C, then components C and A are also co-located, as shown in Figure 46.

As mentioned earlier, if the co-location connectors are defined as a continuum, instead of a discrete relation, the co-location criterion can be used to test assumptions about the amount and richness of communication between closely located components.

3.3.1.6 Open Loop Criterion

The resource allocation, oversight, appointment, and direct report (authority) connectors should not form a closed loop circuit when connected to components (nodes) of the structure. The presence of a closed loop circuit most likely indicates serious dysfunctions or inefficiencies in resource allocation. More research will be required to test this criterion, but it has the potential to be extended to other types of connectors.

3.3.2 CONSISTENCY CRITERIA

Until now, the completeness criteria presented were based on a single connector type at the time. The consistency criteria presented in this section are based on an analysis using multiple types of generic connectors. The consistency criteria analysis is based on a superimposition of multiple layers of different types of generic connectors.

3.3.2.1 Co-Location and Coordination Criterion

The co-location and coordination criterion states that the supersets created through the union of physically co-located components should be connected through strong coordination and information transfer links. Superimposing these two types of generic connectors for the NASA ITA analysis lead to the structure of Figure 51. What can be learned from this exercise is that the NASA Administration-OCE-OSMA triad, located at NASA Headquarters, in Washington, DC, is only connected to the ITA structure through the OCE-STWH link. This indicates a potential information and coordination bottleneck between the Agency Chief Engineer and the System Technical Warrant Holders (STWHs) of the ITA. This potential bottleneck (where STWHs become overwhelmed with communication and coordination tasks that interfere with their ability to carry out other safety-related tasks and responsibilities) may not exist in the real system because of informal ties between the ITA and other HQ-based personnel such as members of the Office of Safety and Mission Assurance, which is responsible for laying down the Agency's safety and mission assurance requirements and

standards. However, the informal ties would still need to be examined for potential gaps and dysfunctional interactions.

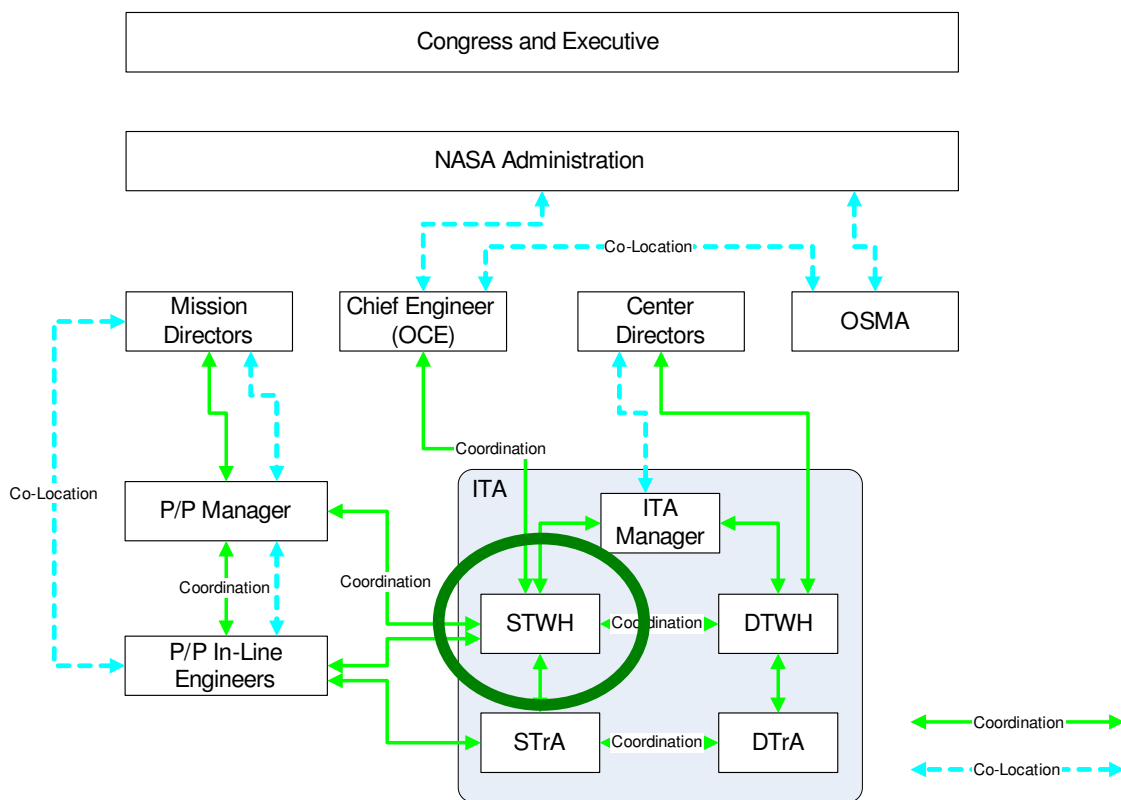


Figure 51: Superimposing physical co-location and coordination links

3.3.2.2 Observability

In control theory terms, observability is a condition necessary to be able to control a process. Observability addresses the issue of sensing and the ability of sensors to capture the dynamical behavior of a system [Belanger, 1995]. The formal definition of observability in linear systems theory, requires that the initial state of a system can be uniquely deduced from perfect knowledge of the inputs and outputs to and from the system for all t between time $t=0$ and time $t=T$. For linear time-invariant systems (LTIS), simple tests have been devised to test the observability property of a system. As we are dealing with complex socio-technical, and mostly non-linear systems that are neither easily linearized nor easily identified using mathematical control theory, the observability criterion has to be relaxed. For our current purpose, it is sufficient to define the observability criterion between Components A and B as the ability for a component A to be able to evaluate the state of a controlled component B

through information inputs and progress reports to component A. The information and/or progress report does not have to be direct, nor does it necessarily have to be formally defined. However, efforts should be made to monitor informal or indirect information transfers to ensure that system states remain observable by decision-makers. The observability criterion will be further defined and discussed when dynamic, component-based modeling is introduced in chapter 4. In the context of static control structures, it is sufficient to state that if component B is connected through a “direct report” link or chains to component A, then component A must be able to observe the state of component B through progress reports or information transfers links, whether they are formal or informal.

3.3.2.3 Controllability

Controllability is another criterion necessary to control a process. Controllability addresses the issue of actuation and the ability of actuators to control the state of a system [Belanger, 1995]. The formal definition of controllability, in linear systems theory, states that a system is controllable if there exists an input function that will make it possible to get the system from state 0 to an arbitrarily chosen state X at time $t=T$. Simple tests are available for controllability in LTIS. These tests are again not directly transferable to the type of systems we are dealing with for the same reason as observability tests. For the purpose of our analysis, it is sufficient to define controllability of component A over component B as the capacity of component A to sufficiently control or influence the behavior of component B through direct authority, oversight, processes, standards, and/or other incentives. The controllability criterion will be further defined and discussed in chapter 4 when dynamic component-based models are introduced. In the context of static control structures, it is sufficient to state that if component B is connected through a “direct report” link or chain to component A, then component A should be able to control the state of component B through authority links or incentives such as resource control and/or performance appraisals.

3.3.2.4 Oversight-Appraisal (OA)

The oversight-appraisal (OA) criterion states that oversight connections that are not accompanied by either a direct report or a performance appraisal link should be examined

carefully because there may be a lack of incentive for the overseen component to follow the guidelines and recommendations of the oversight component.

3.3.2.5 Resource-Report (RR) Criterion

The resource-report criterion states that components that provide resources to other components should also be connected to these other components either through a direct report or a progress report link.

3.3.2.6 Procurement-Funding-Oversight (PFO) Criterion

The Procurement-Funding-Oversight (PFO) criterion states that a component A that provides procurement services to another component B must be provided with oversight from component B, and must receive funding, ideally directly from component B, but at least from a component to which component B reports. More research and data collection will be required to validate this criterion, but from a control theory and system safety perspective, a lack of direct oversight of the safety-related tasks performed by procurement components is highly undesirable.

3.3.2.7 Multiple Loyalties or Reporting-Appraisal Criterion

The multiple loyalties (reporting-appraisal) criterion states that if a component A reports directly to multiple other components, say B and C, then components B and C should participate in the performance appraisal of A, otherwise, there may exist an incentive imbalance that will create problems. For example, while performing a risk analysis for the NASA ITA [Leveson, 2005], a static reporting-appraisal analysis of the control structure was performed and helped identify a possible case of appraisal imbalance in the control structure. As shown in Figure 52, the System Trusted Agents (STrAs) and Discipline Trusted Agents (DTrAs) have two reporting channels. They provide line engineering services and report to the P/P manager, and at the same time, they can be employed (and funded) by the System and Discipline Technical Warrant Holders (STWHs and DTWHs) to perform ITA-related tasks. When performing ITA-related and ITA-funded activities, the Trusted Agents report to the Technical Warrant Holders. However, only the P/P managers participate in the performance appraisal of Trusted Agents. This may create loyalty problems and imbalances when conflicts

arise either in the priority of trusted agents activities or when there are divergences of opinion between P/P Managers and Technical Warrant Holders.

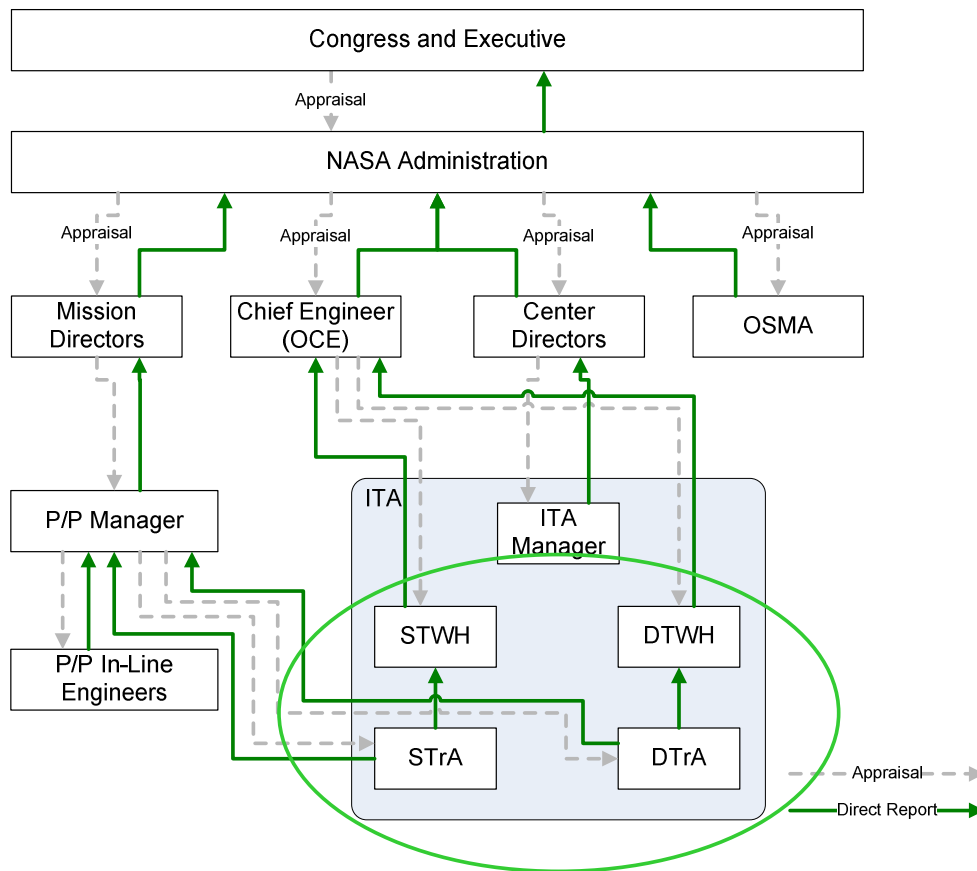


Figure 52: Superimposing performance appraisal and direct report (authority) links

3.3.2.8 Other Criteria

The purpose of the section is to provide a framework and syntax to create a set of completeness and consistency criteria to perform a first-order analysis and check of safety control structures. The list of generic components and connectors should be extended and updated, but it provides a foundation for a formal analysis of control structures. Similarly, the list of completeness and consistency criteria should be also be updated and extended. Future criteria to be developed include criteria to prevent: Conflict(s) of Interest, Misaligned Incentives, Lack of Safety Oversight, Lack of Safety Independence, Inadequate Safety Appointments, Inadequate Safety Funding, Inadequate Coordination (Overlap/Boundary), and Inadequate Safety Information.

3.4 FORMAL AND INFORMAL STRUCTURES AND CONNECTIONS

While many components of the system are often connected through a formal direct report or authority link, a typical command and control structure is only one of the ways safety constraints are enforced. The use of standards and processes, best practice, regulations, informal communication, reviews, testing, validation and verification processes, and even individual decision-making and behavior influenced by the organization culture can constitute effective methods to enforce safety constraints. On the other hand, poor processes and deficiencies in the safety culture of an organization can rapidly lead to an increase in risk through a resistance to the implementation and utilization of otherwise effective risk mitigation strategies.

Consequently, it should be understood that analyzing the formal connections are an important, but only small part of the risk mitigation equation. Softer factors must also be analyzed and included in the safety control structure, such as the attitudes, behavior and culture shared by employees, the accumulated knowledge and knowledge transfer about safety and risk mitigation, safety training, employees with leadership skills, status and credibility of safety and integration people, and the career and promotion process, among others, are even more important in ensuring the safety of a system over time. The inclusion of these factors in a safety control structure factors will be discussed in the next chapter (CHAPTER 4:).

3.5 NEW CONTROL STRUCTURES VS. EXISTING CONTROL STRUCTURES

For the creation and operation of new systems with little operational experience, it may be possible to start with a blank page, do a STAMP analysis, and create a control structure that will enforce safety constraints through the careful distribution and fulfillment of safety responsibilities across the various components of the system. Architecting a new system and safety control structure is the ideal situation, because it allows the mitigation of some risks from the very beginning of the system lifecycle, as well as the creation of a support system and organization with a heavy focus on maintaining the capabilities necessary to ensure safety. Newly created control structures may evolve over time, and the fact that safety constraints are enforced in the original structure does not mean that they will be enforced

throughout the system lifecycle. Consequently, control structures must be re-analyzed periodically to ensure their continuing capacity to enforce safety constraints.

In most cases, a new system will be developed and/or operated by an existing organization where a large portion of the control structure is already in place, including an existing network of formal and informal interactions within and across organizational boundaries, as well as organization culture(s) and norms that make it difficult, if not impossible, to design and implement a completely new safety control structure. The difficulties arise because of bureaucracy, organizational inertia, resistance to change, or political interests in maintaining the status quo. In these situations, a complete STAMP analysis must be performed to identify the hazards and safety constraints to be enforced. Then, an analysis of the existing control structure must be performed, as well as an analysis of the current safety responsibilities for each component of the structure. Finally a gap analysis must be performed to ensure the adequate modifications and additions to the existing control structures are performed to make it capable of enforcing system safety constraints. Whenever possible, modifications and additions to the control structure should be made while respecting the integrity, culture and norms already present in the organization. In some cases, it may not be possible, and major changes in attitudes and cultures will be required. Chapter 4 will discuss how some of the softer factors that make up organizational cultures and norms can be influenced and realigned to contribute to effective safety efforts in a control structure.

3.6 MONITORING SAFETY CONTROL STRUCTURES

As mentioned previously, complex socio-technical systems are dynamic and adapt over time to changes in their environment and under the influence of inevitable performance and resource pressures [Rasmussen, 1997; Leveson, 2004]. Because of these unavoidable changes, a control structure that was perfectly able to enforce safety early in the system lifecycle may not remain effective over time. A well designed and implemented control structure will have features and controls that provide robustness against unsafe changes. In addition, if safety requirements are well documented from the beginning, then potential changes can be evaluated before they are made to ensure they will not result in degraded safety. However, not all changes can be prevented and controlled, so it is necessary to re-

analyze control structures periodically to ensure the hazard control process can still be effectively performed. Unless a periodic system-level structure analysis and correction process is implemented, small well-intended changes made to improve local performance will increase the entropy of systems and make their structure vulnerable to small disturbances that may quickly degenerate to a major loss [Rasmussen, 1997; Leveson, 2004].

A complete STAMP risk analysis is used to identify the critical feedback and control channels required to enforce safety constraints. In order to ensure that these channels remain effective, some monitoring of the critical channels must be performed. One potential method to monitor safety control structures is through the analysis of social networks embedded in the control structure.

3.6.1 FROM COMPONENT TO INDIVIDUALS: MONITORING USING SOCIAL NETWORKS

In social network theory, relationships between individuals are described in terms of nodes (individuals) and ties (relationships). There may be different types of ties between nodes, some of them directional, such as authority ties, some of them non-directional, such as friendship ties. Custom connectors could also be defined to mirror the generic connectors defined earlier. A social network is built by creating a map of the individuals in the system and the ties they have with each other. The networks have various properties described in [Wasserman, 1994] including, among others: “cohesion” which is a measure of the degree of coupling between individuals in the system; “centrality”, a measure of the number of ties connecting an individual to others; “reachability”, a measure of the number of nodes reachable through ties by a specific individual; and “structural holes” [Burt, 1995], that represent potential for individuals to fill a strategic hole in a network by connecting two individuals that would not be linked otherwise.

The analysis of social networks was developed as an alternative to traditional sociological studies that focuses on the attributes of individuals in the network, while the focus of network theory is on how the individuals are connected with one another in a network. Social network analysis has lately emerged as a new lens to look at problems in fields such as modern sociology, anthropology, social psychology, information science and organizational studies [Freeman, 2004]. Research in these various academic fields has shown that networks operate

on many hierarchical levels, from baseball teams to the ties between nations and affect the way organizations are managed, problems are solved, power is attained, diseases are propagated, and individuals find jobs and become successful [Freeman, 2004].

The advent of information technologies on the workplace such as email, video conferences, web-based communities and training, instant messaging and “wiki” communities provides an opportunity to better monitor the informal connections between components of a system that are not depicted in the official organization charts. For example, a STAMP analysis may indicate that a strong coordination is necessary between component A and component B. A coordination link does not involve an authority or direct report connection, so it will not show in official documents, but by using social network analysis, and communication data from various information sources, it is possible to monitor over time the degree of coordination or connectedness between components of a system. Similarly, if a component A is supposed to be overseen by component B, but analysis of the social network shows a decline overtime of the connections and contacts between members of the component A and B, this can be interpreted as an indication of a potentially ineffective oversight link. This information can be used in a control structure analysis to identify weaknesses in the required feedback and control structure.

3.7 SUMMARY

In this chapter, a framework was introduced to facilitate the creation of useful static safety control structures. A preliminary methodology was presented to ensure that the control structures created are complete, and consistent. The framework introduced can also be used to perform semi-automated monitoring of static safety control structures in cases where some changes cannot be entirely prevented. This introduction was necessary to move on to the next chapter, where the focus switches from static control structure to dynamic control structures. A methodology is introduced to facilitate the creation of custom dynamic control structure models based on generic customizable control structure components.

CHAPTER 4: CREATING DYNAMIC RISK MANAGEMENT MODELS USING GENERIC CUSTOMIZABLE DYNAMIC COMPONENTS

In this chapter, a methodology is introduced to create, refine and validate dynamic risk management models. These models are used to identify and prevent dynamic patterns responsible for the migration of systems toward a state of higher risk. This chapter is tightly connected with the repository of generic dynamic components provided in Appendix E. This repository contains a selection of generic components based on two NASA risk analysis projects: an operation-centric risk analysis project focused on the impact of the Independent Technical Authority (ITA) on the safety of the space shuttle program, and a development-centric project focused on safety and risk analysis in the development of the new space exploration system. The components created for these projects were generalized to be used in future dynamic risk analysis and management projects. Model analysis and policy design and testing are addressed in Chapter 5. Chapter 6 provides a detailed case study example demonstrating the application of the methodology to the NASA Exploration System.

4.1 MODEL BUILDING METHODOLOGY

The methodology introduced is based on the assembly and customization of generic dynamic components. Currently available generic components are provided in the repository of dynamic generic components (see Appendix E). Some of these components will be used to illustrate the methodology in this chapter.

4.1.1 USING GENERIC STRUCTURES, COMPONENTS, AND ARCHETYPES

Generic structures are not a new topic. Senge popularized the use of generic system archetypes based on causal loop diagrams to analyze and find non-intuitive solutions to problematic situations [Senge, 1990]. Wolstenholme also supports the use of qualitative generic structures based on causal loop diagramming [Wolstenholme, 1990]. He proposes a framework to classify system archetypes using Senge's archetypes as an example. The ultimate objective is to define a set of core archetypes to be used in system dynamics

modeling and analysis [Wolstenholme, 2003]. Marais [Marais, 2003] customized and extended Senge's archetypes for application to the field of organizational system safety, demonstrating the usefulness and ease of customization of quality generic archetypes. Paich goes one step further toward the definition of formal generic structures created using stock and flow structures [Paich, 1985] as building blocks.

The creation of large dynamic models requires the use of various information sources, including quantitative and qualitative data, as well as interaction with domain experts, usually during small group model building sessions led by a modeling team including a moderator and modeling experts. Creating useful models is a time and resource consuming process that is required because every system and problem is different, so models cannot be easily reused for another application. Generic structures can facilitate the model creation process by reusing parts of a model structure that create behavior patterns that are domain-specific, rather than system-specific.

4.1.2 COMBINING STAMP SAFETY CONTROL STRUCTURES WITH CAUSAL LOOP STRUCTURES

The model-building methodology described in this thesis combines the use of generic system dynamics-based structures and STAMP control structures to facilitate the dynamic risk management model building, validation, and analysis. The generic system dynamics structures are embedded within STAMP control structure components to create generic dynamic STAMP components that can be individually customized, tested and validated before they are combined into a complete STAMP-based dynamic control structure model. The model-building and analysis methodology is iterative and based on the typical system dynamics modeling cycle defined by Sterman (See Figure 53).

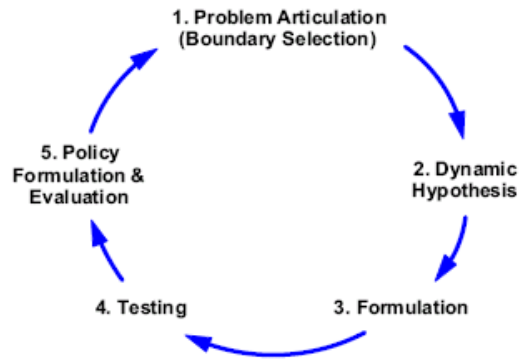


Figure 53: The System Dynamics Modeling Cycle [Sterman, 2000]

An alternative feedback-based view of the model building, scenario analysis, and policy development process is provided in Figure 54. The model is refined and validated as more data becomes available from the analyses performed and from measurement of the system outputs. Model validation and refinement is a continuous iterative process. The methodology necessitates the involvement of system stakeholders and is heavily participative, in the tradition of action research [McKernan, 1996]. Typically, modeling projects often involve the use of consultants acting as facilitators, model-builders, and analysts. Consultants usually approach a problem along the sequence of Schein’s classical process consultation prescription of “engagement-analysis-action-disengagement” [Schein, 1969]. This consultation process is not adequate for lifecycle risk management in complex systems because inevitable changes in the safety control structure may erode its efficacy over time. Consequently, control structures must be analyzed periodically and tracked over time to ensure their continued effectiveness and to keep risk to acceptable levels. One-time consultant engagements are not designed to perform this function. The purpose of the component-based methodology introduced here is to facilitate the model building and analysis methodology to a point where it can be performed by system stakeholders, including engineers, managers, and safety analysts with acceptable levels of training.

Various information sources are used to create the models. Forrester points out that three types of information sources are typically used to create dynamic simulation models: 1) numeric data, 2) textual accounts, and 3) mental models of modelers and others involved in the model building process [Forrester, 1992]. Forrester adds that mental models are the richest source of information for modeling, while numeric data usually accounts for a small

fraction of the data used in modeling. The model will never be perfectly valid, nor will it be entirely complete. As Forrester convincingly argues [Forrester, 1985], the process involved in creating, modifying, validating and testing the model as well as analyzing and interpreting results from experiments, scenario analysis and policy testing may be more valuable and insight-provoking than the simulation results themselves. The methodology presented in this chapter follows this principle and is heavily rooted in insight-creation from both the model and the modeling process.

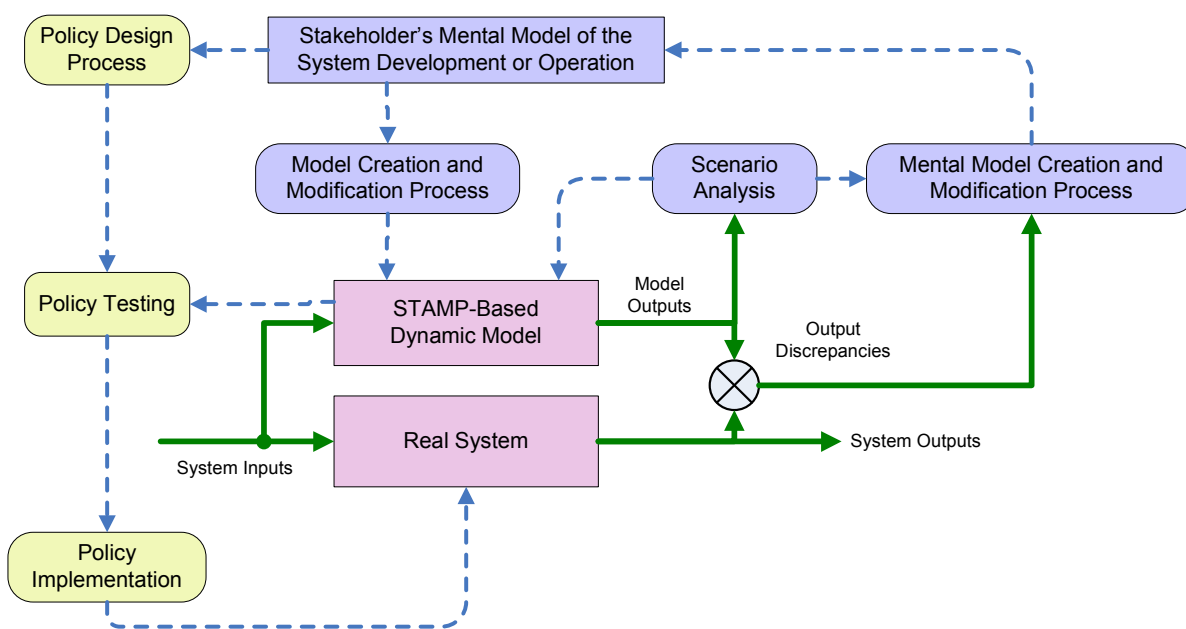


Figure 54: Modeling, Scenario Analysis and Policy Development as a Feedback Process

4.2 CREATING DYNAMIC RISK MANAGEMENT MODELS BASED ON STAMP AND SYSTEM DYNAMICS

In this section, the focus is on activity 6 of the STAMP-based risk analysis process (See Figure 55), that is, system dynamics modeling and analysis based on the STAMP safety control structure and generic dynamic components. The example used to illustrate this

model-creation methodology is the operations-centric model created for the NASA ITA analysis. Policy design, analysis and testing are addressed in chapter 5.

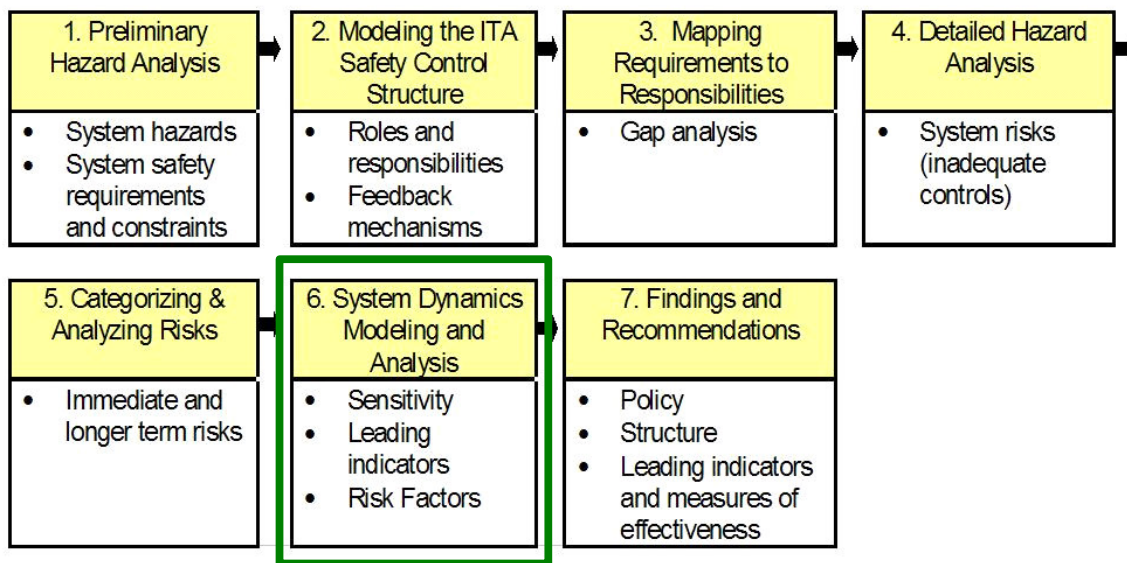


Figure 55: Focus on Step 6: STAMP-Based System Dynamics Model Building

Creating and analyzing dynamic risk management models is an iterative multiple-step process in itself. The modeling part is discussed in this chapter and involves using the information gathered in the previous steps of the STAMP-based risk management process to characterize the system and build the dynamic models that will allow further analysis. The analysis part of the process is covered in chapter 5 and also includes multiple steps that will be dependent on the specific objectives of the analysis and on the specific characteristics of the system. Analysis for the development of a new system is different from analysis of an existing operational system. The model creation process is summarized in the flowchart of Figure 56 using a waterfall-like structure. The process flowchart should be used as a guide to create customized dynamic risk management models. When problems arise at any step of the process, backtracking and iterations are necessary to correct the situation.

The steps shown in the flowchart almost never follow a strict sequential order. In fact, the process is much closer to a continual feedback and convergence process than a linear, sequential process. An alternative convergence-based representation of the flowchart of Figure 56 is shown in Figure 57. This convergence process should continue (around Loop C

in Figure 57) until the model captures adequately the causal structure, relationships and decision-rules used in the real system. In this case, given the same inputs, a good correlation should exist between outputs from the model and the real system. If the correlation is poor, a mismatch between the model and the real system has to be resolved. It is likely that the model causal structure and/or decision rules do not accurately represent those of the system, or that critical feedback loops were omitted. In many cases, this mismatch will be discovered during the analysis phase, when scenarios are created and executed on the model. Consequently, it may (and most likely will) be required to backtrack from model analysis (see link A in Figure 57) to modify the causal structure and decision rules of some components. Nevertheless, following the model creation methodology outlined in this section facilitates and increases the robustness of the model-building process, preventing and lessening the impact of backtracking during scenario analysis.

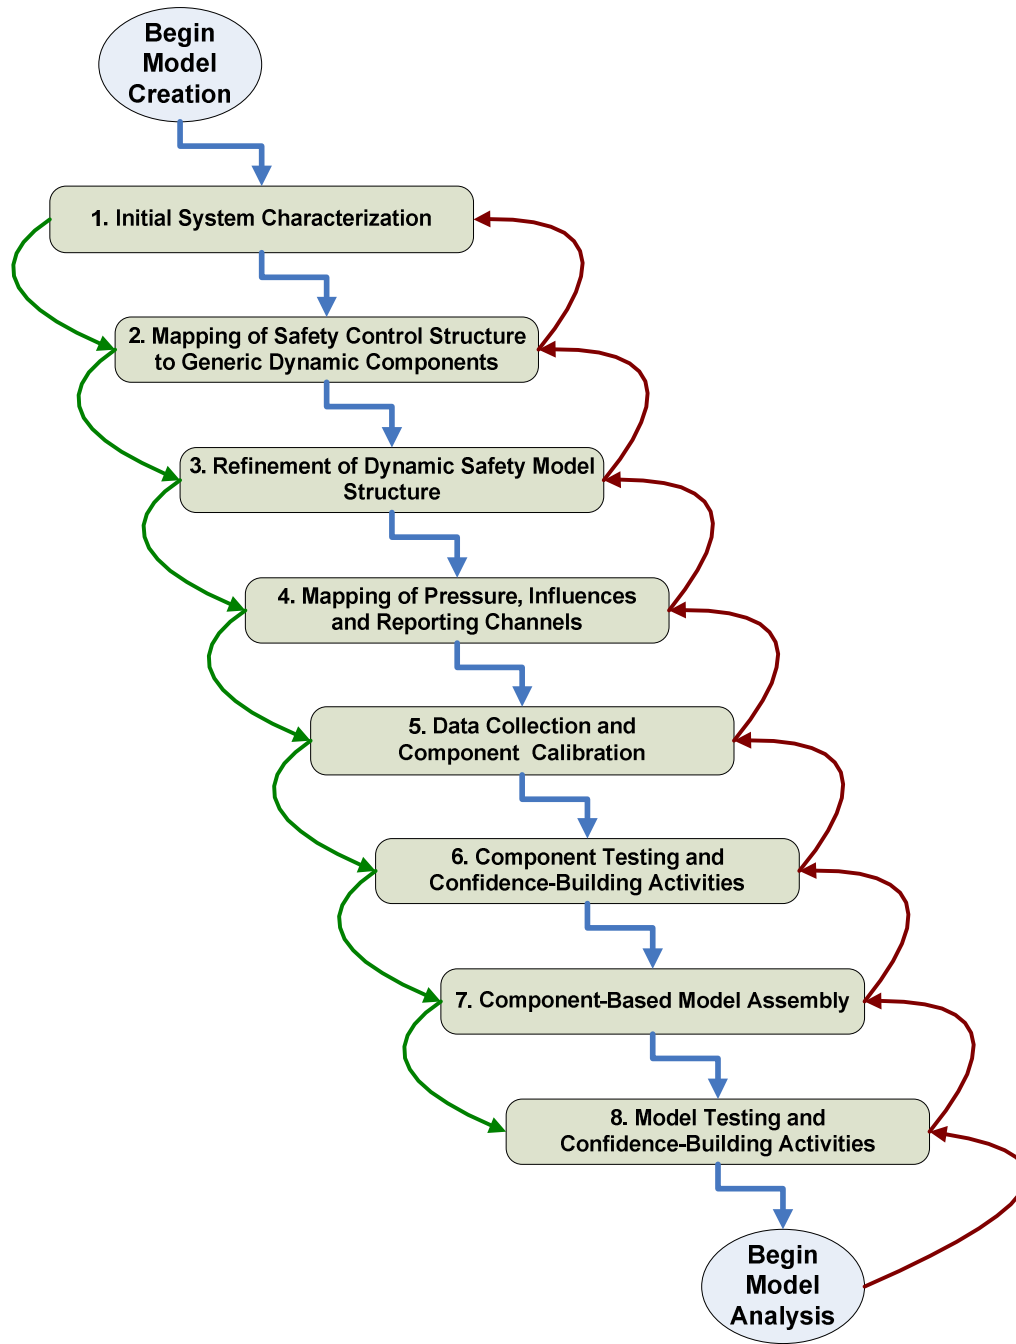


Figure 56: Summary of the component-based model building methodology

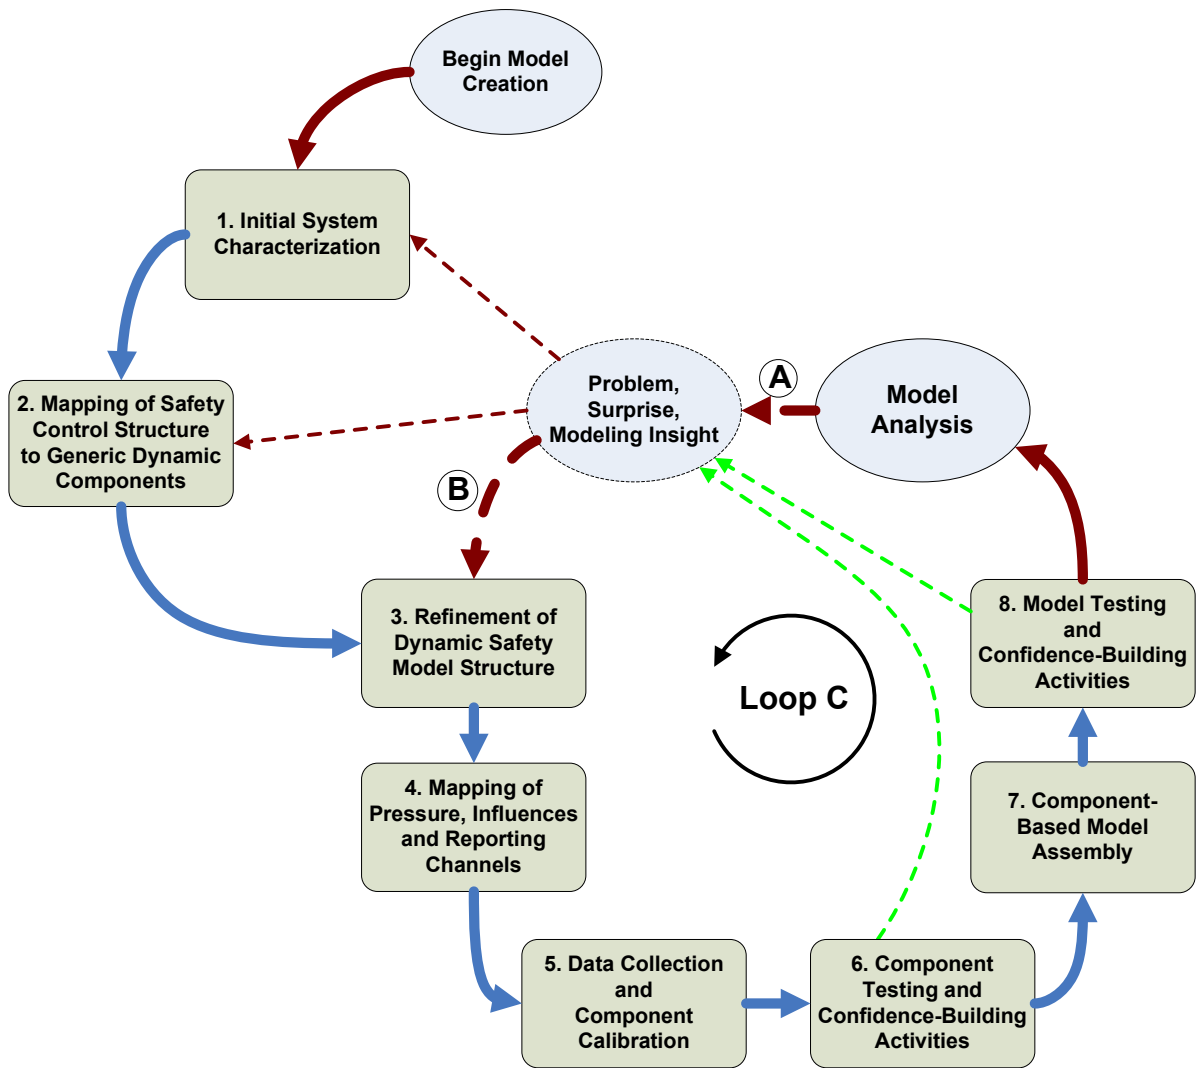


Figure 57: Alternative Convergence-Style Flowchart Structure

4.2.1 STEP 1: INITIAL SYSTEM CHARACTERIZATION

Before the modeling begins per se, it is useful to perform a first order characterization of the system under analysis. The first order characterization does not need to have precise quantitative values associated with it. In most cases, a simple rationale for the characteristic value will be sufficient. This characterization will be useful in subsequent modeling steps and serves multiple purposes. It allows the selection of adequate generic components, as well as the first order selection of exogenous parameters for rapid prototyping of the executable components and integrated models. The list of characteristics is evolving and will be extended and modified as more models are created in various application domains. The

following sections provide an initial list of system characteristics for system development and/or operations along with possible values and a short description. Depending on the type of system modeled, some characteristics will not be useful and/or relevant.

4.2.1.1 Production/Throughput Units

This parameter represents the primary performance measure of the system under analysis. In the case of the NASA ITA, the production unit was defined as shuttle launches per time period. In other production systems, it could be widgets produced per time period, or tons of coal mined per time period, or aircraft taking-off per time period.

4.2.1.2 Initial System Lifetime

This parameter defines the length of the system lifecycle as decided at the beginning of the system operation. The decision to extend or shorten this time period can be taken later during system operation.

4.2.1.3 Tight Coupling

As mentioned previously, tight coupling (where parts of the system are tightly linked to many other parts and therefore a change in one part can rapidly affect the status of other parts) is one of the dimensions defined by Charles Perrow that influences the occurrence and consequence of major accidents in complex systems [Perrow, 1999]. It is very difficult to quantify the tight coupling characteristic. Perrow created a chart for relative ranking of tight coupling of various systems (shown on the vertical axis in Figure 58). However, the chart is not completely accurate because the “tightness” of coupling has more to do with the design and operating characteristics of the system than on the specific system application. A nuclear power plant, for example, could be designed as tightly coupled or loosely coupled.

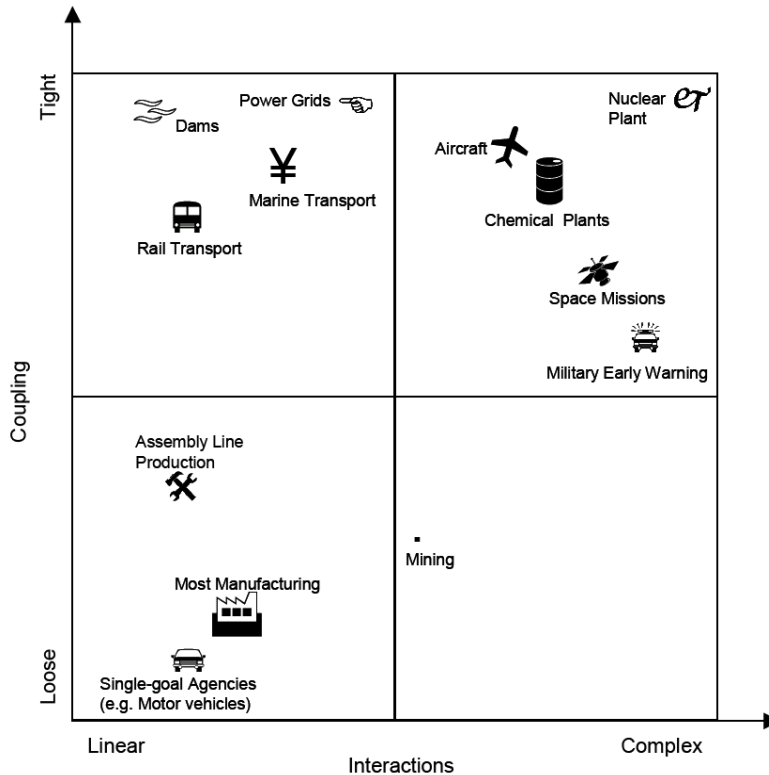


Figure 58: Coupling and Interaction Complexity Chart (adapted from [Perrow, 1999])

4.2.1.4 Interactive Complexity

Interactive complexity (the presence of unfamiliar or unplanned and unexpected sequences of events in a system that are either not visible or not immediately comprehensible) is another dimension defined by Charles Perrow that influences the occurrence and consequence of major accidents in complex systems [Perrow, 1999]. Just as for tight coupling, it is difficult to quantify the interactive complexity characteristic. Moreover, engineers and system designers have a large influence on the level of interactive complexity exhibited by a system, so any ranking based on system application only should be seen as highly subjective and debatable.

4.2.1.5 Maintenance Requirements

The system maintenance requirements are defined as the planned maintenance work required annually to operate the system at its nominal production rate. The maintenance requirements are normalized with respect to the maintenance requirements at initial system deployment.

4.2.1.6 Refurbishment Schedule

The refurbishment schedule of a system allows for possible system overhauls, evolution and upgrades that go above and beyond normal maintenance. The refurbishment schedule is defined as a function of the initial planned system lifecycle.

4.2.1.7 Accident Severity (Negligible, Marginal, Critical, Catastrophic)

The severity or impact of accidents is measured in terms of the loss associated with the accident. Losses are defined with respect to worst case scenarios arising from the hazard realized into an accident. The loss can be defined in terms of human, equipment, mission and/or environmental damage. The exact figures are debatable, and different government agencies (NASA, DoD, FAA...) have different definitions. A typical severity scale used by the Department of Defense [DoD, 2000] is shown in Figure 59. An alternative scale for accident severity based on DoD and NASA standards was introduced in [Dulac, 2004].

Description	Category	Environmental, Safety, and Health Result Criteria
Catastrophic	I	Could result in death, permanent total disability, loss exceeding \$1M, or irreversible severe environmental damage that violates law or regulation.
Critical	II	Could result in permanent partial disability, injuries or occupational illness that may result in hospitalization of at least three personnel, loss exceeding \$200K but less than \$1M, or reversible environmental damage causing a violation of law or regulation.
Marginal	III	Could result in injury or occupational illness resulting in one or more lost work days(s), loss exceeding \$10K but less than \$200K, or mitigatable environmental damage without violation of law or regulation where restoration activities can be accomplished.
Negligible	IV	Could result in injury or illness not resulting in a lost work day, loss exceeding \$2K but less than \$10K, or minimal environmental damage not violating law or regulation.

Figure 59: Suggested mishap severity categories (From MIL-STD-882D [DoD, 2000])

4.2.1.8 Length of Pause after Accidents

The length of time required to bring the system back to a nominal operational state after a loss impacts the dynamics of the system. It ranges from none, such as a car accident that only slows down the throughput of the system, to long, such as the time required to fix problems on a space shuttle (over two years after the Columbia accident).

4.2.1.9 Outsourcing Ratio

This parameter reflects the fraction of hardware and personnel necessary to operate a system that is outsourced through contracting agreements.

4.2.1.10 Employee Attrition and Turnover Fraction

This parameter indicates the historical average of employee attrition and turnover for the organization in charge of direct system operation or development. The fraction is defined in terms of fraction of total employees per month. When a significant portion of the work is outsourced, the average attrition and turnover rate of contractors should also be recorded and documented.

4.2.1.11 Quality and Availability of Lessons Learned

Lessons learned allow the avoidance of past mistakes and repeat-history accidents. Organizational learning is discussed in the organizational risk theory section of Appendix C. Many difficulties can hinder learning, especially for new systems with little operational experience, systems with no problem reporting process, systems where problem reporting is discouraged or systems where problem investigation is deficient. High-quality lessons learned may exist but be unavailable because they are buried in gigantic opaque databases or in the form of tacit knowledge possessed by employees who may be nearing retirement, in which case, the knowledge may be lost.

4.2.1.12 Requirements Waivers Allowed

In some organizations, requirements waivers may be granted to allow the system to operate despite some requirements (including safety requirements) not being fulfilled. At the time of Columbia, over 3200 high-criticality (1/1R) items had associated requirement waivers [Gehman, 2003].

4.2.1.13 Dedicated Safety Organization

Large systems often have a component of the organization dedicated primarily to system safety activities safety analyses. In other systems, safety may be integrated with other engineering and maintenance activities.

4.2.1.14 Independence of System Safety Organization

If safety-related decision-making has a dedicated reporting path up to the highest decision-maker, the safety organization is said to be independent.

4.2.1.15 Geographic Dispersion of Organization

Organizations where operations are physically co-located have low geographic dispersion. Organizations where similar systems are operated in different countries or where centers in different locations are responsible for different operation phases exhibit high geographic dispersion.

4.2.1.16 Political Uncertainty

Political uncertainty affects the stability of national objectives. Systems depending on government funding can be affected by changing national objectives. There are many sources of political uncertainty, including uncertainty in local political dynamics, the number of earmarks and specific budget allocation constraints, as well as foreign policy obligations.

4.2.1.17 Criticality of Program

Government-sponsored Programs that are critical to current national objectives have a lower likelihood of getting cancelled, or being the victim of budget cuts.

4.2.1.18 Leadership and Vision for Program

A clear and precise vision for a government-sponsored program reduces uncertainty and likelihood of cancellation and budget cuts.

4.2.1.19 Coherence and Consistency of Program Objectives and Policies

A program exhibiting high coherence and consistency in policies and objectives reduces the likelihood of misinterpreted requirements and requirements changes, as well as unrealistic

production objectives. It may also improving the morale of program employees as requirements are more stable and allow employees to bring their projects to completion before moving to something else.

4.2.1.20 Congress and Executive Ability to Market Program

The ability of congress and the executive to market a program to their constituents affects the support allocated to the program, and thus the stability of the program and of its funding.

4.2.1.21 Congress and Executive Risk Tolerance for Program

Risk tolerance of congress and the executive affects how they react to an accident or major incident. High risk tolerance will prevent a major re-evaluation of the program's objectives and resources after a major incident or accident, while low risk tolerance may prompt such reassessment.

4.2.2 STEP 2: MAPPING OF STATIC SAFETY CONTROL STRUCTURE TO GENERIC DYNAMIC COMPONENTS

The objective of this step is to map each components of the static safety control structure (created earlier in the STAMP-based risk analysis process) to associated generic dynamic component(s). A partial list of generic components created and stored in the generic component repository is provided in Table 2. In some cases, new dynamic components will have to be created (and added to the repository of generic components) to ensure that every critical component of the safety control structure is represented in the dynamic model. Usually, it is neither necessary nor useful to use every type of generic dynamic component. For example, if congressional and executive decision-making has a negligible effect on the operations of the system, the Congress and Executive component can be left out. Chapter 3 provided some guidelines and criteria to decide which components should be used in creating a static safety control structure. If the guidelines are followed, mapping the static control structure to generic dynamic components should be trivial. Otherwise, the boundary of some components may have to be redrawn to ensure the behavior and influences of each component are properly included. Figure 60 shows a mapping of the initial ITA structure created during the ITA risk analysis [Leveson, 2005] to generic dynamic components taken from the

repository of generic components (see Appendix E). The mapping was performed by matching the overall ITA structure, as well as the individual responsibilities of each ITA component (as documented in the ITA implementation plan) to the structure and processes of individual generic operations components in the repository. For example, the responsibilities of the Congress and White house components in the ITA structure match those of the Congress and Executive generic component (see Figure 60). Table 2 provides a summary of the mapping between control structure components and generic components takes from the repository.

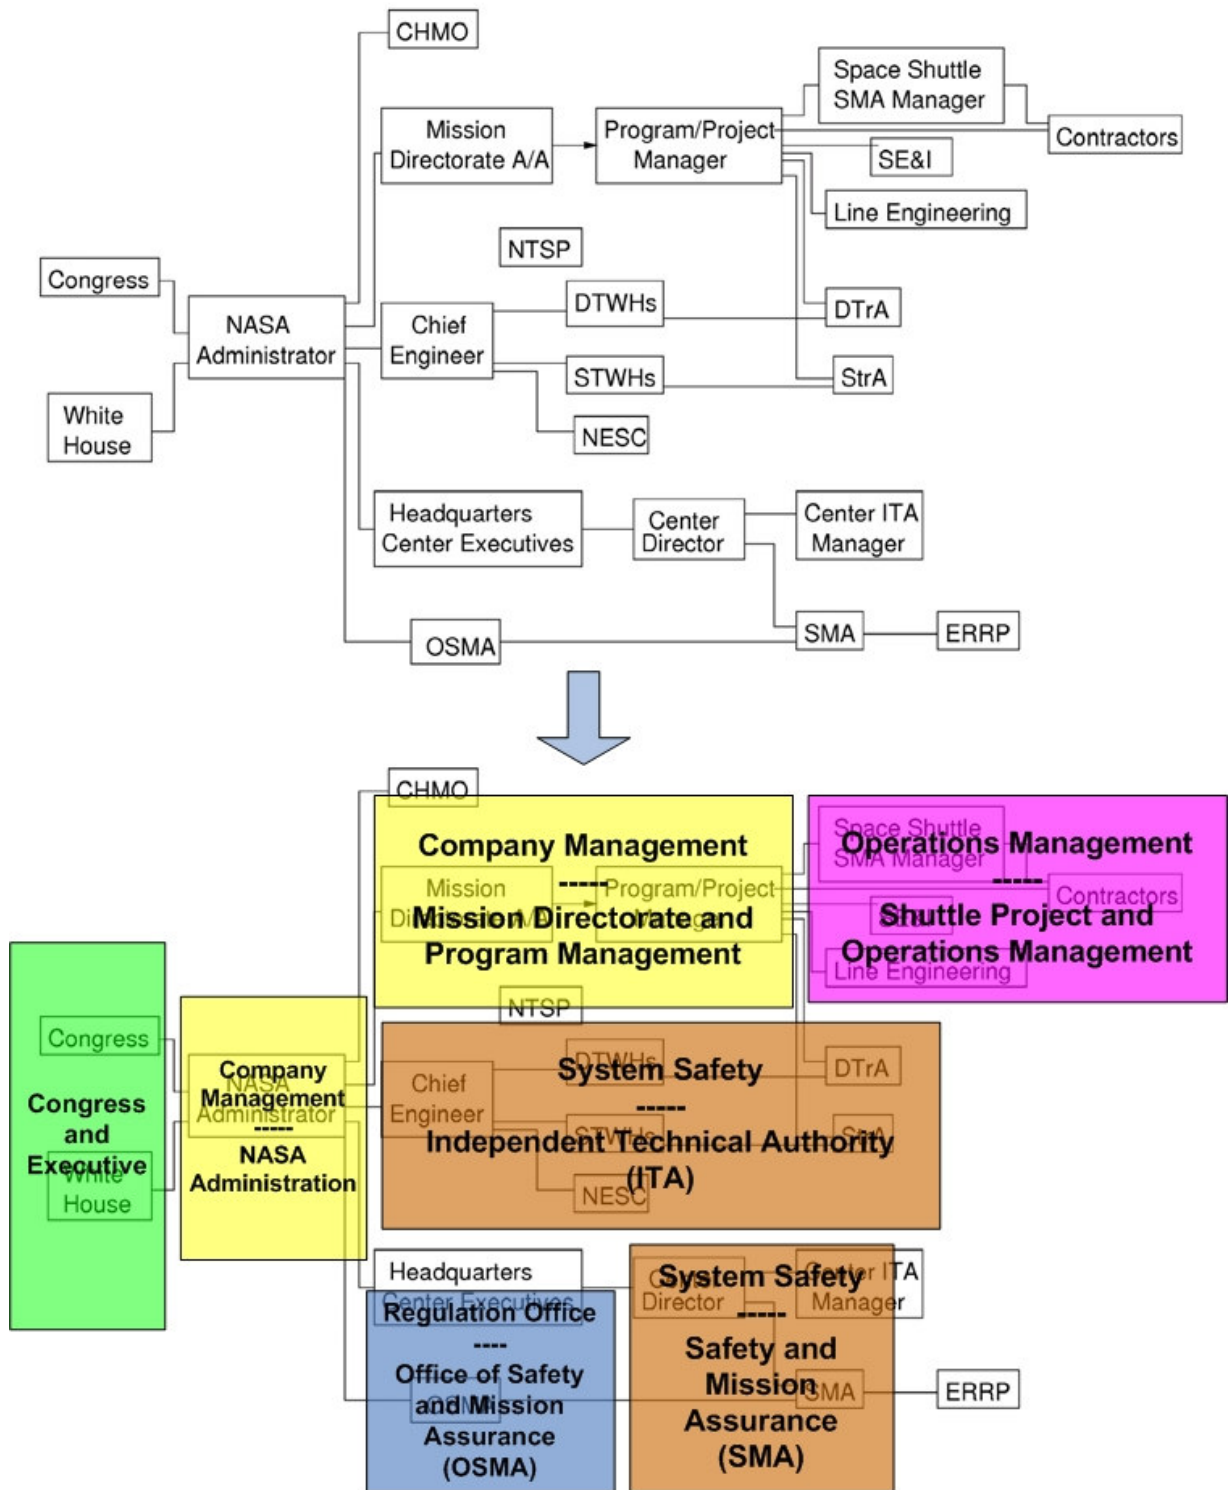


Figure 60: Step 2 - Mapping of the ITA Static Safety Control Structure to Generic Dynamic Components

Socio-Technical Safety Control Structure	Generic Component
Congress and White House	Congress and Executive
NASA Administration	Company Management
Mission Directorate and Program Management	Company Management
Office of Safety and Mission Assurance	Regulatory Agency
Safety and Mission Assurance	System Safety Activities
Chief Engineer and ITA	System Safety Activities
Shuttle Maintenance and Evolution	System Maintenance and Evolution
Shuttle Project Management	Operations Management

Table 2: Summary of mapping between ITA control structure components and generic components

4.2.3 STEP 3: REFINEMENT OF DYNAMIC SAFETY MODEL STRUCTURE

Once the mapping from the original safety control structure to the generic dynamic components is performed, the components must be rearranged to form the structure of the dynamic safety model. An example of this rearrangement is shown in Figure 61. The continuous links on the model structure diagram indicate continuous or regular influences. For example, in the case of the space shuttle system, Congress annually allocates budgets and sets performance objectives; in return, it receives regular reports from NASA administration. These influences affect the system operations on an ongoing basis. Dotted lines on the diagram indicate discrete decisions that influence system operations indirectly. Decisions made early in the system development, such as the scope and purpose of the system as well as initial decisions about system design for safety and maintainability have a deep impact on system operation even though they were one-time decisions made years ago. It is necessary to capture some of these early developmental decisions in order to understand their long-term impact on the safety of the system. Additional components that reflect some of the tasks necessary to operate the system may be added if necessary. For example, the performance of the space shuttle program relies heavily on the maintenance of its vehicle fleet, as the vehicles are re-used and must be prepared, maintained and upgraded. Consequently, a component called “Shuttle Maintenance and Evolution” was added to the newly created control structure

(see Figure 61). It is not necessary to define precisely the nature of the influences connecting every component pair at this point, however, as much information as possible should be collected as it will facilitate the rest of the process.

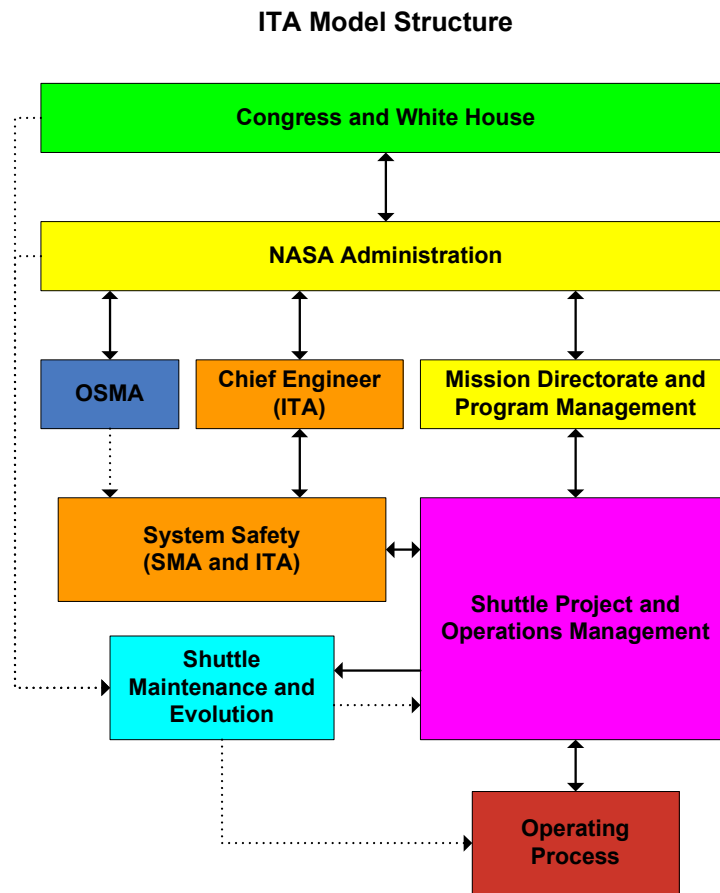


Figure 61: Step 3 - Re-structuring of safety control structure of the ITA

4.2.4 STEP 4: MAPPING OF PRESSURES, INFLUENCES AND REPORTING CHANNELS

This step involves the mapping of dynamic pressures, influences, and reporting channels throughout the safety control structure obtained in step 3. The following subsection provides a short description of a list of dynamic connectors that should be used as a guideline to map the dynamic influences. The mapping is done in a similar way to the static connector mapping introduced in chapter 3. In fact, some of the dynamic connectors are similar to static connectors. In that case, the mapping should also be similar to that introduced earlier. Not all connectors are relevant to every control structure model. Moreover, the list is not exhaustive;

it was created based on a limited set of example models and should be augmented with additional items as necessary.

4.2.4.1 Performance Pressure (Production, Throughput, Schedule)

This Performance Pressure connector transfers performance pressure information across dynamic components. Performance pressure is defined as the urgency to improve performance (production, throughput, schedule) relative to current system state. The performance pressure is felt differently by various components and has varying impact.

4.2.4.2 Quality of Imposed Safety Processes and Standards

The Quality of Imposed Safety Processes and Standards connector can be static or dynamic, depending on whether safety processes and/or standards are frozen, or change over time. The source of this connector is usually a regulatory body or office or a standard and processes office or department, either governmental or within a company. The final sink of this connector is the component where the standards or processes are used or applied.

4.2.4.3 Quality and Quantity of Safety Oversight

The Quality and Quantity of Safety Oversight connector can also be static or dynamic, depending on whether the oversight effectiveness changes over time. This connector often follows the “safety standards and processes connector”, but not always, as the component responsible for the creation of safety standards, and that responsible for the safety oversight or standard enforcement may be different.

4.2.4.4 Safety Pressure

The Safety Pressure connector transfers safety pressure information across components. Safety pressure is defined as the urgency to improve safety relative to current system state.

4.2.4.5 Resource Pressure (Safety and Operation)

Resource pressure is a dynamic connector that usually follows the resource allocation static connector defined in chapter 3. Resource pressure is a relative connector defined with respect to a baseline amount of necessary resources. Two types of resource pressures are defined. Safety resource pressure directly impacts safety-related activities, including safety headcount,

quality and experience of safety personnel, safety analyses, testing, etc. Operations resource pressure is more general and affects normal system operation activities that may indirectly affect system safety, such as system maintenance and evolution, system integration activities, hardware manufacturing and testing, etc.

4.2.4.6 Problem, Incident and Accident Reports

Incident report is a dynamic connector that follows the reporting path of problems, incidents and accidents encountered during the operation of a system. Problems are defined as events or states that do not require a pause in the system operation, but may necessitate immediate or urgent attention and corrective actions. Incidents are close-call events or states that could have resulted in a loss (accident) and may necessitate a pause in the system operation. Accidents involve a loss that necessitates some investigation and repair time, usually requiring a pause in the system operation. In many cases, the reporting channels will be similar for the three types of events. If the channels are significantly different, they should be de-coupled and treated separately. If accident visibility is very high, the accident report channel may not be needed as every component becomes immediately aware of a major loss through other means (media, etc.).

4.2.4.7 Performance Reports

The performance report connector provides information about the current performance of the system. It usually runs in a direction opposite to the performance pressure connector.

4.2.4.8 Cost Reports

The cost report connector provides information about the current operating cost of the system. It is defined in relative terms with respect to planned and agreed upon operating costs.

The following figures (Figure 62, Figure 63, and Figure 64) illustrate the mapping of generic dynamic connectors to the components used in the Independent Technical Authority (ITA) model created in step 3.

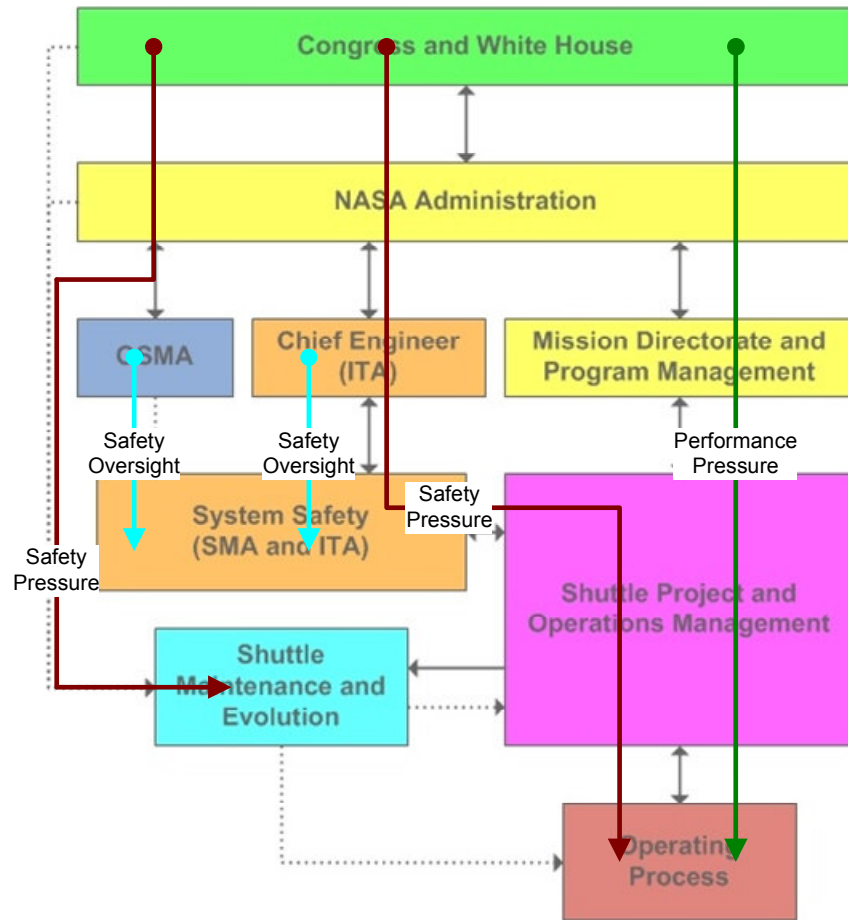


Figure 62: Mapping of Safety Oversight, Safety Pressure and Performance Pressure on the ITA dynamic control structure

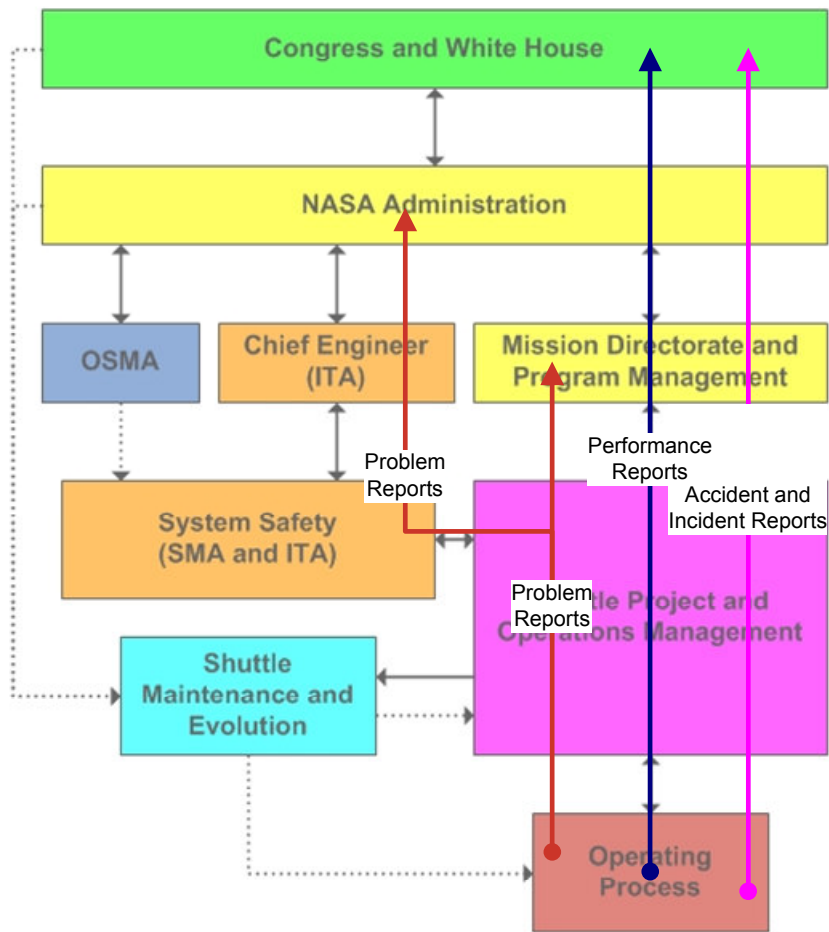


Figure 63: Mapping of Performance Reports, and Problems, Incidents and Accident Reports on the ITA dynamic control structure

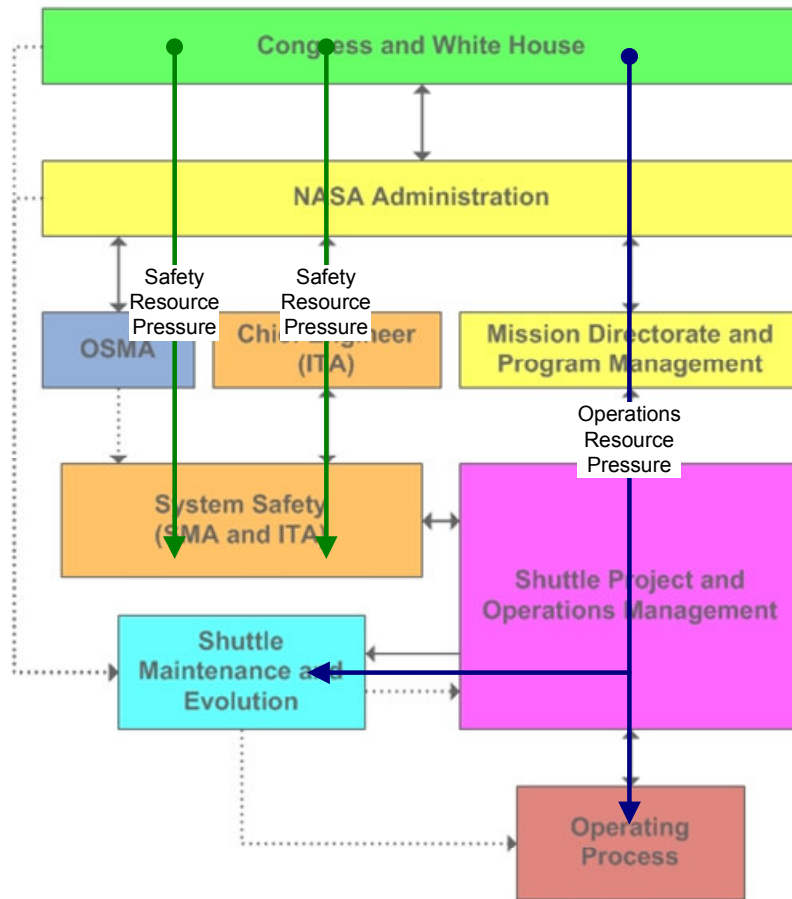


Figure 64: Mapping of Safety and Operations Resource Pressure on the ITA dynamic control structure

4.2.5 STEP 5: DATA COLLECTION AND COMPONENT CALIBRATION

In this step, the internal causal structure and decision rules of the generic components are customized, modified, and calibrated through data collection and interaction with domain experts and system stakeholders. In step 5, the focus is at the individual component level. As recommended by Morecroft [Morecroft, 1985], partial-model testing is used throughout the process both to improve the robustness of models, and to show that model components provide a good representation of the intended rationality of decision-makers within components. Multiple sub-steps are documented to guide the component customization process.

4.2.5.1 “Free-Component-Diagramming” Method

Each component is an independent dynamic model by itself. Inside each component are state variables that are computed as a function of the values of inputs to the component. Each component outputs values that are computed as a function of both the state variables and the inputs to the component. This process of converting inputs to outputs through dynamic state variables is shown in Figure 65. The outputs created are used as inputs to other components.

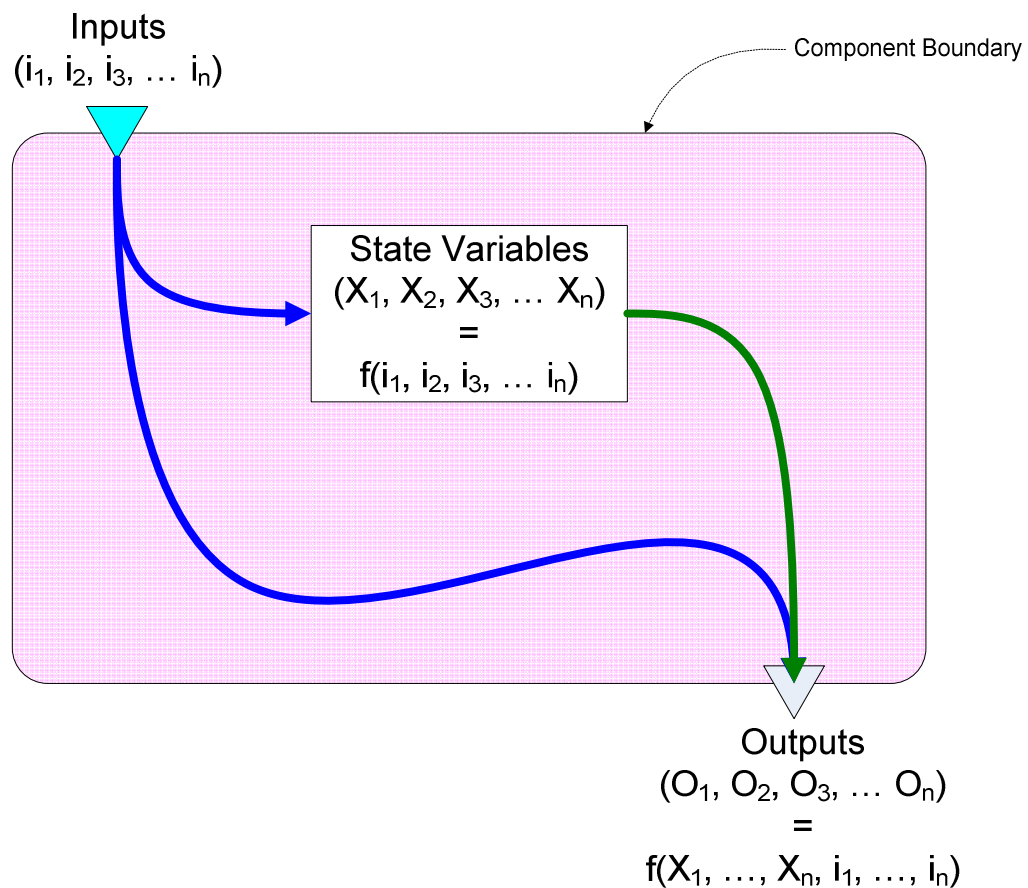


Figure 65: Component inputs, state variables, and outputs.

In order to improve the robustness of the model creation process, the generic components selected in step 2 must be verified to ensure that their internal structure is well aligned with that of the system control structure model and influence map created in steps 3 and 4. A technique was developed to perform a first assessment of the alignment of the component and integrated model structure. The technique involves the isolation of each component by “cutting” all the links created from the mappings performed in step 4. Figure 66 shows all the

dynamic connectors superimposed from step 4 (see Figure 62 to Figure 64). In Figure 66, a “virtual container” is created by “cutting around” the NASA Administration component. Each cut connector becomes a generic input or output to the newly created “virtual container”. The final result is a component “virtual container” that communicates with the outside through generic connector points (see Figure 67).

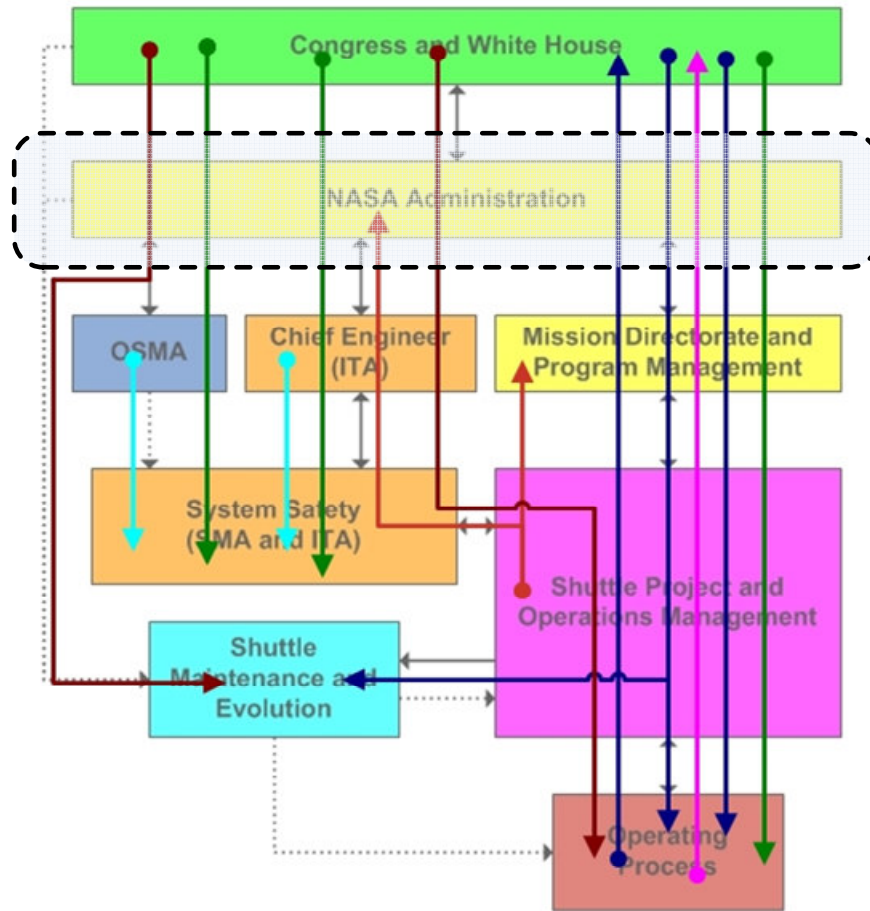


Figure 66: Isolating a component by "cutting out" dynamic connectors

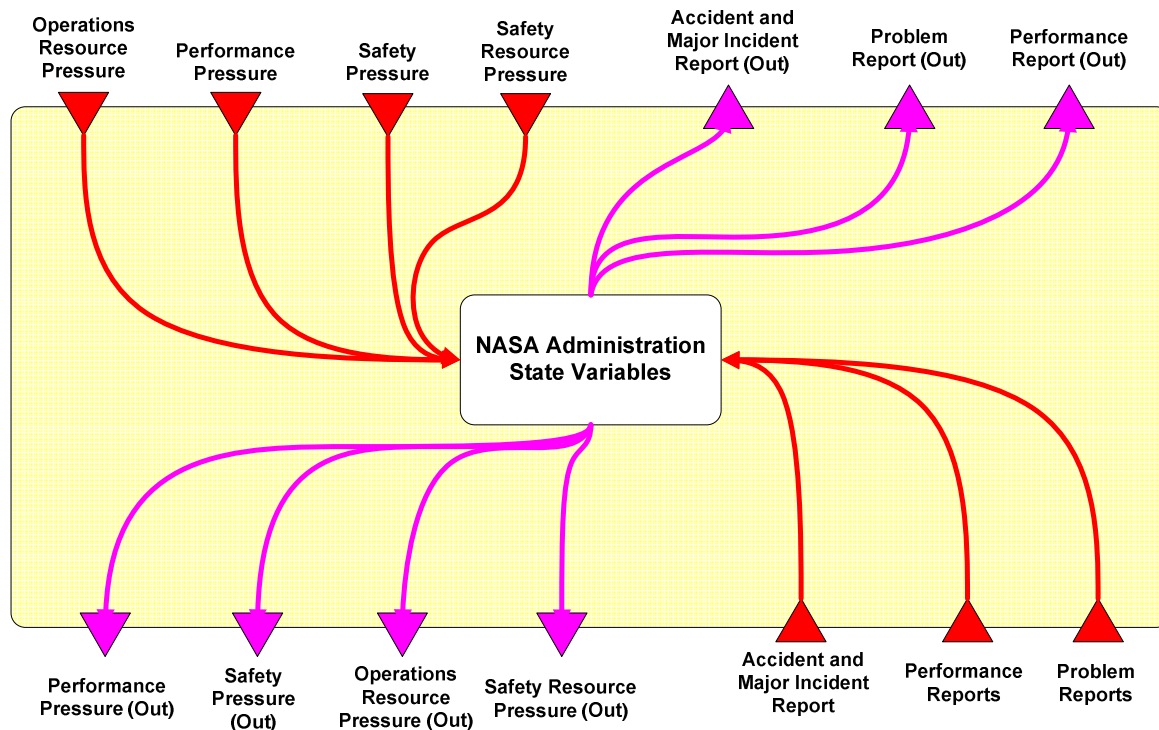


Figure 67: Input-Output virtual container created using the "Free-Component Technique"

Additional information may be required for the component to function properly. This information comes from inputs transferred from other components in the control structure model. Additional inputs can create additional dynamic connectors that must be further documented. The technique presented in this section provides a means for verifying that the internal structure of components selected from the repository of dynamic generic components to create an integrated control structure model is properly aligned with the system-level structure of interconnected components. This verification process can be done by matching the inputs and outputs of the I/O virtual container created in this step with that of the generic dynamic components selected from the repository (see Appendix E). If there is a large mismatch between the I/O structure of virtual containers and the I/O structure of generic dynamic components, then either: 1) the wrong component was selected, 2) the component should be modified to better match the container I/O structure, 3) the component boundaries should be re-drawn, 4) the influence mapping is inaccurate and should be re-assessed, or 5) a new component should be created and added to the repository. The process of creating virtual containers and comparing their I/O structure to that of generic components from the repository can be automated as soon as a control structure is created and mapped (step 4).

Automated tools can even be created to scan through the repository and suggest components that better match the I/O structure of the virtual containers created.

4.2.5.2 Data Collection and Component Customization

For each of the components used in the model, as much data as possible should be collected to align the behavior of the model components to that of real-world subsystems. The generic causal structure and variables provided in the components (see Appendix E) should be used to guide data collection. Many system-specific variables in the model must be quantitatively estimated to align the component behavior to that of the real system. The generic structure variables are usually expressed in a non-dimensionalized format to make the model “parametric” and account for a variety of different systems and applications. System data should be collected in order to convert the critical variables into a dimensionalized form. The objective is to anchor critical model variables to quantitative variables and characteristics of the real-world system. Examples of this anchoring process will be provided in the case study of Chapter 6. A list of the system-specific variables to be estimated are provided along with the generic component structure in Appendix E. Sources of data to be used in order to estimate those variables include interviews with system participants and other stakeholders. Data collection should be guided by these interviews and interactions. Chapter 6 provides a detailed example of interview-driven component-specific data collection using the Exploration Systems Mission Directorate project as an example.

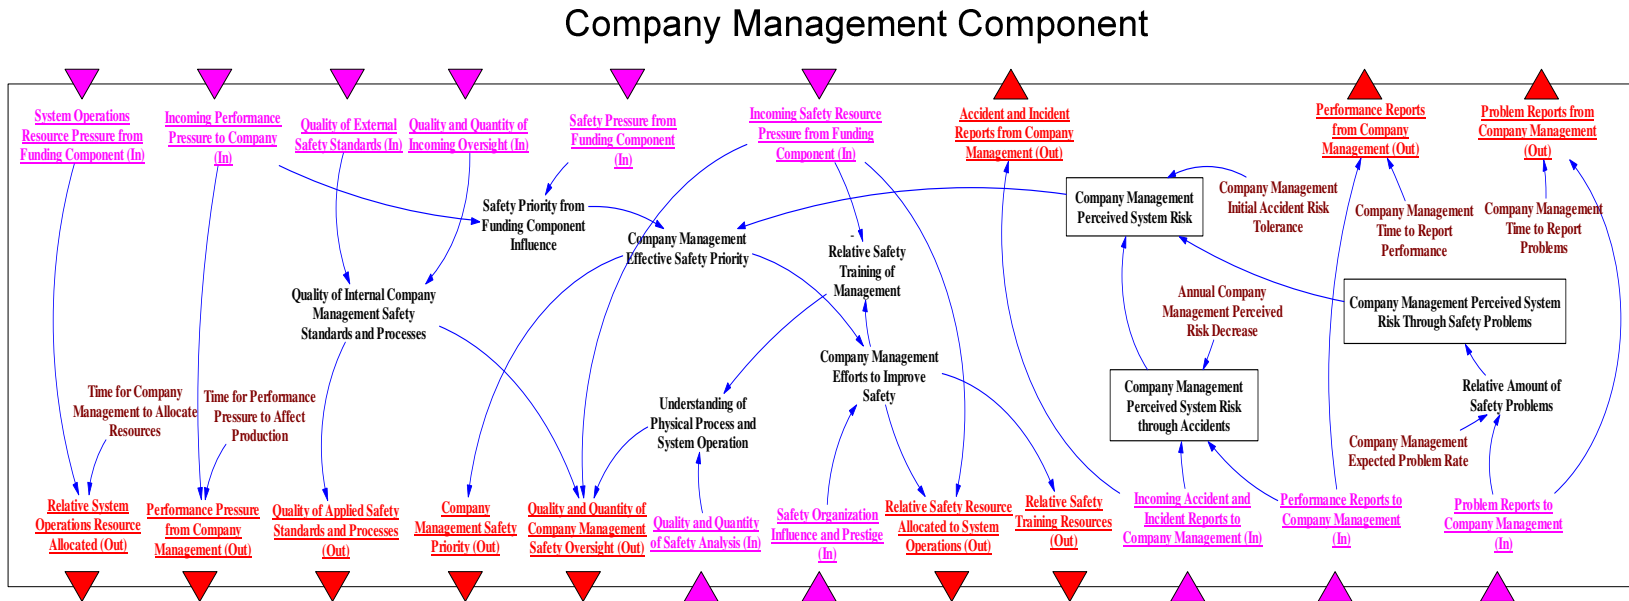
4.2.5.3 Component Causal Structure Validation Process

A causal structure validation methodology was created to assist this process. The first step of the methodology involves a discussion of the control structure with interviewees. The objective is to validate the model control structure and identify the component(s) that better fits the position and expertise of each interviewee. Once an agreement has been reached on which component to discuss, an enlarged, poster-sized copy of the initial component structure and connections is provided to facilitate discussion. Interviews follow a semi-structured format, where the initial component structure provides a map to guide the discussion, but the interviewees are given the freedom to spend more time discussing specific parts of the causal structure. Decision rules associated with critical causal links are discussed. Quantitative

variables are estimated based on expert opinions, and interviewers ask for additional sources of data as necessary.

This validation process is performed through multiple rounds of interactions with system participants. If the causal structure is found by participants to be inadequate to represent the behavior of a specific system or subsystem, three solutions are possible, ranked in order of implementation difficulty: (1) In some cases, a simple redrawing of the component boundary can be sufficient to capture the relevant causal structure. In this case, parts of a component causal structure can be borrowed from another component, or parts of the structure can be removed from a specific component. (2) A second solution involves a modification of the structure of generic components. This solution creates a new modified component that better matches the causal structure and associated decision rules of a component. The modifications can be as simple as adding an input to a component to communicate with another. The modifications can also be more involved, including an internal re-wiring of individual components or the addition (or removal) of micro-structure within a component. (3) A third solution is the creation of an entirely new component that can later be added to the repository of generic components as needed. Micro-structure from other components can be borrowed as necessary from other generic components. An example generic component is shown in Figure 68. Note the similarity between the I/O structure of the NASA Administration component of Figure 67 and the I/O structure of the Company Management component shown in Figure 68. This similarity facilitates the customization of generic components for their use in system-specific dynamic control structures. A detailed example for a complete model is provided in Chapter 6.

Figure 68: Example generic component (Company Management) and associated I/O structure



4.2.5.4 Quantification of Component Relationships and Decision-Rules

One of the main difficulties in creating simulation models based on STAMP safety control structures is the elicitation and quantification of decision-rules used in the component models. This difficulty is common to every behavioral mathematical model. Sterman provides guidelines for the elicitation of decision rules in chapter 14 of his book [Sterman, 2000]. These guidelines include the use of reference modes elicited in the form of visual table functions easily understandable by system decision-makers, and the decoupling of decision rules until a single relationship is captured by the decision rule. Ford and Sterman [Ford, 1998] describe an expert knowledge elicitation technique that combines the knowledge of several experts into a single reference mode for the relationship. This technique is divided in three phases: 1) the positioning phase, where context is provided for every single relationship and examples from similar settings are provided, 2) the description phase where a story is created to vividly illustrate the relationship, using verbal, textual, visual and graphic descriptions of the story and associated relationship, and 3) the discussion phase, where individual descriptions from each expert are examined, compared and combined into a single relationship to be used in the model.

A similar process is used in the model-building methodology introduced here. Relationships between various model factors are examined individually and discussed with domain experts. Table functions and reference modes are used to facilitate the coding and further modification of relationships. Standard formulations are used to simplify data entry and relationship definition. For example, Figure 69 shows a sample component that defines the size of the system safety workforce in a company. The component receives four inputs from other components of the system: 1) Relative Employee Productivity, 2) Relative Budget Available, 3) Safety Priority, and 4) Average Workforce Age. These four input variables affect two intermediate variables (shown in hexagon in Figure 69): 1) Target Number of Employees, and 2) Attrition Fraction. These intermediate variables are defined relative to a “normal” value (shown in circles in Figure 69). The normal variables are typically exogenous variables that need to be quantitatively estimated in order to customize the component to a specific system. For example, while interviewing a human resource manager at NASA Headquarters, it was easy to obtain preliminary estimates for the normal hiring time (an average of two months for a civil servant employee, from the time a position opens to the time the new employee starts

working) and the historical average of the attrition fraction (five percent of the workforce per year). The typical number of employees is also easy to estimate from human resource data. The exact number depends on which types of employees are considered, e.g. only system safety analysts at the program level. The variables starting with “Effect of” get multiplied by the normal/typical/baseline values to obtain intermediate values. For example:

$$\text{Target Number of Employees} = (\text{Typical Number of Employees}) * (\text{Effect of Employee Productivity on Target Number of Employees}) * (\text{Effect of Budget on Target number of Employees}) * (\text{Effect of Safety Priority on Target Number of Employees})$$

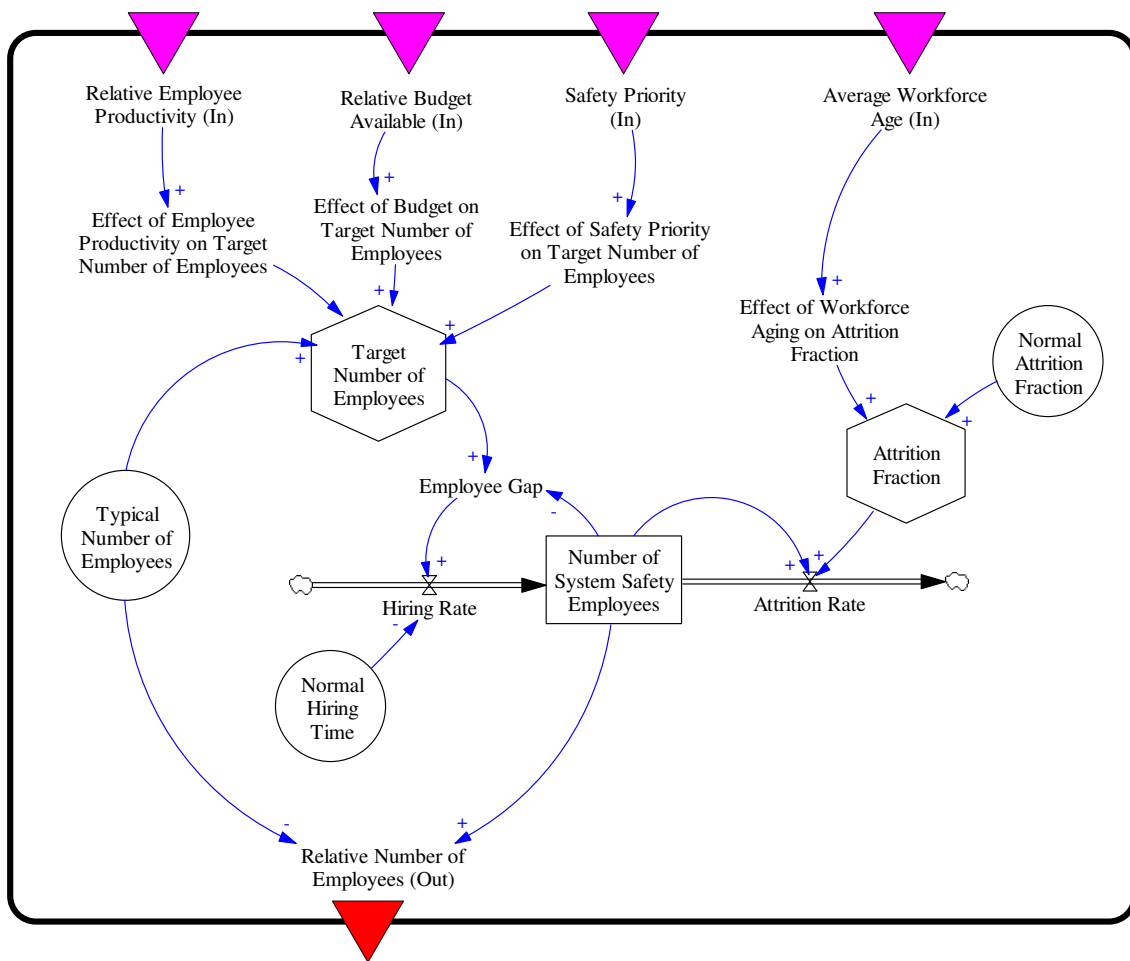


Figure 69: Example Human Resource Component Structure

The “Effect of” variables use table functions and/or equations to transform inputs into an impact on an intermediate variable. These are the type of relationships that must be discussed

with domain experts. Using lookup tables (see Figure 70 for an example) to illustrate and discuss relationships between variables facilitates the discussion and relationship elicitation process. In these tables, the input is on the x-axis, and the value of the “Effect of” variable is on the y-axis. In addition, a component-based approach allows the isolation of concepts and domain areas that can be individually discussed by appropriate experts. Anchor points are used to facilitate the definition of the relationships by anchoring the values to hard real world constraints. For example, if there is no budget available, the Target Number of Employees should be zero (see the relationship on the left of Figure 70). In addition, it is useful to define tables with respect to a non-dimensional reference value. If the input is at the reference value, that is, the relative value is 1, then the table or relationship should output the “equilibrium” value 1 for the “Effect of” variable (see the next section for a discussion on this topic).




Figure 70: Sample relationships used to define intermediate variables

When multiple values are combined into a single intermediate variable, some widely used and accepted decision analysis techniques such as Multi-Attribute Utility Theory (MAUT) [Keeny, 1993] and the Simple Multi-Attribute Rating Technique (SMART) [Edwards, 1977] can help analysts to combine and quantify the preferences and decision rules of several experts into a single function that defines the value of the intermediate variable.

In cases where domain experts are unable or unwilling to describe a relationship, simple relationships such as linear relationships (see right side of Figure 70) and averages are used until more accurate relationships can be elicited. Sensitivity analyses to model parameters,

relationships and decision rules are used throughout to ensure policy robustness, that is, to ensure that the outcome of recommendations made based on model behavior is not sensitive to uncertainty in model parameters and relationships. Policy robustness is discussed in more detail in the context of model analysis (see chapter 5).

4.2.5.5 Equilibrium Boundary Conditions

In order to alleviate the complexity associated with the use of large simulation models with hundreds of variables and relationships, a technique was developed to standardize and facilitate model analysis and component-based model creation. The technique involves the use of “equilibrium component boundaries”. The main operating principle is that individual components should be able to exhibit equilibrium behavior if they are undisturbed from the outside. For example, when left completely undisturbed, the state variables and outputs of the Company Management component shown in Figure 68 are in equilibrium. This equilibrium behavior of the Company Management is demonstrated in Figure 71. The equilibrium condition is critical for building large models because it forces the creation of standard interfaces which allows components to be connected through generic connectors.

In order to obtain well-defined interfaces that are able to exhibit equilibrium behavior, it is often necessary to non-dimensionalize the parameters transferred in and out of components. This is done by defining baseline system-level dimensionalized values that act as a key to re-dimensionalize a parameter when needed. For example, if the baseline expected production for a manufacturing system is 10,000 widgets per month, the desired production from management and the production report to management will be transferred across components in a non-dimensional form. If company managers suddenly desire a production of 12,000 widgets per month, the Relative Desired Production connector will transfer a non-dimensional value of 1.2 to the production unit component. If the current value of production is 9,000 widgets per month because of supply delays, for example, the production unit component will output a value of 0.9 through the Relative Production Report connector. Non-dimensional values can be easily re-dimensionalized and anchored to the real system by multiplying them with baseline values. In the equilibrium behavior mode, non-dimensional values of 1 are passed across most generic connectors. In the previous example, it would mean that company managers desire a production of 10,000 widgets per month, which is exactly what is produced

and reported by the Production Unit component, thus no production adjustment is necessary and the behavior is in equilibrium.

In addition, equilibrium behavior allows partial models to be executed piecewise, facilitating testing and validation. Finally, it allows some model component influence to be turned on and off, which facilitates analysis by limiting the number of simultaneous model influences to be considered. For example, if the structure or relationships in one component has high associated uncertainty, the component can be put in “artificial equilibrium”, thus removing its influence on the larger model. This greatly facilitates model and policy robustness analysis.




Figure 71: The Company Management component exhibiting equilibrium behavior

4.2.6 STEP 6: COMPONENT TESTING AND CONFIDENCE-BUILDING ACTIVITIES

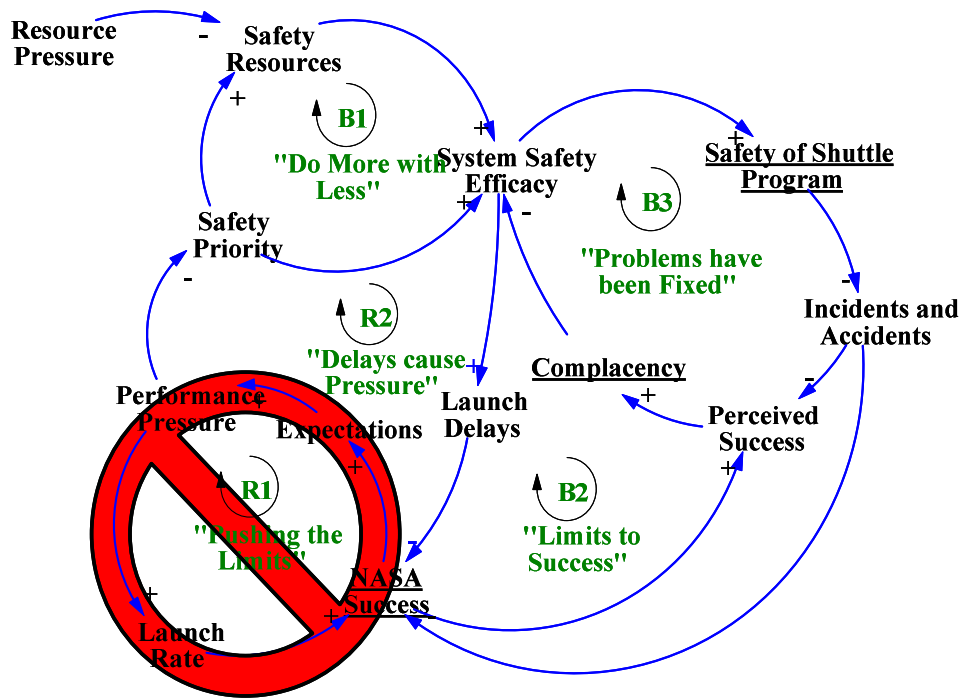
Step 6 involves the quasi-validation of component behavior through testing and confidence-building activities. Building components with a well-defined I/O structure and interfaces facilitates this validation process. This section discusses component testing and validation techniques tailored to the component-based model creation methodology.

4.2.6.1 Stress Testing

The stress testing condition states that given inputs to a component that are within the bounds defined and documented in the component, then component outputs created should also fall within the bounds defined for the component outputs. This can be easily verified (but not proved) using brute force testing. Random inputs can be generated and injected using modern system dynamics modeling packages including Monte-Carlo simulation toolkits. Given a combination of randomly generated inputs within the bounds of input ranges, the outputs should also fall within acceptable documented ranges.

4.2.6.2 Equilibrium Condition Test

The Equilibrium Condition test is closely related to the component equilibrium behavior discussed earlier. In an operational system, the equilibrium behavior is achieved by “cutting” the “Pushing the Limits” feedback loop (see Figure 72). Unless this reinforcing loop is inactive, equilibrium will not be achieved because of increased expectations. As mentioned previously, the baseline values of variables for each component should be chosen so that connectors transfer the non-dimensional value of 1 when the “Pushing the Limit” loop is inactive. Components should be able to exhibit this equilibrium behavior given that they are initialized with equilibrium values. In practice, the “Pushing the Limits” loop is usually active because, in most systems, increasing performance is necessary for an organization to remain competitive and economically viable over the long term. However, this theoretical “loop-cancellation” exercise is useful to ensure that each component has the ability to exhibit equilibrium behavior in cases where the necessary conditions are fulfilled. This condition is necessary (but not sufficient) to allow the inter-operability and inter-connectivity of various components of the model. The equilibrium condition can greatly facilitate model integration and analysis.



Inactive Pushing the Limit Loop

Figure 72: Deactivating the "Pushing the Limits" Loop

4.2.6.3 No-Accident Condition Test

This partial-model rational behavior test is performed to determine if the components behave according to expected rational partial-model behavior, given a theoretical ultra-safe system where accidents and major incidents never occur, regardless of the safety efforts and resources deployed. Generally, if the “Pushing the Limit” Loop is inactive, the components should be in an equilibrium position. If the “Pushing the Limit” loop is active, more and more resources and efforts will be put toward production versus safety. Components should exhibit a “Do More with Less” behavior: Safety efforts and efficacy are reduced to a minimum while resources and efforts are allocated toward increasing the throughput of the system. Figure 73 shows the result of the no-accident condition with the active “Pushing the Limits” reinforcing loop. Because accidents do occur in the real world and we are mostly interested in safety-critical systems with large accident consequences, this “no-accident” condition will not be achieved in a real system. However, if the theoretical “no-accident” situation did occur, decision-makers and their decision-rules should adapt their behavior to be consistent with the context.

Company Management Component Critical Variables




Figure 73: The Company Management component tested under the no-accident condition

4.2.6.4 Accident and Major Incident Response Condition

The objective of this test is to ensure that the response of individual components of the socio-technical system to accidents and major accidents is consistent with the rational behavior of decision-makers given the context. If the individual open-loop response of each component to a system accident or major incident is consistent, when the model is assembled, the closed-loop behavior of the model should also be consistent. In order to verify the rational behavior under the accident response condition, an accident is artificially “injected” in the component at time t=60 months. As can be seen in Figure 74, perceived risk and safety priority go down until the accident occurs, then go up for a short period after an accident, and then starts to go down again when the system goes back to normal operation. A similar behavior mode was obtained independently by Salge and Milling [Salge, 2006]. The behavior of every component should match the mental model of component decision-makers when an accident or serious incident “shocks” the system. The magnitude and duration of the shock is system-dependent and should be discussed with domain experts.

Company Management Component Critical Variables

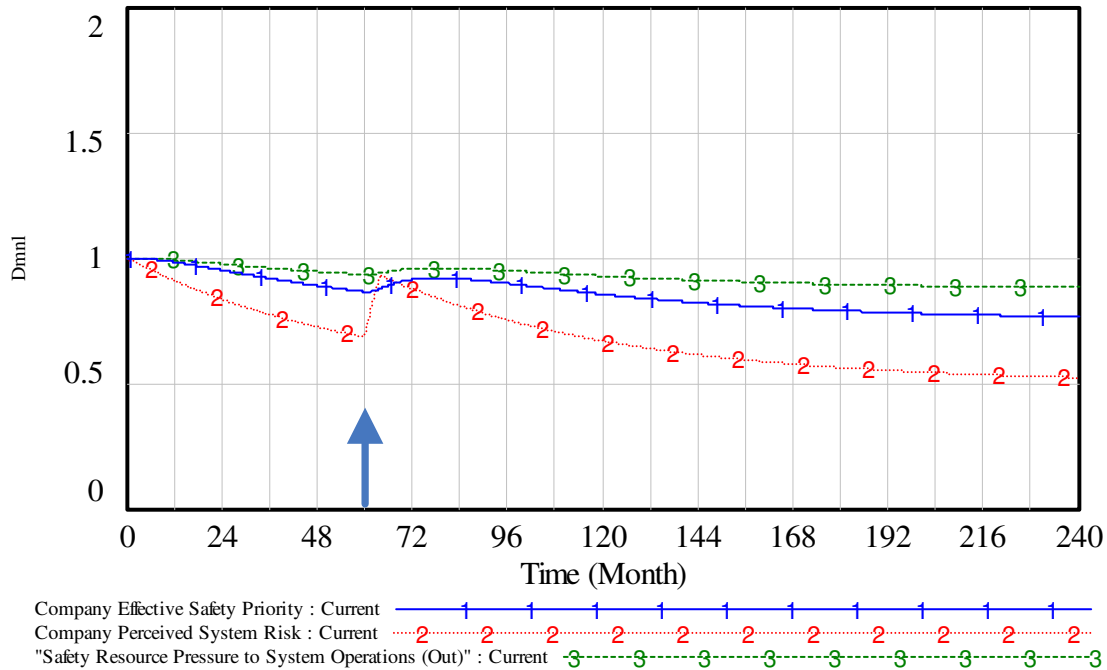


Figure 74: The Company Management component tested under the accident response condition. An accident occurs at t=60 months.

4.2.6.5 Component Intent Rationality Test

The component intent rationality test states that system experts should be able to predict the output of a model component, given that a single input is manipulated at a time. Domain experts should be asked to describe the expected behavior of component outputs when a single variable is moved from one extreme value (low bound) to the other (high bound). If there is a mismatch between the prediction of experts and the outputs of the component, then the discrepancies should be resolved.

4.2.7 STEP 7: COMPONENT-BASED MODEL ASSEMBLY

If the previous steps were performed correctly, the assembly of individual components into complete models should be greatly facilitated. The latest re-structured control structure obtained in step 3 is used as a guide to visualize the hierarchical context of each component.

The generic connectors mapping performed in step 4 provides a map of the necessary connections. Finally, the standard interfaces and connection points allow the seamless interconnection of components. The use of non-dimensional connectors is of great use for this purpose. The upper and lower bounds of a component output should match those of the component input to which it is connected. Figure 75 shows an example of the connections created between a Congress and Executive component and a Company Management Component. These connections match those between the Congress and White House component and the NASA Administration component shown in previous figures. It is not necessary to connect every interface of every component because the equilibrium condition across component boundaries allows partial models to be executed and tested. In some cases, the interfaces may be created for completeness and not necessary connected in the model. However, component interfaces that are left unconnected should be examined carefully and rationale should be documented to explain their non-use.

4.2.8 STEP 8: MODEL TESTING AND CONFIDENCE-BUILDING ACTIVITIES

The last step in the methodology is closely related to the model analysis techniques described in the next chapter (chapter 5). Step 8 is where the cyclical nature of the modeling effort becomes important. Problems, surprise behavior and insights mostly occur at the system-level testing and analysis phase. When those occur during analysis, they may trigger a reassessment of the model structure and formulations. However, before system-level simulations and analysis begins, it is important to perform a few more system-level tests on the model to build confidence in the system behavior. The following subsections provide a short description of some tests that can be performed to correct errors introduced by component interactions and connections, while simultaneously building confidence in the system-level model behavior. Some of the tests are similar to component-level tests, but are performed with the integrated system-level model.

4.2.8.1 System-Level Stress Testing

While component level stress testing is performed using component inputs, system-level stress testing is performed by randomly selecting exogenous variables within their documented bounds, and by recording the behavior of key stocks in each component. The objective is to identify and correct potential instabilities, out-of-bounds, numerical integration and other types of errors. Random generation of exogenous variable can also be done at the component level, but some errors may only appear at the system level where all components are connected into an integrated model.

4.2.8.2 System-Level Equilibrium Test

If all components are able to exhibit equilibrium behavior when taken in isolation and when the “Pushing the Limits” loop is inactive, then the integrated model should also be able to exhibit system-level equilibrium behavior when undisturbed. Perfect equilibrium may not be achievable because some variables have a stochastic component and because of initial conditions and transient behavior. For example, some components such as the problem resolution structure embedded in the Shuttle Project Management (Operations Management)

Sample Key Variables from the NASA Administration Component

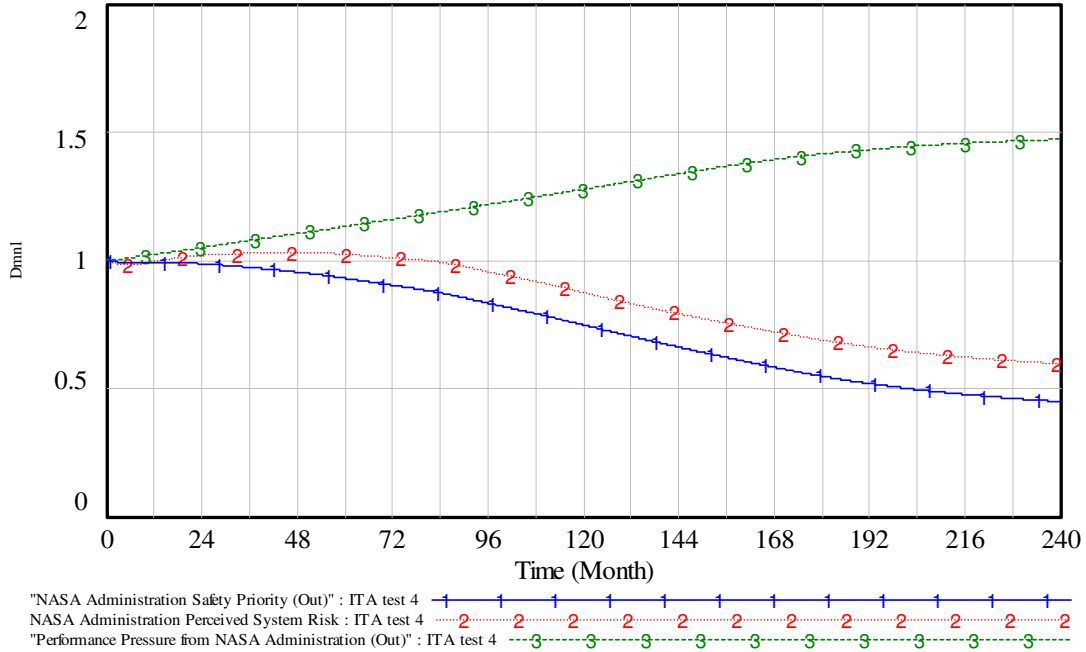


Figure 77 : Sample no-accident behavior for integrated model

4.2.8.4 System-Level Accident and Major Incident Response Condition Test

The last two test conditions are located at the boundary between system-level testing and preliminary model analysis, as they can be used to generate insight for future analysis and scenario generation. The objective of the system-level accident and major incident response condition test is to ensure that the response of the integrated model to accidents and major incidents is consistent with the rational system-level behavior given the response of every individual component. Figure 78 shows an example of accident-response behavior for the integrated system when an accident is “injected” at time t=60 months. The system-level response may seem unintuitive at first and it may only be understood in the context of individual component interactions. Nevertheless, if the system-level response is completely counter-intuitive or unexpected, it should be investigated thoroughly by backtracking through individual variables to ensure that analysts and experts understand and accept the mechanisms behind the generation of individual behavior modes before further model analysis is performed.




Figure 78: Sample accident response behavior for the integrated model

4.2.8.5 System-Level Intent Rationality Test

The system-level intent rationality test is both a test and a first step toward model analysis. The test involves the variation of one exogenous variable at the time, from one extreme value to the other within the defined bounds of the given variable. The resulting behavior for critical component state variables should be analyzed. Step functions are a great tool for this test because the sudden change in variable value can help the identification of transient model problems such as overshoot, overcompensation, oscillations, initial values, and numerical integration errors. This test should be performed in the presence of domain experts and the ensuing discussion, explanation and surprises should be documented as a great source of insight and problem identification.

4.3 MOVING FORWARD: MODEL ANALYSIS

This chapter presented an iterative methodology that should be used in combination to the generic components collected in the repository (see Appendix E) to facilitate the creation of

dynamic STAMP-based risk management models. Once the last step of the component-based model creation and validation methodology has been performed and enough confidence has been accumulated about the structure and behavior of the model, analysis and policy design and testing can begin. The next chapter (chapter 5) describes some analysis techniques that can be used in designing controls for unsafe behavior. Chapter 6 provides a complete case study covering every step of the model creation methodology (chapter 4) and analysis (chapter 5) using the NASA Exploration Systems project.

CHAPTER 5: ANALYSIS OF DYNAMIC SAFETY AND RISK MANAGEMENT MODELS

In previous chapters, a component methodology was introduced to facilitate the creation of custom STAMP-based dynamic risk management models. The methodology is iterative and once enough confidence has been accumulated in the structure and behavior of the model, analysis can be performed to fulfill four main objectives: 1) Improve the quality of the mental model used to make safety-related decisions, 2) Analyze the risks identified by system analysts and stakeholders, 3) Improve the robustness of systems against time-dependent risk increase, and 4) Improve risk monitoring to detect and correct potential migration toward higher risk levels. This chapter is divided in four sections associated with each objective. Several techniques are presented to achieve each objective. The objectives are not mutually exclusive and some of the techniques presented to address a specific objective can also be used to address others.

5.1 IMPROVING SAFETY/RISK DECISION MAKING

One of the objectives of the modeling process and resulting model is to improve the mental model of decision-makers. The assumption is that improving mental models will consequently improve the quality of safety-related decision-making. Improving mental models of decision-makers is a common topic in many fields including engineering, business and policy-related decision-making. Moreover, one of the long-term objectives of system dynamics research (and modeling in general) is the improvement of the mental model of decision makers in order to improve the performance of organizations and systems [Forrester, 1969; Forrester, 1985; Morecroft, 1985; Morecroft, 1988; Sterman, 1989; Sterman, 1989; Wolstenholme, 1990; Sterman, 2000]. Three related techniques are discussed in this section: 1) The creation and use of custom risk management tools and simulators based on the dynamic models created; 2) Improved visualization of model structure and behavior; and 3) Interactive scenario-based learning and decision support.

5.1.1 MANAGEMENT SIMULATORS

Management flight simulators based on system dynamics models have been used extensively to help decision-makers in improving their mental model of specific business phenomena. Examples of extensively used flight simulators include [Sternan, 2000]: 1) The People Express simulator, where users attempt to guide a new airline toward financial success, 2) The B&B (Boom and Bust) Enterprises flight simulator, where users manage a new consumer product from launch to maturity, 3) The F&B (Food and Brands) Enterprises flight simulator where users manage a consumer brand and product through its entire lifecycle within a competitive environment, and 4) The Beer Distribution simulator, a simulation-based version of the “Beer Game”, where users manage the production and distribution of a product.

Using a similar approach, a prototype risk management simulator was created [Friedenthal, 2006] based on a model of the NASA space shuttle system built by the author. The pictures in Figure 79 and Figure 80 show two interfaces of the risk management simulator. As with other simulators based on system dynamics models, the risk management simulator allows the simulation to be divided in equally spaced time intervals. At each interval, users have to make management-type decisions such as production (launch) objectives, and the amount of resources allocated to areas such as system safety, outside contracting, problem investigation, and system maintenance. After each decision, the simulator runs the model for a time interval, and displays the results of the model through indicators such as the Average NASA Safety Experience, Safety Training, Reported Problems, Fraction of Problems Investigated, Investigation Quality, Investigation Workload, Launches Delayed because of Safety Problems, etc. These indicators can be displayed in numerical format (Figure 79), or in a 4x3 risk matrix used for traditional risk management at NASA (Figure 80).

More research will be required to evaluate principles and guidelines for the design of effective custom risk management simulators, interfaces and features. However, preliminary informal experiences with the manned space program management simulator shows that it can make large dynamic models accessible to managers, help them better understand the dynamics of

the systems they are controlling (managing), and, more importantly, allow them to analyze the impact and potential unintended effects of safety-related decisions.

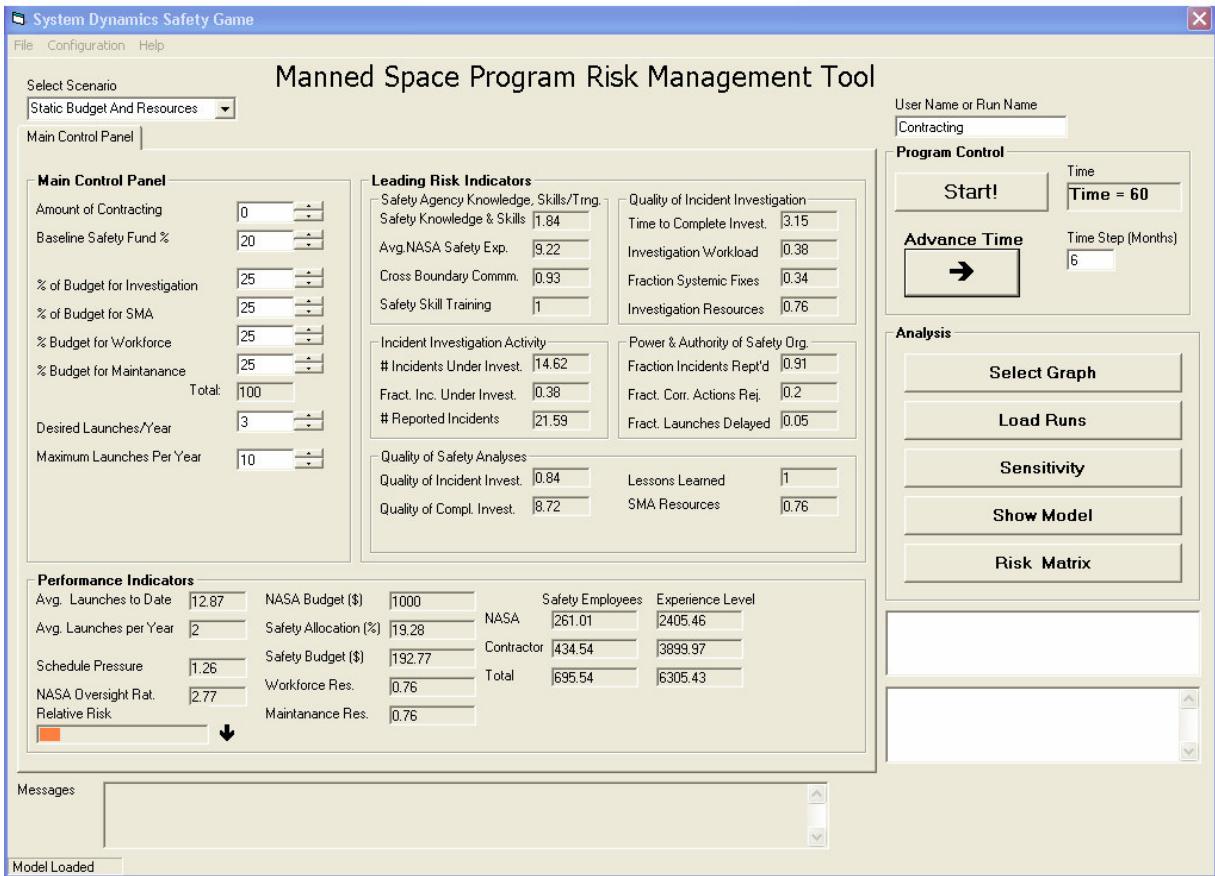


Figure 79: Main interface of the Manned Space Program Risk Management Simulator [Friedenthal, 2006]

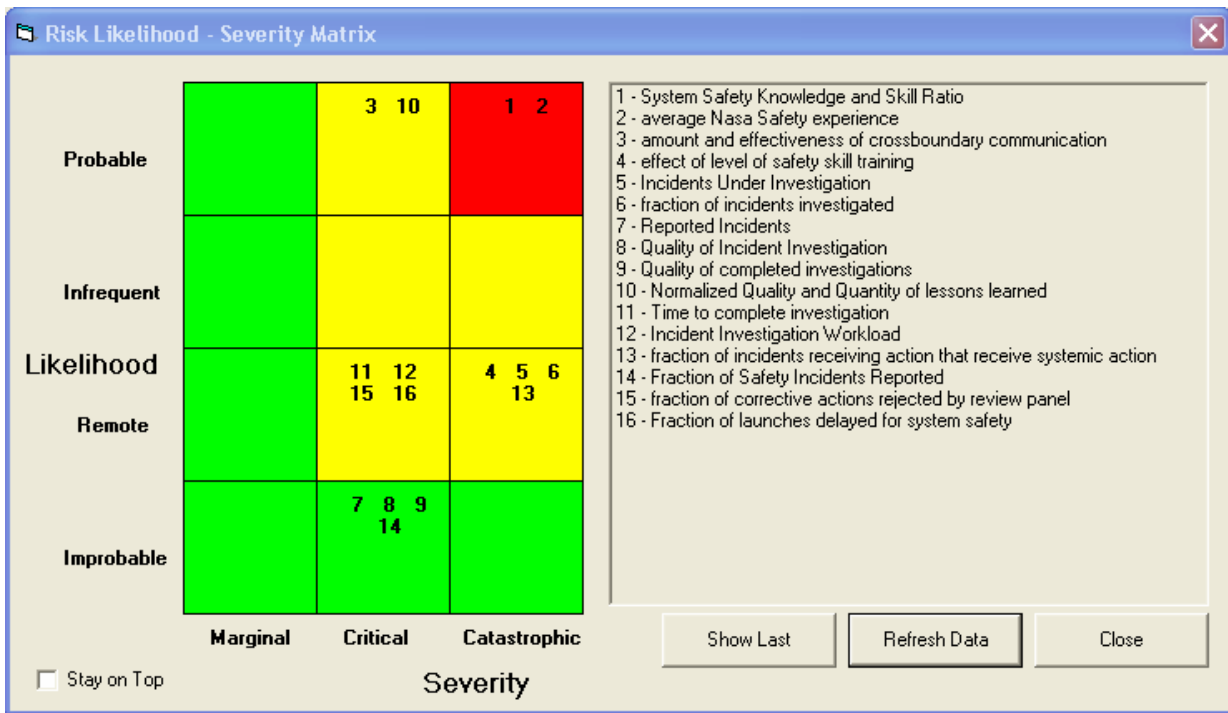


Figure 80: Interface showing a 4x3 risk matrix for tracking the risk scores of the leading indicators [Friedenthal, 2006]

5.1.2 VISUALIZATION OF MODEL STRUCTURE AND BEHAVIOR

A potential obstacle to the acceptance and use of STAMP-based dynamic risk management models (and system dynamics models in general) is the complexity associated with the dozens or hundreds of variables and feedback loops embedded in models. The hierarchical decomposition of the models created based on the STAMP-based technique presented in chapter 4 facilitates the breaking down of large models into smaller models that are more intellectually manageable. For example, it is possible to “zoom” into each component and see the hierarchical context within which it functions and its purpose within the larger system model. In the ITA model, for example, zooming into the ITA components shows the context and structure within which the ITA subcomponent is embedded (Figure 81). Model creation tools implemented in the future must be able to take advantage of the natural hierarchical decomposition provided by the combination of STAMP safety control structures with generic dynamic components and structures.

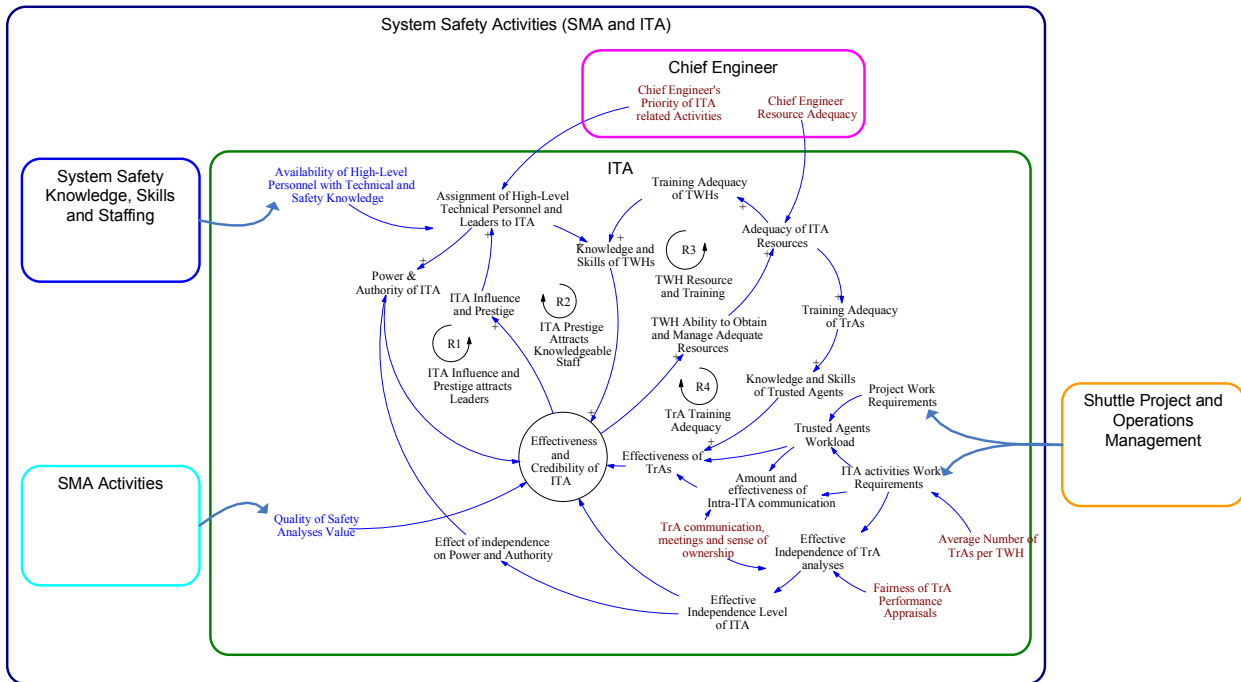


Figure 81: Context of the ITA subcomponent within the larger ITA model

In addition to the visualization and understanding of models through a more effective use of hierarchical levels and decompositions, visualization tools can be created to better understand the dynamics associated with individual components and/or individual feedback loops. As an example, the prototype risk management simulation tool documented in [Friedenthal, 2006] provides a user interface that allows the visualization of individual feedback loops, their effects and the presentation of related documentation. Individual loops, behavior and documentation can be easily added and removed from the display, which facilitates model understanding. Figure 82 and Figure 83 show screen captures of the risk management tool interface that demonstrate the ability to superimpose individual loops and understand the model structure in “layers”.

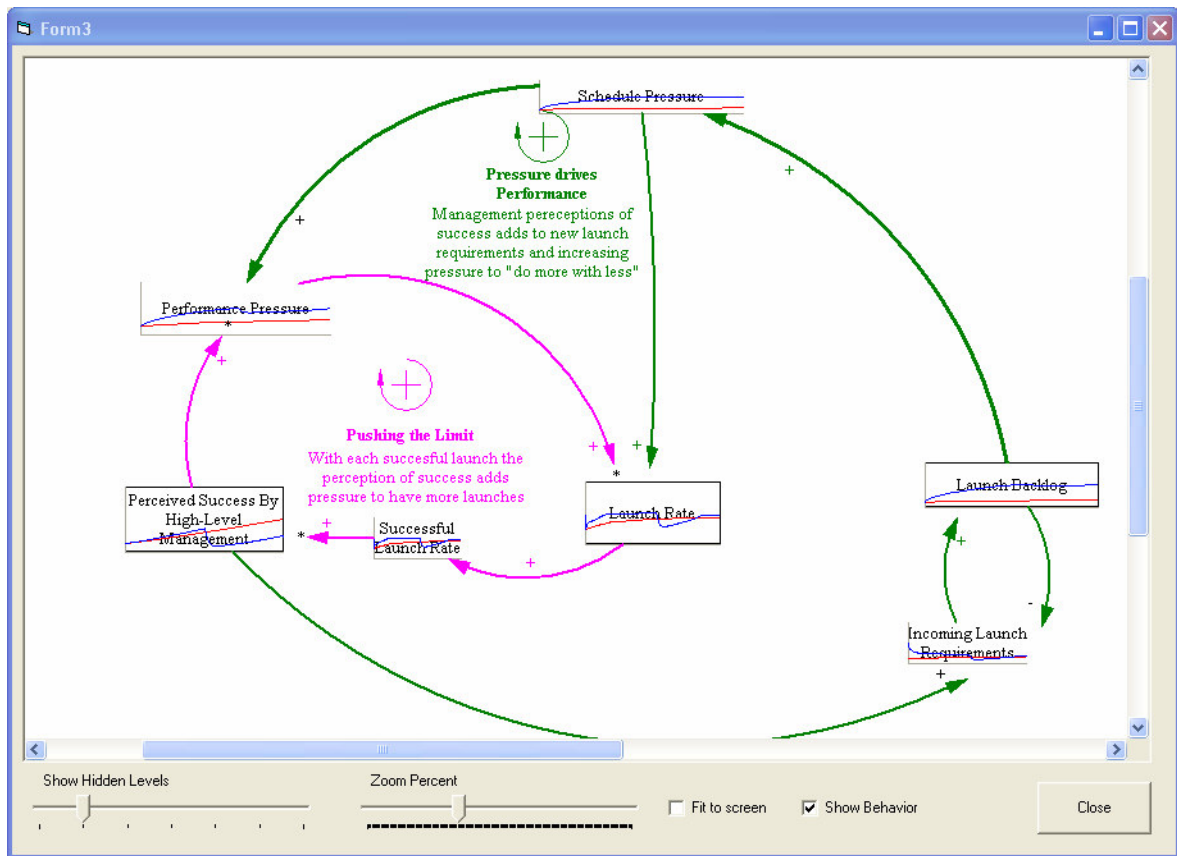


Figure 82: Screen capture of three feedback loops of the ITA model [Friedenthal, 2006]

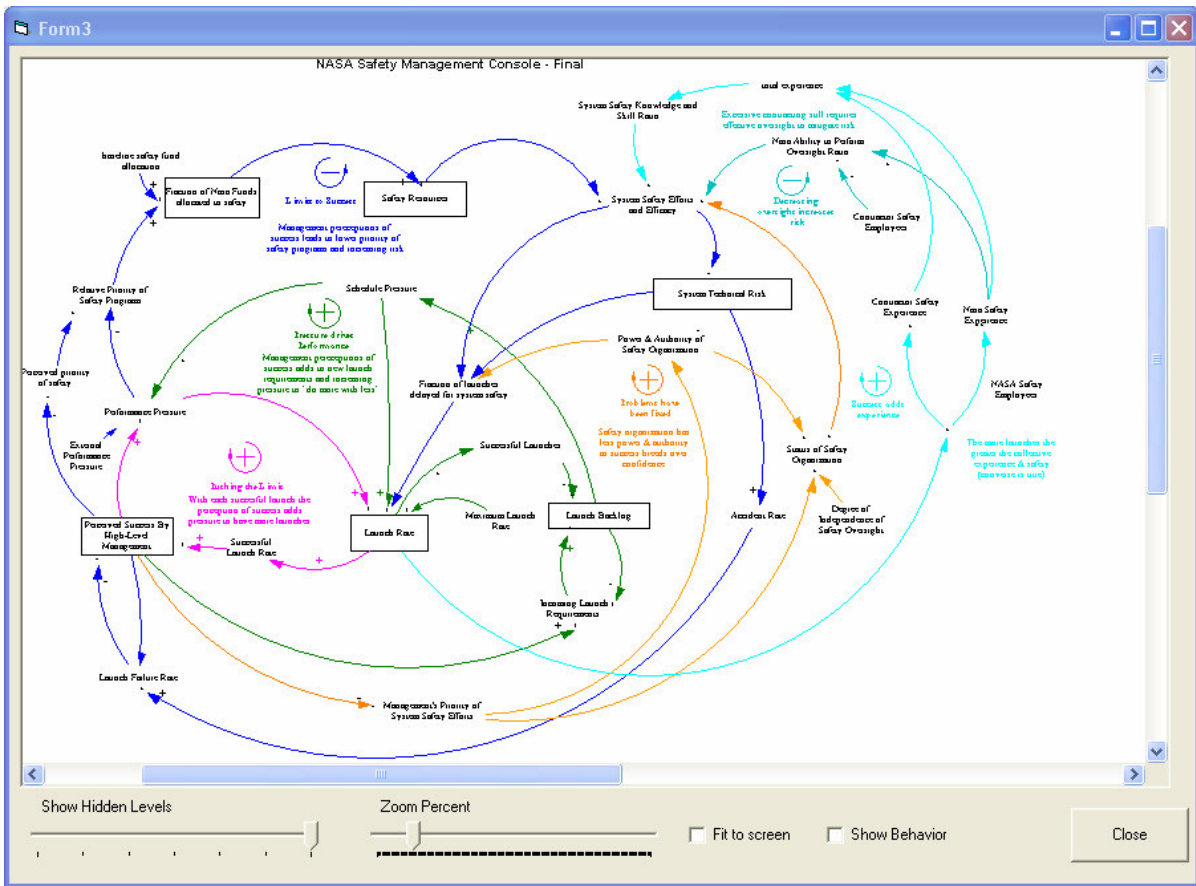


Figure 83: Screen capture of multiple loops of the ITA model [Friedenthal, 2006]

5.1.3 INTERACTIVE SCENARIO-BASED LEARNING AND DECISION-AID

The interactive features of the risk management tool allow the users to learn about the model dynamics by running individual scenarios and guiding the users in his/her understanding of the system dynamics. The tool includes a message window (see bottom of Figure 79) used to communicate the occurrence of external events to the users, for example, a simple scenario could be:

“You have been given a 20% decrease in your budget starting from the next fiscal year because of a change in corporate priorities.”

Given such a scenario, the user has to decide which parameters to modify in order to continue operating the system safely within the budget allocated. The scenarios do not have to be (and

most of the time should not be) hypothetical. A properly calibrated model can be used as real-time decision aid for managers faced with situations like that described above. Another scenario discussed in chapter 6 involves a governmental hiring freeze at NASA that was a real management concern. The model can be used to help decision-makers in maintaining the necessary amount of in-house capabilities given the governmental hiring constraints and the increasing amount of civil servants eligible for retirement.

5.2 ANALYZING IDENTIFIED RISK SCENARIOS

Scenarios are one of the main ways to analyze potential safety risks using the dynamic models created. This section discusses the identification of risks and the creation and analysis of associated risk analysis scenarios. The scenarios are used to better understand, prevent and control risk in the development and operation of complex systems.

5.2.1 RISK IDENTIFICATION

Traditionally, risks are identified in an ad hoc fashion. In some organizations, engineers and domain experts identify risks in their own area of expertise and often report them through a risk database for tracking and management. Risk identification can be facilitated through a combination of interviews with domain experts and brainstorming sessions, but it does not improve the completeness and rigorousness of the risk identification process. A more effective and rigorous risk identification method is provided by Leveson's STAMP model [Leveson, 2004]. The risks generated can be used in risk analysis scenarios both to prioritize the risks identified and to better understand them in order to devise mitigation measures.

5.2.1.1 Risk identification using the STAMP analysis

The STAMP analysis performed in steps 3, 4 and 5 of the STAMP-Based Risk Analysis process (See Figure 29) is used to systematically identify safety risks based on hazards and a safety control structure that includes the detailed safety-related responsibilities of each component [Leveson, 2003; Leveson, 2005]. During the ITA analysis [Leveson, 2005], the STAMP model was used to identify 250 safety risks based on a single high-level hazard, namely:

H1: Poor engineering and management decision-making leading to an accident (loss)

In order to avoid this hazardous state, the safety control structure has to enforce constraints on system behavior. Four system safety requirements and constraints (and associated subconstraints) were identified for H1, one of which being:

- **SC1:** Safety-related technical decision-making must be done by eminently qualified experts with broad participation of the full workforce
 - **SC1.1:** Technical decision-making must be credible (executed using credible personnel, technical requirements, and decision-making tools)
 - **SC1.2:** Technical decision-making must be clear and unambiguous with respect to authority, responsibility, and accountability
 - **SC1.3:** All safety-related technical decisions, before being implemented by the Program, must have the approval of the technical decision-maker assigned responsibility for that class of decisions
 - **SC1.4:** Mechanisms and processes must be created that allow and encourage all employees and contractors to contribute to safety-related decision-making

In order to identify risks, the four general high-level risks used in the STAMP process were customized for the ITA analysis, namely:

1. Unsafe decisions are made by or approved by the ITA
2. Safe decisions are disallowed (i.e., overly conservative decision-making that undermines the goals of NASA and long-term support for the ITA)
3. Decision-making takes too long, minimizing impact and also reducing support for the ITA
4. Good decisions are made by the ITA, but do not have adequate impact on system design, construction, and operation.

These four general risks were used in combination with the detailed control structure of the system including safety-related responsibilities for each component (as provided in the ITA implementation plan) to identify a list of detailed risks associated with the behavior of each component. A detailed description of the process and a complete list of identified risks can be found in [Leveson, 2005]. As an example, detailed risks associated with the responsibilities of the shuttle program/project manager are shown in Table 3.

PROGRAM/PROJECT MANAGER		
#	Safety-Related Responsibility	Inadequate Control (Risk)
1	Ensure that a full understanding of ITA is communicated through the program/project team.	Does not ensure a full understanding of ITA is communicated throughout the team and responsibilities for interfacing are not assigned or not communicated adequately.
2	Working with ITA, ensure that documentation, including the CoFR, is updated to reflect the required TA signature blocks.	Documentation is not updated to reflect required TA signature blocks.
3	Acquire STWH's agreement before applying technical standards and requirements or altering them.	Decisions made affecting safety are not communicated to ITA perhaps because p/p manager does not "buy into" ITA program or does not respect TWHs.
4		Applies or alters technical standards without appropriate engagement from STWH and DTWHs.
5	In event of a disagreement with TWH, explores alternatives that would allow achieving mutual agreement and, if cannot, raises issue up chain of command.	Interacts inappropriately with TWH or does not raise issues of disagreement up chain of command.
6	Obtain TWH agreement on technical decisions affecting safe and reliable operations prior to the Program or Project's application of technical standards and requirements and any alternation thereof.	Does not incorporate ITA fully in technical decision making perhaps because does not "buy into" ITA program or does not respect TWHs.
7	Provide the TWH with complete and timely access to program technical data, reviews, analyses, etc.	Does not comply with warrant holder requests and controls.
8		Does not allow (limits) complete and timely access to program technical data, reviews, analyses, etc. to technical warrant holders.

9	Support Trusted Agents and others and provide access to aspects of the project (reviews, etc.) necessary to perform their duties.	Penalizes employees for raising safety issues or handling safety concerns in performance appraisals or impose other career impacts.
10	Ensure safety has priority over programmatic concerns among those who report to him (line engineering, Shuttle SMA Manager, etc.).	Places safety second and pressures those reporting to him/her to do the same. Inaccurate understanding of current risk (complacency and overconfidence).
11		Abdicates responsibility for safety to Chief Engineer and Technical authority; does not adhere to safe engineering practices.

Table 3: Shuttle Program/Project Manager Risk [Leveson, 2005]

Some of the risks identified using the STAMP process can be analyzed directly using a static STAMP analysis as described in [Leveson, 2004]. Others are good candidate for the creation of dynamic risk analysis scenarios.

5.2.2 SCENARIO CREATION AND ANALYSIS

Once risks have been identified and selected as candidates for dynamic analysis, custom risk analysis scenarios have to be created. There are many ways to create risk analysis scenarios and the specifics are different depending on the type of risk to be analyzed. In most cases, the first step in creating a risk analysis scenario is to identify the model variables linked to the risk to be analyzed. Usually, this identification is straightforward because the model creation methodology is based on a STAMP control structure, which is also used to identify safety risks. If the model does not include the necessary variables, it may be an indication that either the risk can be better analyzed using a static model, or critical variables have not been included in the model and some backtracking and review is necessary. Similarly, the model creation and validation activities based on interactions with system stakeholders should naturally include variables that are highly connected to the concerns of interviewed stakeholders. Once variables have been identified, the key to scenario creation is to create a coherent and well-documented dynamic description for each risk analysis scenario.

5.2.2.1 Example Scenario: Contracting Effects on Safety and ITA Effectiveness

During the ITA project, one of the risks identified was the potential system safety and integration problems created by high levels of outside contracting for the space shuttle

program. The NASA structure is built upon civil servants providing technical oversight of contractor work. Consequently, one of the concerns is that high levels of outside contracting will drain the technical knowledge out of the civil servant workforce. If the civil servants and other in-house NASA employees do not perform enough technical work themselves, they will gradually lose their ability to perform technical oversight of contractor work, which should result in a risk increase. Furthermore, as the Technical Warrant Holders (TWHs) and Trusted Agents (TrAs) are in-house NASA employees, it was necessary to investigate the impact of high levels of contracting on the effectiveness of the ITA program.

In order to evaluate these risks, a scenario was developed to assess the effect of contracting on system technical risk and on the ability of ITA to remain credible and effective. A scenario was created where the fraction of work contracted out varied linearly from 4% to 96%. The results show that increased contracting does not significantly change the level of risk until a “tipping point” is reached where NASA in-house employees are not able to perform the integration and safety oversight that is their responsibility. Once that “tipping” point is reached, risk escalates substantially (see Figure 84).




Figure 84: Effect of Increased Contracting on Risk Level

Higher levels of contracting also affect the credibility and effectiveness of the ITA (see Figure 85). Initially, the effect is seen by the shortage of high-level technical leaders with safety

knowledge who can be assigned to the ITA (see Figure 86). As the shortage becomes acute and the Technical Authority (TA) is unable to find high-quality people to fill ITA positions, credibility and effectiveness start to suffer. Eventually, the gain of the reinforcing loops at the core of the ITA model becomes smaller than 1 (the tipping point is reached), resulting in a rapid deterioration of the situation with negative ripple effects all across the system. As the situation deteriorates, the TA will have difficulties in recruiting high-quality experienced employees, while contractors will have a relative abundance of high-level experienced employees whose safety knowledge and skills could prove very useful. However, without a strong in-house structure to coordinate and integrate the efforts of both NASA and contractor safety activities, effective risk mitigation is compromised.




Figure 85: Effect of Increased Contracting on ITA Effectiveness and Credibility




Figure 86: Effect of Increased Contracting on Availability of High-Level Technical Personnel

5.2.3 RECOMMENDATIONS AND POLICY DEVELOPMENT BASED ON SCENARIO RESULTS

In some cases, enough quantitative evidence may be available to make very specific quantitative recommendations based on the results of scenario-based risk analyses. Often, the insights and associated recommendations are qualitative and will include a claim for further data collection to better understand and quantify a specific phenomenon. For example, the model structure and scenario analysis uncovered the presence of a “tipping” point in the ITA model that could be triggered by the use of high levels of contracting. The ITA model did not have enough resolution and calibration to provide precise quantitative recommendations regarding the required level of contracting. However, the scenario analysis performed identified the potential problem and the need for further study on the amount and type of contractor oversight necessary to avoid reaching the tipping point identified in the risk analysis scenario. A final recommendation to NASA was:

“For projects in which significant contracting is anticipated, careful study of the types and amount of oversight needed to avoid reaching the tipping point will help with NASA’s planning and staffing functions. The answer may not be a simple ratio of in-house expertise to contracting levels. Instead, the type of project as well as other factors may determine appropriate expertise as well as resource needs.[Leveson, 2005]”

More examples of scenario formulations, as well as recommendations and policy formulations are provided in chapter 6.

5.2.4 POLICY ROBUSTNESS AND SCENARIO SENSITIVITY ANALYSIS

Before recommendations are made and/or policies are suggested based on risk analysis scenarios, it is critical to ensure that the recommendations and policies are robust to uncertainty in model parameters and decision rules. Monte-Carlo based sensitivity analyses are used throughout scenario analyses to ensure that the recommendations and policies have a positive impact on safety in a large majority of cases, regardless of uncertainty in model parameters and decision rules. More examples of sensitivity analyses and outcome distributions are provided in the case study of chapter 6.

5.3 PREVENTING RISK INCREASE AND IMPROVING ROBUSTNESS TO SAFETY EROSION

Designing and operating safe systems is the main objective the STAMP model and of the dynamic risk management modeling techniques introduced in this thesis. In classical hazard analysis, the goal is to eliminate and control hazards in the design and operation of complex systems. The author previously developed a methodology to evaluate system architectures from the conceptual design phase according to their inherent risk mitigation potential [Dulac, 2004]. One of the objectives of this research is to further assist hazard mitigation by helping to prevent and control the safety erosion and the associated risk increase that can happen in complex systems as changes occur in the structure of the system and the behavior of its components and participants. The most effective way to avoid safety erosion is to develop inherently safe systems where the possibility for risk increase over time is eliminated from the beginning of the system lifecycle. In some systems and situations, it is impossible to completely prevent a risk increase from the beginning through risk elimination and mitigation. In fact, many of the risks identified during an analysis with a focus on organizational factors cannot be completely eliminated from the system. For example, one of the risks identified during the ITA analysis was having insufficient NASA in-house capability to perform quality hazard analyses. For this risk, mitigation involves ensuring that employees continuously have

adequate knowledge and training. This risk cannot be eliminated entirely. For these types of risks that require time-dependent mitigation, the objective is to control risk increase and at the same time, improve the robustness of the system to some levels of risk increase.

5.3.1 PREVENTING TIME-DEPENDENT SAFETY EROSION

Figure 87 shows a summary of some possible outcomes of risk increase, along with associated mitigation strategies (ranked in order of effectiveness). Static control structure analysis and dynamic scenario-based risk analysis (as presented in [Leveson, 2004; Leveson, 2005] and in this thesis, respectively), combined with the implementation of effective risk mitigation strategies and policies is a preferred method when it is not possible to completely eliminate potential risk increase. The model creation and analysis methodologies and guidelines provided in this thesis all contribute to a better understanding, prevention and control of the migration of systems toward high-risk states.

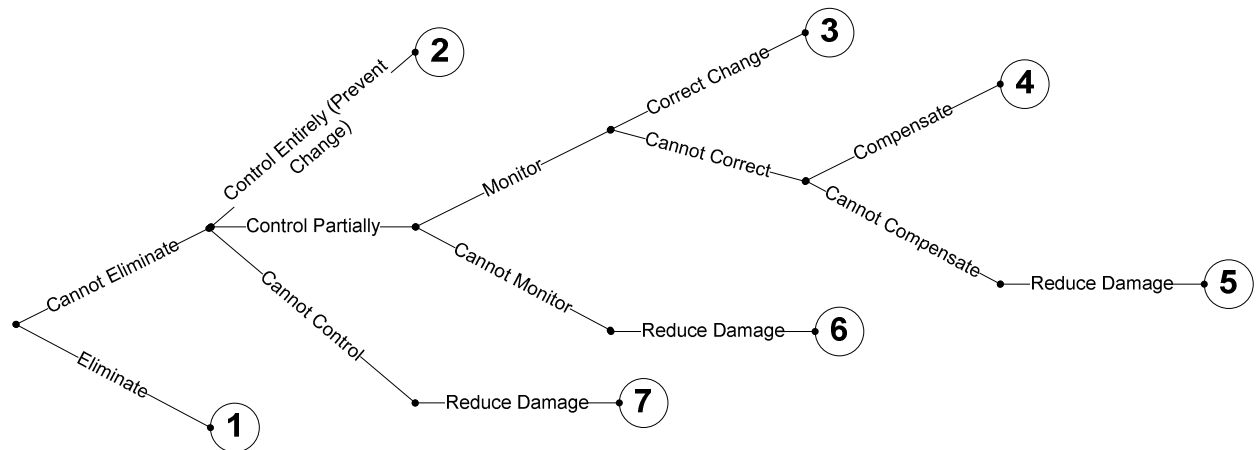


Figure 87: Possible outcomes of the migration toward high risk process

In addition to preventing the migration of systems toward high risk through risk analysis and associated mitigation strategies, it is possible to improve the inherent robustness of systems to some level of risk increase.

5.3.2 IMPROVING SYSTEM ROBUSTNESS TO RISK INCREASE

Improving the robustness of the system to risk increase can be done different ways. In this section, three distinct types of robustness are discussed: 1) passive robustness, 2) semi-active robustness, and 3) active robustness. Robustness should be interpreted in the larger context of the mitigation of time-dependent risks that cannot be completely eliminated in the system design.

5.3.2.1 Passive (Safety Factor) Robustness

Passive (safety factor) robustness is the typical way by which safety was traditionally implemented in engineering systems. The main operating principle is a two step process: 1) System designers estimate the largest loads to which the system may be subjected. 2) The system is designed to withstand this load multiplied by a chosen safety factor.

It should be emphasized that this approach can be used for the technical design of a system, but also for some of the organizational aspects. For example, in some systems, the maximum number of problems under investigation at one time may dictate the size of the system safety workforce. This type of robustness is the least likely to be effective against time-dependent risk increase and erosion of safety margins because the pressures for more performance and efficiency work against the use of large static safety margins.

5.3.2.2 Semi-Active Robustness

Semi-active robustness (or resistance) is a type of robustness where changes that go counter to the safety of a system are resisted. Semi-active robustness does not require a precise action from decision-makers, as the system has built-in protection (whether it was built-in intentionally or not) against changes that could reduce the safety margins of the system. This process can be pictured as damping in a mechanical system, where any disturbance from the equilibrium position is countered by friction from a damper (see Figure 113 for a pictorial). The faster the change, the more resistance the system provides. There are advantages and disadvantages to this type of robustness. Unless the semi-active robustness processes are carefully planned in the system design and are well-focused toward safety, the resistance to change may become a general characteristic of the system and changes that can be beneficial to safety (and other performance criteria) will also be resisted. Creeping bureaucracy is an

example of a non-specific resistance to change process that may be beneficial to safety in some very specific situations (for example in the difficulty at NASA to transfer safety employees toward other functions), but that may hinder performance (and safety) improvements in most other situations. Another disadvantage is that while semi-active robustness may work great against brute force change attempts (ex: firing 20% of the NASA civil servant workforce), it may not be very effective against slow and steady change attempts (ex: not replacing the employees that retire), which is usually the way safety margins are eroded.

5.3.2.3 Active Robustness

Active robustness (or active control, compensation) is a type of robustness where the system is monitored, and changes that go counter to system safety are detected and actively compensated. There are three necessary conditions for a system to be able to exhibit active robustness. The first condition is that the system can be monitored with a sufficient level of accuracy (the observability criterion discussed previously). The second condition is that the system monitoring be able to provide early warning of a system heading toward a hazardous state *before* it is reached. The third condition is that there exist control actions that will allow the system to head back toward safe levels of risk (the controllability criterion discussed previously). The last condition is assumed to be fulfilled, given that enough analyses were performed to understand the system functioning and the risk mitigation strategies available. However, it is also assumed that the longer the early warning is available, the more likely it is that control actions directed toward risk decrease will be successful. Consequently, monitoring and early warning are important topics and are discussed in the next section.

5.4 IMPROVING RISK MONITORING IN COMPLEX SYSTEMS

As risk in complex systems is neither directly observable nor directly measurable, it is necessary to use proxy metrics and early indicators to monitor and evaluate the level of risk in a system. The objective is to create an early warning system, or a “canary-in-the-coal-mine” [Leveson, 2004] that alerts analysts and decision-makers when the system is heading toward a high-risk state.

5.4.1 IDENTIFYING LEADING INDICATORS

Leading indicators are needed in some systems to detect risk increase before a hazard state is reached and an accident occurs. Identifying effective leading indicators for risk increase is a matter of searching for model variables that are observable in the real world, and that provide good correlation with risk increase dynamics (assuming that early indicator variables were identified as important to the system dynamics and consequently included in the model). Once potential indicators are identified, their usefulness must be validated through data collection and analysis in the field.

During the ITA analysis, it appeared that many good indicators of increasing risk were available in the model. However, many of these indicators become useful only after a significant risk increase has occurred, i.e., they are lagging rather than leading indicators. For example, one of the potential early indicators investigated was the number of outstanding accumulated requirements waivers. Throughout the life of the shuttle system, a large number of system requirements waivers were granted. Waiver issuance became the path of least resistance in order to continue flying the system on schedule when all requirements could not be met. Once issued, the waivers were rarely revisited, so a large number of those waivers started to accumulate and over 2300 high-criticality waivers were on file at the time when the Columbia accident occurred [Gehman, 2003]. Consequently, the number of waivers seemed to be a potentially effective risk indicator. However, after analyzing the behavior of associated model variables, it appeared that the requirements waiver accumulation pattern is indeed a good indicator, but only becomes significant when risk starts to rapidly increase (see Figure 88), thus casting doubt on its usefulness as an effective early warning.




Figure 88: Waiver Accumulation Pattern for ITA Risk Increase

Alternatively, the number of incidents/problems under ITA investigation appeared to be a more responsive measure of the system heading toward a state of higher risk (see Figure 89). A large number of problems under investigation results in a high workload for trusted agents, who are already busy working on project-related tasks. Initially, the dynamics are balancing (See Figure 90), as ITA personnel are able to increase their problem investigation rate to accommodate the increased investigation requirements. As the investigation requirements become higher, corners may be cut to compensate, resulting in lower quality investigation resolutions and less effective corrective actions. If investigation requirements continue to increase, the TWHs and trusted agents become saturated and simply cannot attend to each investigation in a timely manner. A bottleneck effect is created that makes things worse through a fast acting reinforcing loop (see Figure 90).

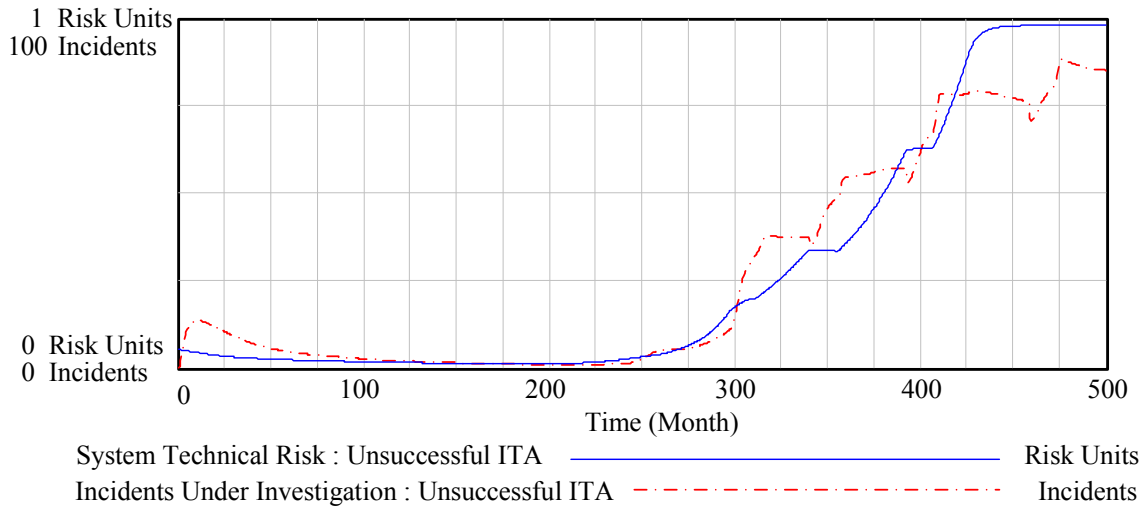


Figure 89: A Potential Early Indicator: Problems under Investigation

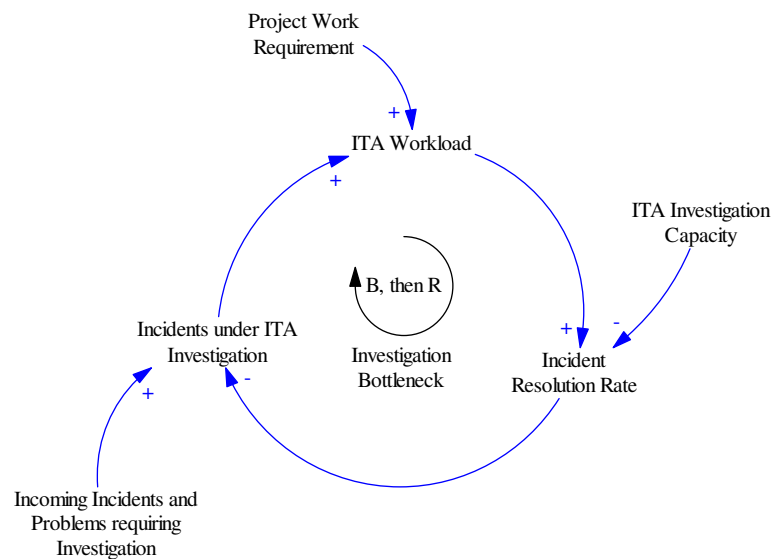


Figure 90: The Balancing Loop Becomes Reinforcing as Workload Increases

5.4.2 LEADING INDICATOR SENSITIVITY

The more sensitive a metric is, the more likely it is to be able to identify risk increase early, when there is more time for the risk increase to be addressed and mitigated. However, very sensitive indicators and alerting systems may have a tendency to trigger more false positive warnings. Tradeoffs have to be made between the sensitivity of indicators and their ability to

provide early warning. In order to further improve accuracy of early warning systems and reduce type I and type II errors, smaller dynamic models can be used. In fact, once potential metrics have been identified and data collection has begun, smaller, more focused dynamic models should be used to help understand and calibrate the “metrics dynamics” to that of the real system. For example, Figure 91 shows a sample model that can help in understanding the dynamics of problem resolution and waiver accumulation. The model is derived from the “Safety Capability Trap” [Senge, 1990] and Repenning and Sterman’s model of the capability trap in process improvement [Repenning, 2002]. Further research will address the problems of data collection and calibration of early indicators, as well as sensitivity of warning systems.

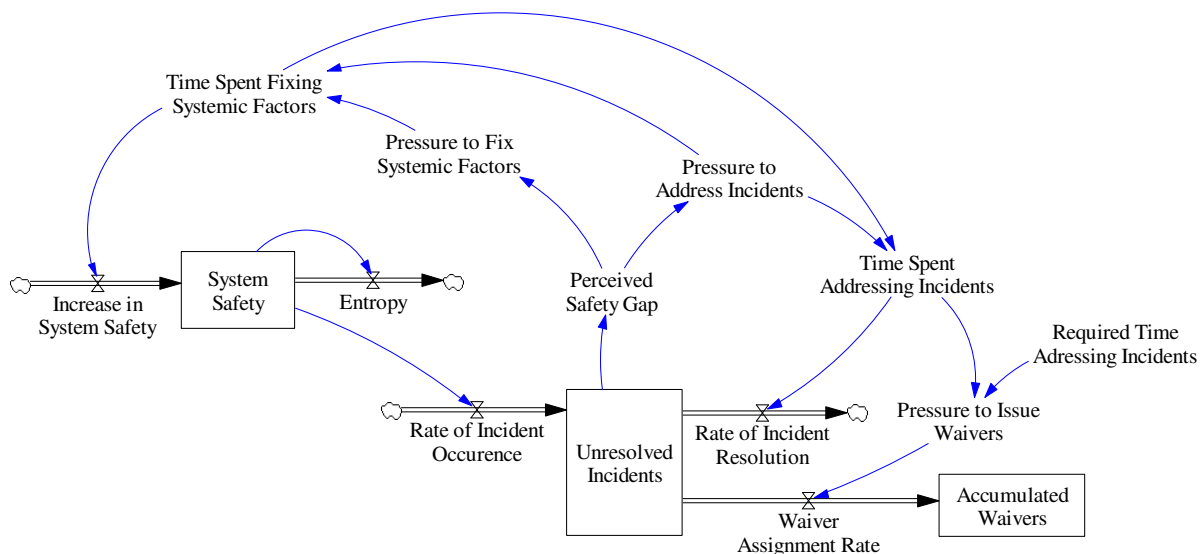


Figure 91: Example of a Smaller "Metrics Dynamics" Model

5.5 SUMMARY

This chapter built upon the model creation methodology chapter by providing an overview of some of the results and benefits that can be obtained from the models, along with methods to achieve these results and benefits. Some of the results used in this chapter and previous ones were based on the ITA analysis performed for the Office of the Chief Engineer at NASA. The next chapter (chapter 6) provides a demonstration of the entire model building and

analysis methodology using a completely different example, namely the risk analysis of the new space exploration enterprise performed for the NASA Exploration Systems Mission Directorate.

CHAPTER 6: CASE STUDY - RISK MANAGEMENT IN THE DEVELOPMENT OF NASA'S SPACE EXPLORATION SYSTEM

The case study presented in this chapter provides an example of the model-building methodology introduced in this thesis. The methodology was used to perform a risk management study for the newly formed Exploration Systems Mission Directorate (ESMD) at NASA. ESMD is the directorate in charge of the development and implementation of the space exploration system as defined by the President's Vision for Space Exploration (VSE). Among other objectives, the Vision for Space Exploration will require NASA to develop human-rated launch and landing systems for the first time since the 1970s, as well as a versatile Crew Exploration Vehicle (CEV) adaptable for Lunar and Martian exploration. Readers interested in more detailed history and background for the NASA ESMD project should refer to Appendix G.

6.1 ESMD MODEL-BUILDING METHODOLOGY

In this section, the steps of the model-building methodology introduced in chapter 4 are followed to demonstrate their use in facilitating model development. The model-building steps are very data intensive. A combination of quantitative and qualitative data taken from interviews with experts, NASA documents and existing literature was used throughout the model-building steps. Detailed description and references for the data sources used in this example can be found in Appendix H.

6.1.1 STEP 1: ESMD INITIAL SYSTEM CHARACTERIZATION

Before the modeling begins, step 1 is used to define the high-level characteristics (and boundaries) of the system and the model to be built. The following is an example of such system/model characteristics:

- The ESMD model is a safety-centric development model. Consequently, the associated production units are system development completion rates and fractions.
- The focus of the ESMD model is on the first flight of the Crew Exploration Vehicle, scheduled for the 2012. The planned development time for CEV is 8 years. In order to capture the entire development time period, the start date of the model and simulation is January 1st, 2004, and the end date is July 1st 2016.
- The workforce used in the model is that of the Exploration Systems Mission Directorate. Data on workforce size, distribution, hiring and attrition rates, experience, etc. are derived from the data publicly available on NASApeople [NASA, 2006].
- The budget for system development is exogenous and derived from budget request forecasts for FY2004 to FY2007.
- The outsourcing ratios are derived from procurement statistics using the data available for FY2002 to FY2004 (the only procurement data publicly available)

Additional model characteristics and conventions are defined in Appendix I. Further model characteristics are defined later in the model creation process.

6.1.2 STEP 2: MAPPING OF STATIC SAFETY CONTROL STRUCTURE TO GENERIC DYNAMIC COMPONENTS

Before performing step 2, a draft of the system control structure must be available. For the ESMD study, the initial control structure was created during discussions with project sponsors (see Figure 138 in Appendix). The safety control structure was mapped to the generic dynamic components taken from the repository of generic components provided in Appendix E. The result of mapping the generic dynamic components to the NASA/ESMD control structure is show in Figure 92. This mapping was slightly modified a few times during model refinement and validation based on the input of ESMD domain experts interviewed (see H.1.1 for interview details).

Table 4 provides a summary of the mapping between NASA/ESMD control structure components and generic components taken from the repository.

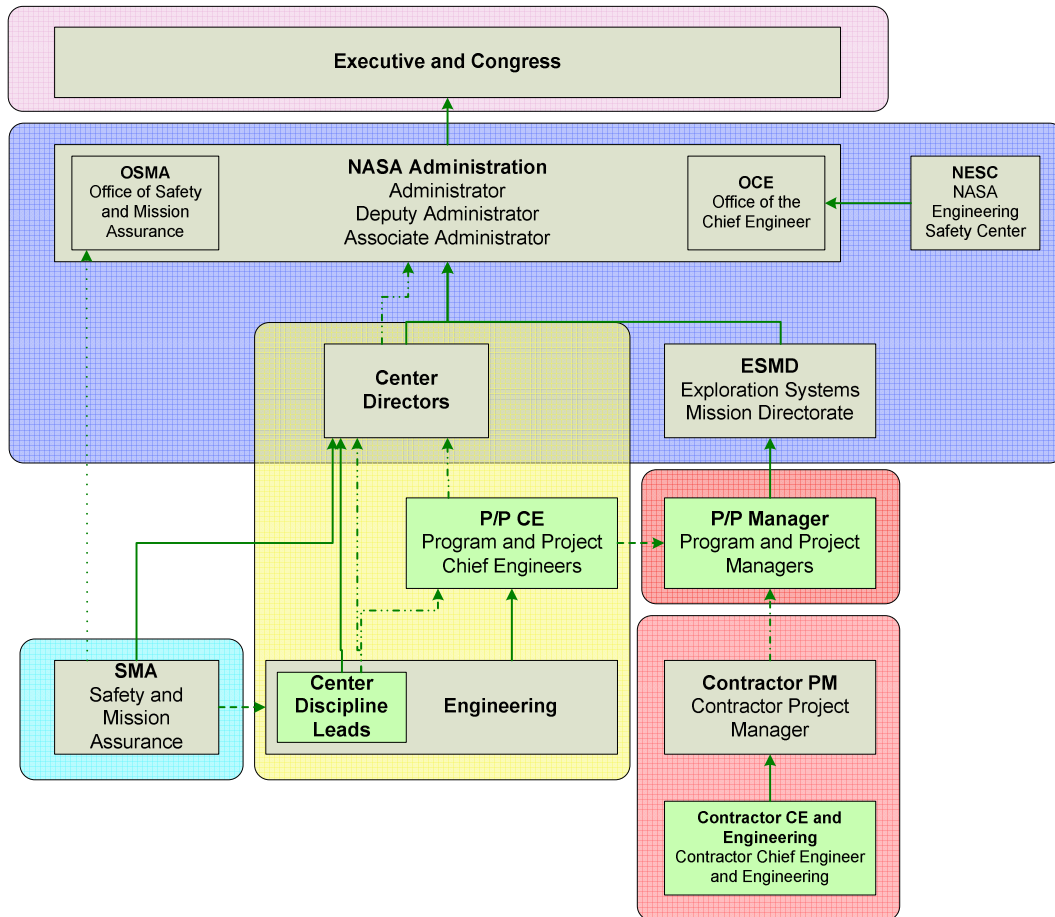


Figure 92: Mapping generic components to the NASA/ESMD structure

NASA/ESMD Control Structure	Generic Component
Executive and Congress	Congress and Executive
NASA Administration/NESC	Administration
Exploration System Mission Directorate	Administration
Center Directors	Administration/Engineering Management

Safety and Mission Assurance	Safety and Mission Assurance
Engineering and Chief Engineers	Engineering Management
Program/Project Manager	Program/Project Management
Contractor PM/CE/Engineering	Engineering Procurement

Table 4: Summary of mapping between NASA/ESMD control structure components and generic components

6.1.3 STEP 3: REFINEMENT OF DYNAMIC SAFETY MODEL STRUCTURE

Once the mapping from the NASA/ESMD control structure to the generic dynamic components is performed, the generic components must be assembled to form the structure of the new dynamic safety model. The components and key variables were renamed to better match the initial NASA/ESMD structure. The result of the rearrangement and renaming of the generic components is shown in Figure 93. Sub-components related to the processes performed within each generic component are also shown (see boxes within boxes in Figure 93).

Generally, upward arrows represent feedback channels, while downward arrows represent control actions. The operating principle follows that of the STAMP accident model, that is, safety is achieved by performing the control actions necessary to ensure that safety constraints are enforced throughout the system lifecycle. Horizontal arrows represent lateral information transfer between components.

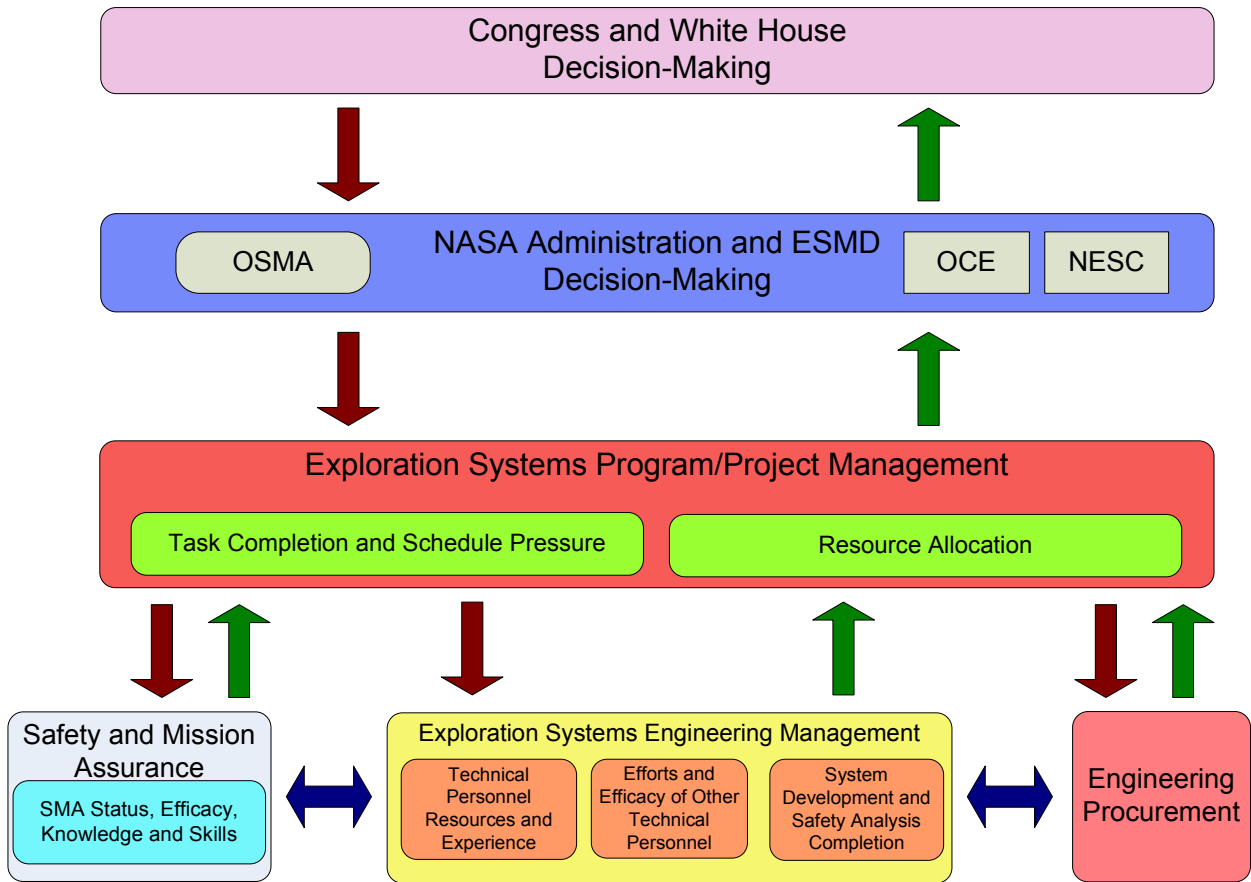


Figure 93: Components and Structure of the NASA/ESMD Model

6.1.4 STEP 4: MAPPING OF PRESSURES, INFLUENCES AND REPORTING CHANNELS

As documented in chapter 4, the mapping of pressures, influences and reporting channels facilitates the final assembly of the components into an integrated model. It also allows the verification of the structure and I/O interface of the generic components and their selection before additional modeling effort is deployed. A sample of the mapping performed for the NASA/ESMD model based on the technique introduced in chapter 4 is shown in Figure 94 and Figure 95. Figure 94 shows the mapping of performance, resource, and safety pressure on the ESMD model structure. For example, resource pressure starts from the *Congress and White House* component, and flows down to the *program and project management* component through the *NASA Administration and ESMD* component. Once the program and projects are subjected to resource pressure from above, pressures are “distributed” in every component below as resource scarcity affects the system development, safety activities, as well as procurement from outside contractors.

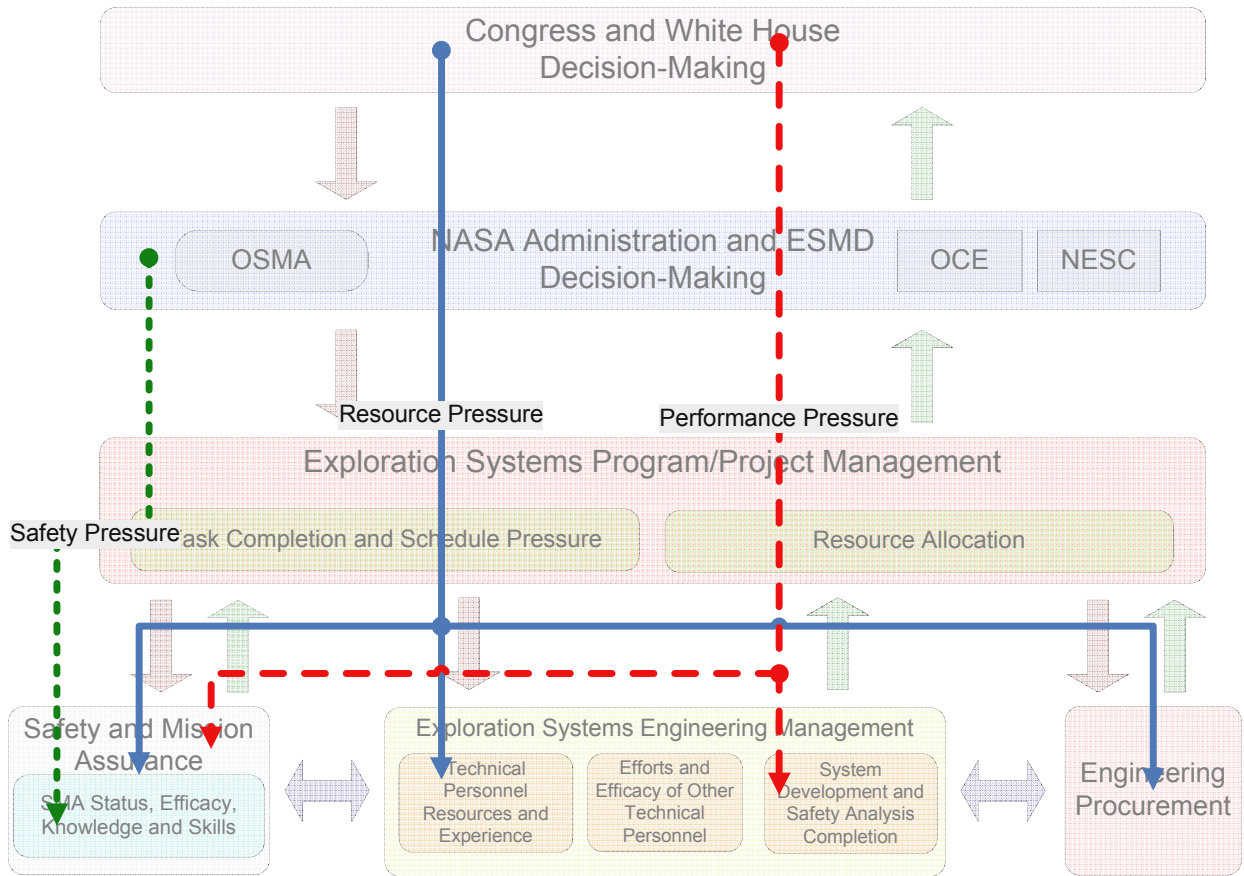


Figure 94: Mapping of performance, resource and safety pressures on the ESMD model structure

Figure 95 shows the flow of progress and problem reports in the ESMD model. For example, progress reports start within the system development subcomponent, and flow all the way up to the *Congress and White House* component, through the *Program/Project Management* and *NASA Administration* components. The format of the information transferred in progress reports changes on the way up to congress. Generally, at the bottom, the information is very detailed (ex: which specific design tasks were completed by which deadline), while the progress reports at the administration and congress level abstracts away the low-level details, but the essential reporting flow remains.

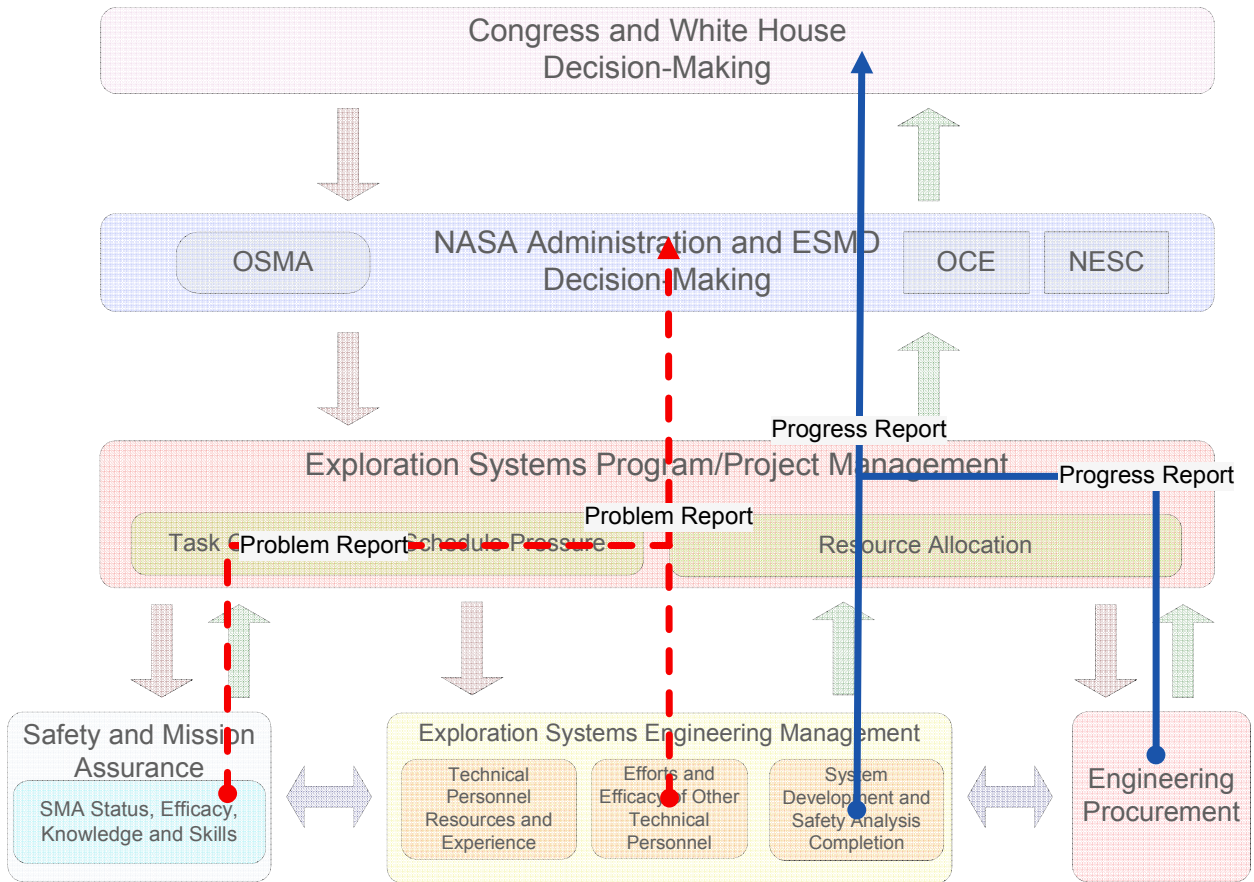


Figure 95: Mapping of progress and problem reports on the ESMD model structure

6.1.5 STEP 5: DATA COLLECTION AND COMPONENT CALIBRATION

Interviews with domain experts were extensively used to refine and validate the model components. The guidelines for final component creation and calibration introduced in chapter 4 were used throughout the component customization process. In step 5, the internal causal structure and decision rules of the generic components are customized and calibrated through data collection and interaction with domain experts and system stakeholders (complete data collection protocols and source are documented in Appendix H). Many iteration cycles were used for each component. After each round of interview, the components were updated before a following round was scheduled. Early interactions with domain experts focused on refining the high-level component causal structure. Later interactions focused on the stock and flow structure of the model, and the specific decision rules and equations used to define model behavior. The components were updated until we

converged to a structure deemed acceptable to a majority of domain experts interviewed. Disagreements between experts become the subject of further model tests and sensitivity analyses. Multiple sources of data (e.g. quantitative budget and employment data, accident reports, risk management literature, etc.) are superimposed to ensure that agreement or disagreements between domain experts are supported by empirical evidence. The detailed documentation for the executable component ultimately used in the model is provided in Appendix K.

6.1.5.1 “Free Component Method”

The “Free Component Diagram” (see chapter 4) method was used to create and verify the component interfaces. Figure 96 shows a sample generic (Administration) component taken from the repository. As can be observed, the interfaces of the component in Figure 96 exhibits very high correlation with the “virtual container” created for the NASA Administration and ESMD component using the influence and pressure mapping shown in Figure 94 and Figure 95. Additionally, the initial components and the final customized components can exhibit equilibrium behavior, which later facilitates the assembly and testing of the model.

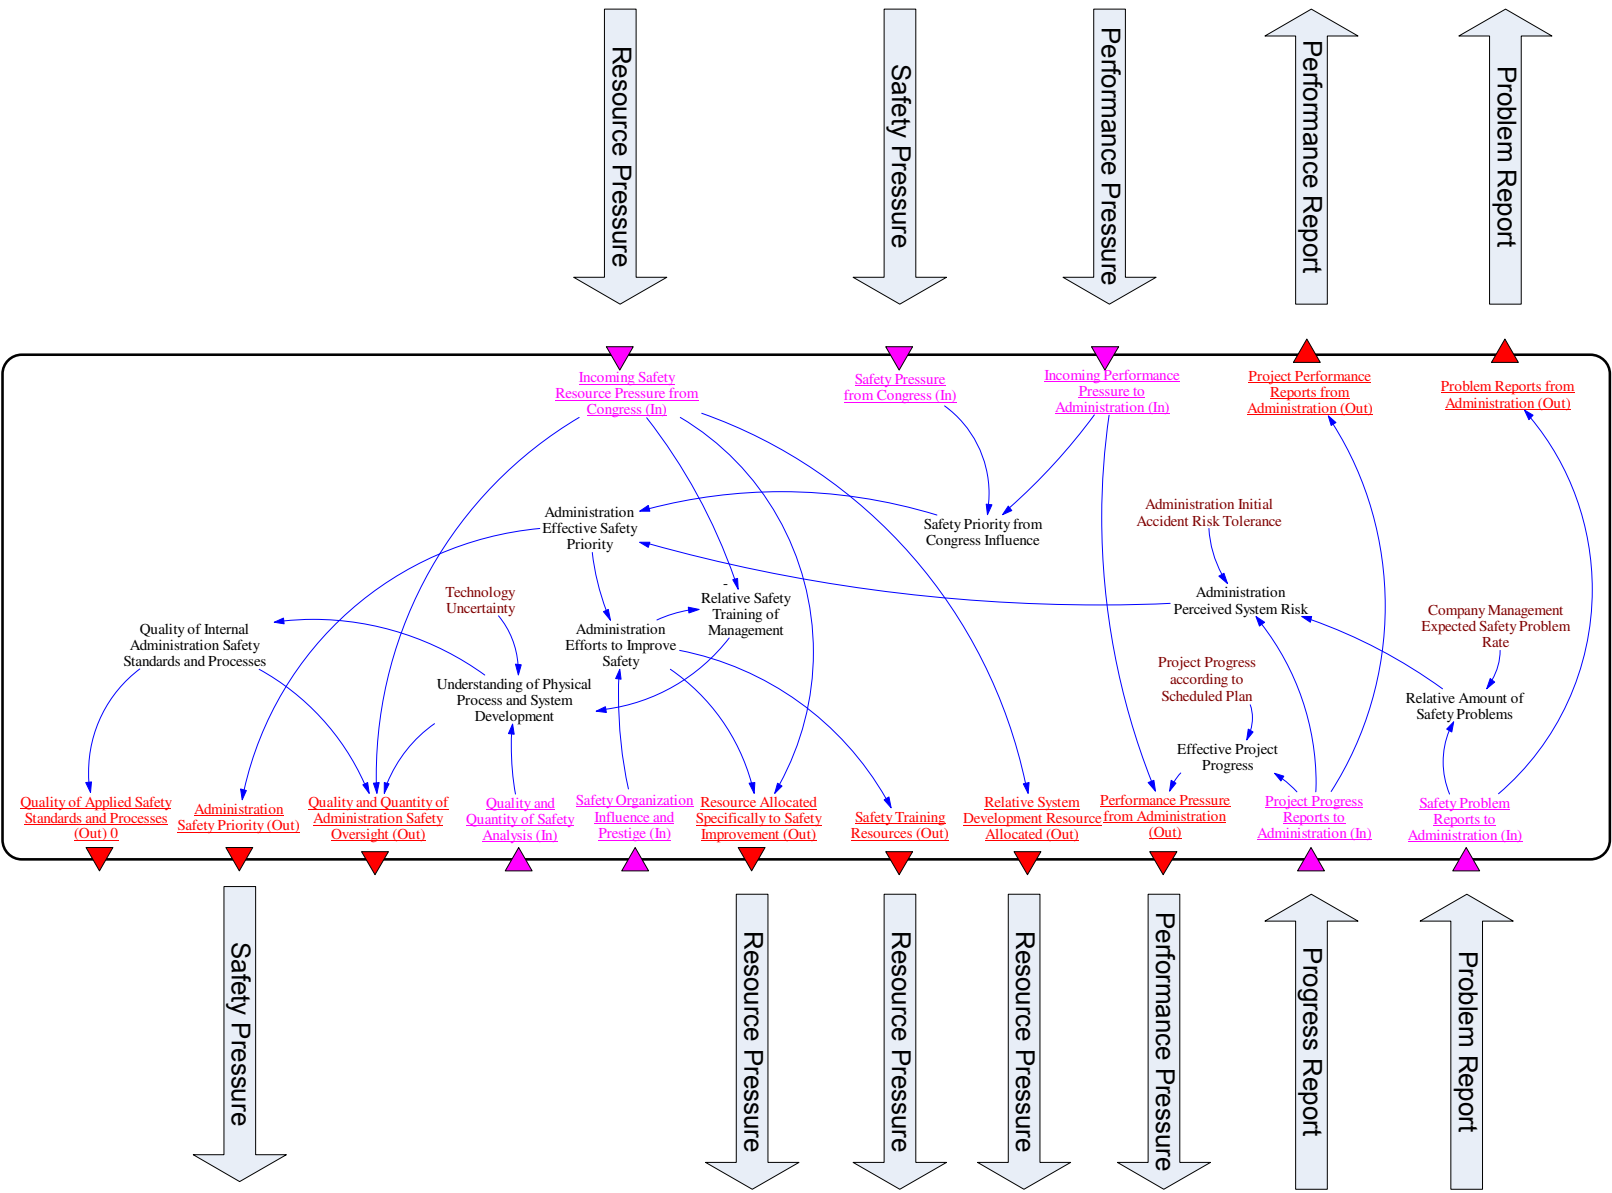


Figure 96: Generic Administration Component for Development Organization

6.1.5.2 Example customization: Engineering – Technical Personnel Resource and Experience

As mentioned previously, the initial components were extremely useful to guide the discussion and validation, but still had to be modified to fit the system and the purpose of the project. Figure 97 and Figure 98 show a sample initial component and the result of some customization based on expert interviews and data collection. Figure 98 highlights some of the specific changes made. Namely:

- 1- Initially, the component was divided into two distinct structures. The in-house employee structure keeps track of the flow of in-house technical employees and the contractor structure keeps track of the flow of contractor technical employees. However, NASA is different than typical development companies as it uses support contractors extensively. A more appropriate division is that of civil servant and support contractors (procurement contractors are handled in a different structure). The ratio of support contractor to civil servant varies depending on programs (the ratio can go up to 8:1 in some programs), but support contractors are also in-house employees and work directly with civil servants at NASA centers. They are defined in NASA slang as “blue badges” (their only apparent difference from civil servants). In most cases, civil servants will be required to oversee support contractors as civil servants are usually more permanent, but oversight is not a large portion of their daily activities. In effect, the work of civil servants and support contractors is highly integrated. The use of support contractors increases NASA’s flexibility because civil servants are difficult (almost impossible) to lay off, and recent government hiring freezes have increased the difficulty in hiring additional civil servants. Support contractors, on the other hand are very easy to hire and lay off as long as budget is available. Consequently, the model structure was modified to consider support contractors as in-house employees (See number 1 in Figure 98).
- 2- Based on recommendations from NASA interviewees, the hiring rate for in-house employees was disaggregated into inexperienced and experienced employees (see number 2 in Figure 98). The rationale for this change is that according to NASA employees, an experienced employee hired takes only a couple months to adjust and

become productive. On the other hand, recent graduates can take up to a couple years before they become fully productive.

- 3- Other minor changes were made and exogenous factors were added based on the inputs from NASA interviewees (see number 3 in Figure 98). For example, the “hiring freeze” policy at NASA (and other government agencies) limits the number of civil servants that can be hired. Data was gathered through interactions and additional documents about the hiring and attrition (lay offs, retirement, early retirement) rates, demographics of NASA workforce, etc. The data was used to calibrate the model.
- 4- Another change based on interactions with NASA interviewees involves the linking of early retired NASA employees with the experienced contractor hiring rate (see number 4 in Figure 98). A significant number of NASA employees eligible for early retirement leave NASA and go to work for contracting companies, as the positions available in the private sector offer higher compensation. Consequently, a large number of early retirements at NASA increases the availability of highly qualified employees to contracting companies.

Changes as those documented above were performed on every component based on interviews and further data collection.

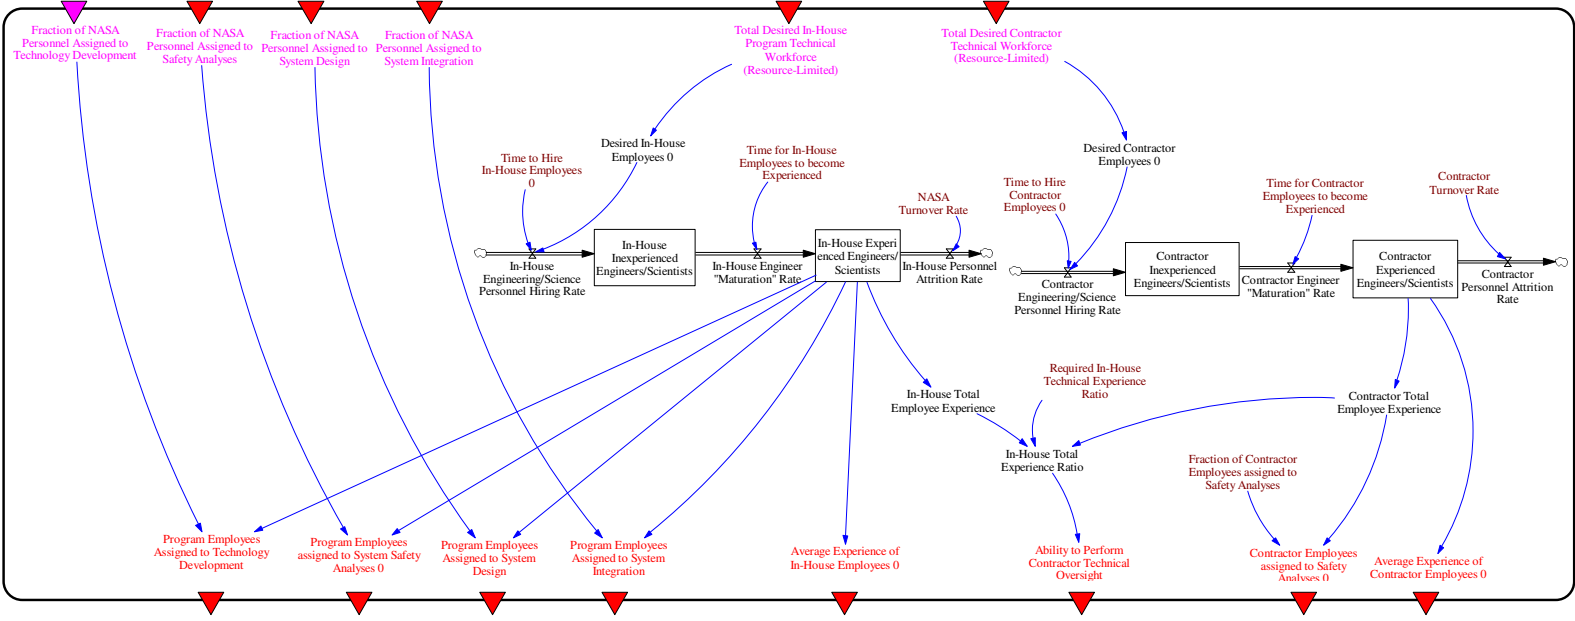


Figure 97: Initial Technical Personnel Component

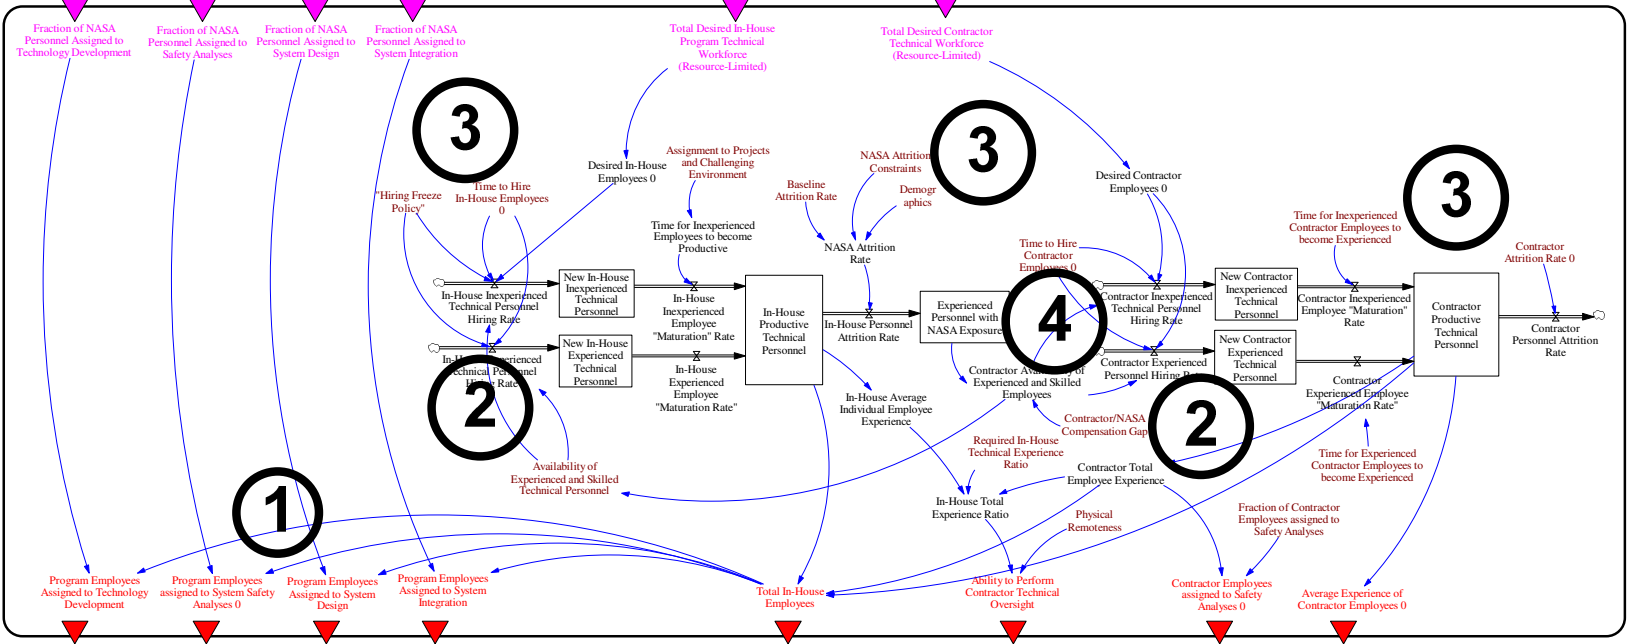


Figure 98: Technical Personnel Component after two iteration cycles

6.1.6 STEP 6: COMPONENT TESTING AND CONFIDENCE-BUILDING ACTIVITIES

Once an executable model was created for each component, various tests were used to improve confidence in the behavior of the components. The tests used were described in the methodology of chapter 4. Stress-testing is used throughout by randomly varying the inputs and exogenous values of each component and monitoring the output values. Stress testing does not provide a formal analytical proof of the correctness of the differential equations used, but it improves confidence in the model behavior, which is critical.

Accident dynamics does not apply to system development modeling, so the accident conditions introduced in the methodology (chapter 4) cannot be used. However, components are calibrated to ensure that equilibrium behavior is possible at the boundary of each component (given equilibrium inputs). An example of such equilibrium behavior for the Administration component is shown in Figure 99. Specifically, the component equilibrium behavior shown Figure 99 is obtained by using constant inputs for the resource and safety pressure to the NASA Administration component, as well as progress reports from program management that are exactly on schedule. Intent rationality tests are also performed by injecting changes in component inputs and ensuring that the consequences are consistent with the intended rationality of decision-makers within the component.




Figure 99: Sample equilibrium behavior for the Administration component

6.1.7 STEP 7: COMPONENT-BASED MODEL ASSEMBLY

Step 7 involves the assembly of components into an integrated model. The refined structure obtained in earlier steps (see Figure 93) provides a map of the necessary connections. The standard interfaces and connection points created for each component facilitate the assembly of executable, previously-tested and semi-validated components into an integrated model. Appendix J provides an example of a three-component assembly created by connecting component interfaces through generic connectors. The components used in the example are: *Congress and Executive*, *NASA Administration and ESMD*, and *Exploration Program and Project Management*. The entire model is composed of nine different components and subcomponents and was assembled in this fashion.

6.1.8 STEP 8: MODEL TESTING AND CONFIDENCE-BUILDING ACTIVITIES

As mentioned previously, the last step of the model-building methodology is highly related to model analysis. This is especially true of development-centric models, because the accident response scenario testing does not apply (see chapter 4). Problems with model formulations are often found during preliminary analysis, and backtracking is necessary to correct problems before further analysis can be performed. The main drawback to this analysis-testing-analysis cycle approach is that preliminary analyses have to be re-done when problems are found and corrected. Fortunately, modern software packages allow the saving of analysis parameters, which facilitates the re-run of analysis scenarios when changes are made. Nevertheless, system-level stress-testing was performed extensively on the ESMD model by randomly varying exogenous variables (within their allocated ranges) while monitoring state variables and component interfaces. In addition, intent rationality tests were performed by varying a single exogenous variable at a time and analyzing model results. Surprise behavior was investigated and corrections were made accordingly. Time and budget considerations for the project limited the number of iteration cycles possible once the model was functional. Nevertheless, results were shown to a number of NASA experts to perform a first order validation of model behavior.

6.2 INDIVIDUAL ESMD COMPONENT DOCUMENTATION

Model documentation can be performed at various levels, all the way from high-level documentation of the functions performed by the component, down to the individual differential equations and decision-rules used. In creating the ESMD model, the final components were obtained by customizing the generic components available in the repository of generic components of Appendix E. For large models, documentation can be very extensive. Consequently, this section provides a very brief description of the function of each component and subcomponents used in the ESMD model. Extensive component-based model documentation including model structure and detailed description can be found in Appendix K. Full software-generated model documentation is available from the author upon request.

6.2.1 CONGRESS AND EXECUTIVE (WHITE HOUSE) COMPONENT

The Congress and Executive component is responsible for defining the vision for the US space exploration enterprise, as well as providing the level of funding necessary to develop and operate a safe exploration system. Many external factors affect the ability and willingness of the Congress and Executive to define and implement a realistic (and safe) system. Some of these external factors include: *Political Uncertainty*, *Time Horizon of Political Objectives*, and the *Executive Branch Initial Leadership and Vision for Program*.

Congressional and Executive dynamics are extremely complex. In this model component, we did not attempt to precisely quantify the relationships between different variables. Instead, we merely tried to improve our confidence in the existence of these relationships. In the baseline model, the variables in this component are in equilibrium, that is, unless the values of external variables in this component are modified, the component will have negligible effect on the dynamics of the integrated model. Nevertheless, all the relationships have been implemented in the model, thus allowing us to test Congress and Executive-related policies as well as scenarios where external events affect national priorities and Agency funding.

6.2.2 NASA ADMINISTRATION AND ESMD COMPONENT

The purpose of the NASA administration and ESMD component is to define the agency level policies, requirements, and guidelines that will enable the development of a safe and

successful exploration system. The Agency receives directives and funding from Congress, and allocates resources according to program needs. The primary function of the NASA Administration and ESMD component is to allocate resources (human and material/financial) to different programs while respecting the constraints imposed by Congress and presidential administrations.

6.2.3 EXPLORATION PROGRAM AND PROJECT MANAGEMENT COMPONENT

The purpose of the Exploration Systems Program Management component is to reproduce the behavior of program and project managers during real system development. Program managers have to ensure that the system under development meets technical requirements (including safety and performance requirements) while remaining within budget and on schedule. Program managers use multiple control actions to achieve these objectives, including reshuffling schedules, reallocating resources (human, financial and material), and applying various pressures to lower-level managers, engineers and other technical workers.

6.2.4 SYSTEM DEVELOPMENT COMPLETION AND SAFETY ANALYSES COMPONENT

The System Development Completion and Safety Analyses (SDCSA) component is at the core of the ESMD model. It includes three different task completion flows that have to be synchronized and coordinated to produce a final integrated product. The three flows are: 1) Technology development, 2) System development, and 3) Safety analyses. The task completion flows depend on many factors including the size, experience, and overall productivity of the workforce allocated to perform technology, design, and safety-related activities. The timing of the flows is critical. Late technologies cannot be used in design without significant development delays. Similarly, late safety analyses might delay design or might simply not be used in design decisions, resulting in an unsafe system. The SDCSA component regulates the flow of task completion for the three types of activities mentioned above. Figure 100 provides a schematic of the three task completion flows merging into an integrated product.

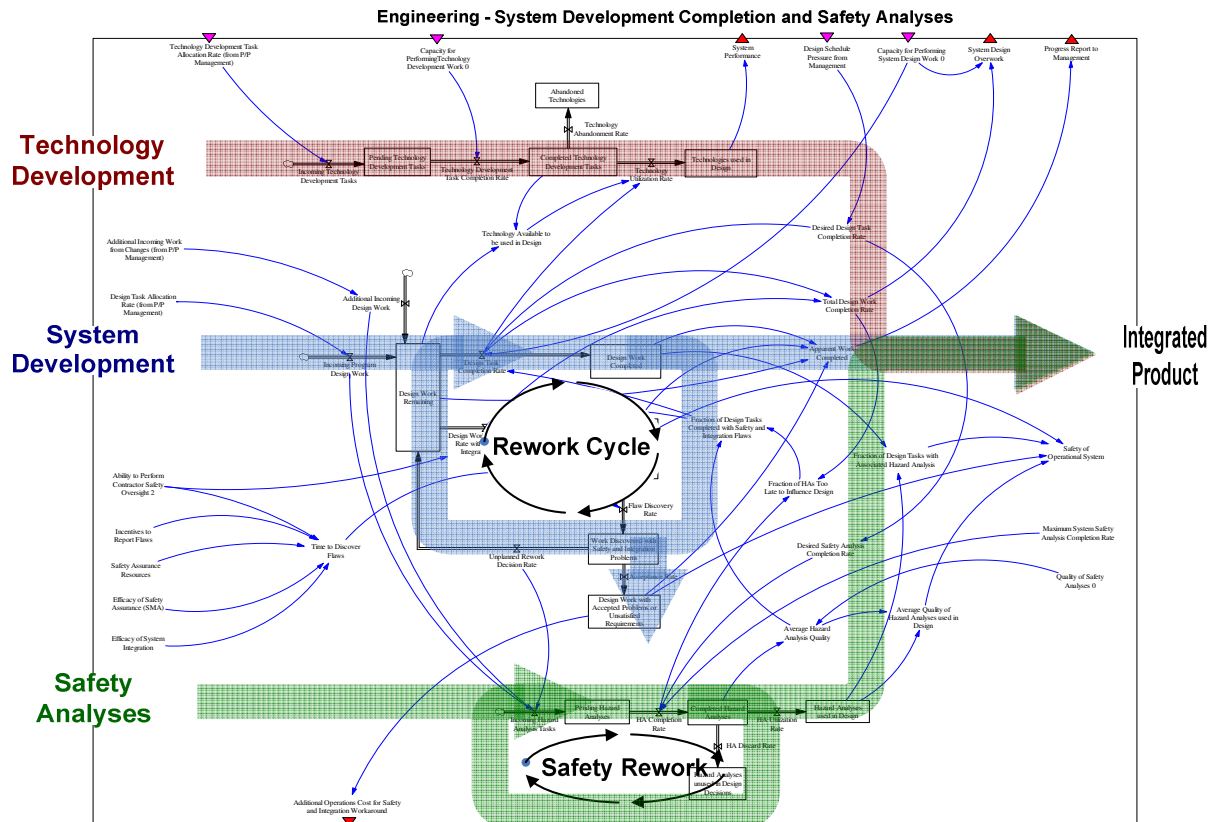


Figure 100: Schematic of task completion flows

6.2.5 ENGINEERING – TECHNICAL PERSONNEL RESOURCES AND EXPERIENCE COMPONENT

The purpose of the Engineering (Technical Personnel Resources and Experience) component is to keep track of the human resources working on ESMD projects. This component was initialized and calibrated using NASA employment data available online [NASA, 2006]. The objective is to monitor the availability and characteristics of the technical workforce responsible for the development of the exploration system. The component considers the number of people hired for entry-level positions and for experienced positions, as well as transfers between ESMD and other NASA directorates such as the Space Operations Mission Directorates (SOMD). It also keeps track of the experience of NASA technical employees as well as attrition rates, retirements, early retirements, and potential transfers of NASA employees to private contractors.

6.2.6 ENGINEERING – EFFORT AND EFFICACY OF TECHNICAL PERSONNEL COMPONENT

The purpose of the Engineering (Effort and Efficacy of Technical Personnel) component is simply to collect information from various sources in the model and output the total capacity of the in-house workforce to perform development work in areas of technology development, system integration, and system development.

6.2.7 SAFETY AND MISSION ASSURANCE - EFFORT AND EFFICACY OF SYSTEM SAFETY ANALYSTS (EESSA) COMPONENT

The focus of the Safety and Mission Assurance component is on the effort and efficacy of in-house employees working on safety analyses. The purpose of the component is to determine the capacity of safety analysts to work hand-in-hand with other engineers and technical people in order to produce high-quality, useful safety information to be used in making design decisions.

6.3 INTEGRATED ESMD MODEL OVERVIEW

The final model is made up of a large number (many hundreds) of feedback loops that cut across individual model components. Some of those feedback loops are generic and were introduced in chapter 2 (see Figure 25 for an example). For example, Figure 101 shows a schematic explaining how the management pressure feedback loops influence system development as progress reports are communicated upwards through feedback channels, and multiple layers of management and administration attempt to ensure that the project meet requirements while remaining within acceptable budget and schedule.

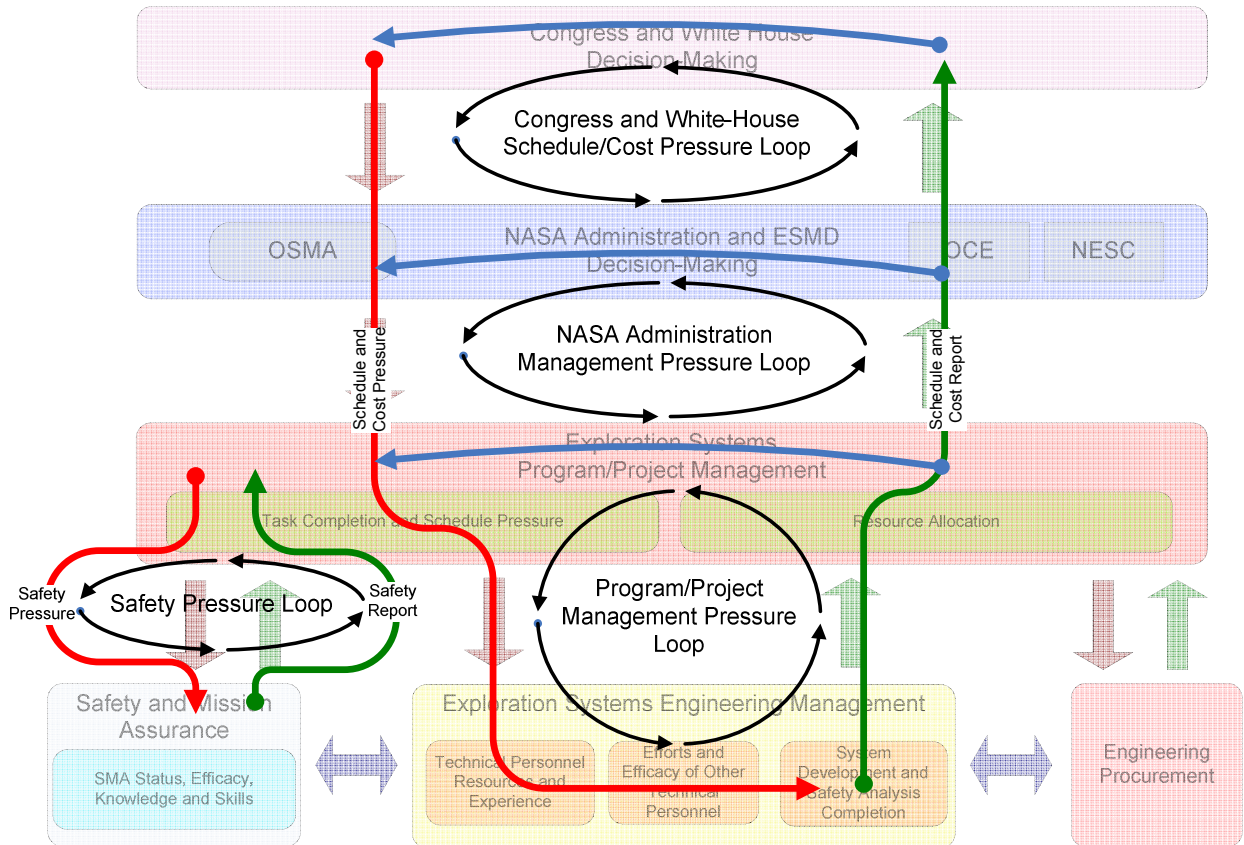


Figure 101: Management pressure feedback loops reproduced through component interactions

As another example, Figure 102 shows a schematic explaining how some of the previously introduced feedback loops are reproduced across components and impact system safety. The management pressure loops of Figure 101 are very important to ensure that the system is developed without catastrophic disruptions in cost and schedule. However, as Figure 102 shows, excessive management pressures have negative side effects that may impact system safety. As too much pressure is applied toward getting development done on time, safety becomes an effectively lesser priority. Resources may be re-allocated from safety analysis toward design completion, budget cuts may be directed toward safety or safety analyses that can delay design completion may be ignored. All of these factors end up negatively impacting the final safety of the system. One of the interviewees nicely summarized the interaction of schedule pressure and design flaws in this comment:

“Schedule pressure is not a bad thing if it’s applied right. Schedule pressure is necessary, so it’s a positive thing too. People don’t produce as well without schedule pressure. It’s a matter of when the schedule pressure goes over the edge, and your fraction of tasks with flaws goes up too high. It’s almost like an exponential curve, for a long time, the effect is not too bad, but when schedule pressure is too high, people just give, and they say, whatever you want, you got, you want that thing out the door, you got it. Productivity increases with schedule pressure, but flaws increase too.”

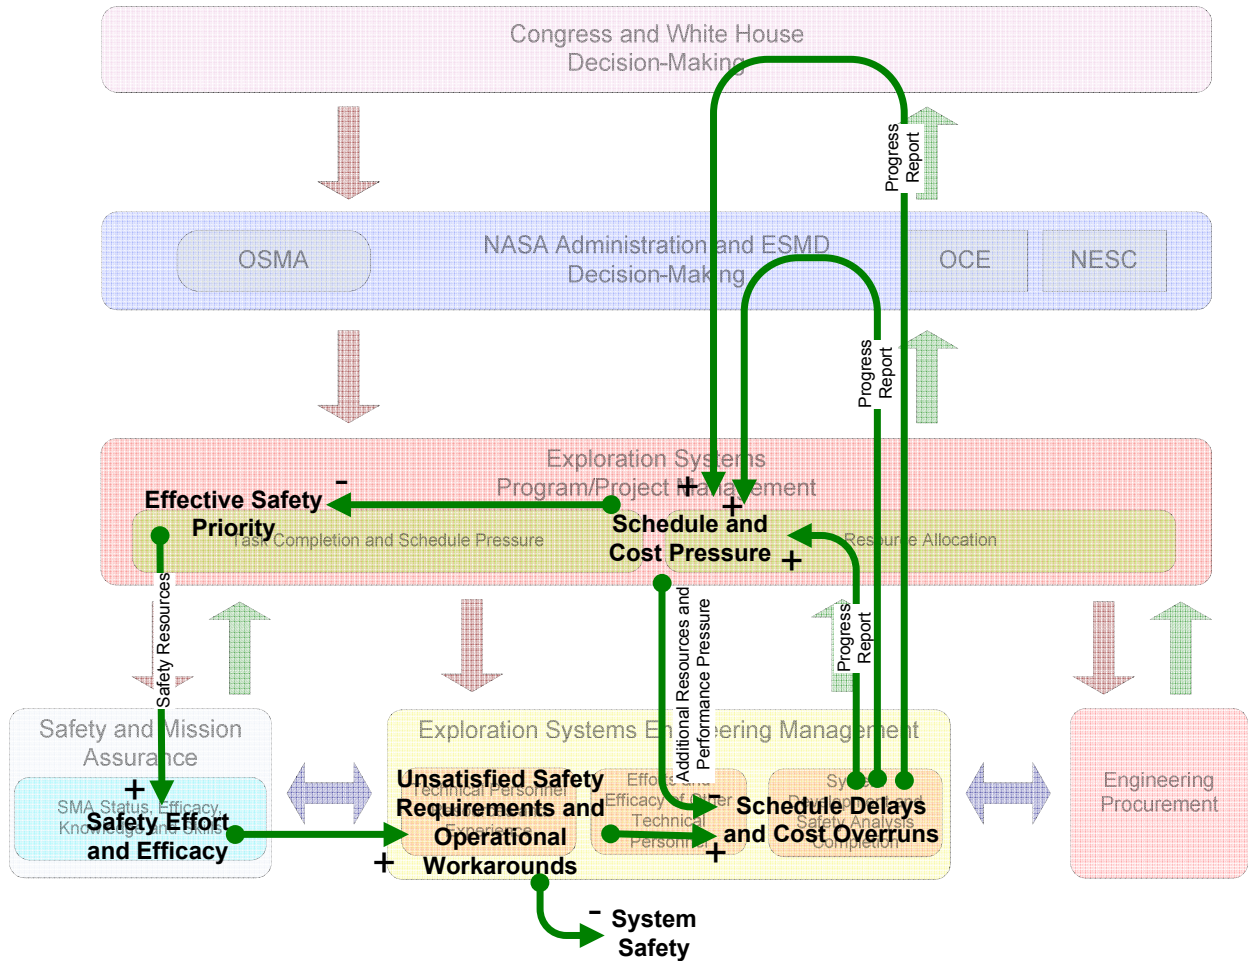


Figure 102: Interplay of schedule pressure and system safety

Another interviewee explained how effective safety priority may simply change over time:

“That (Management Safety Priority) variable does change over time. For the folks working system safety, the priority does not change but as an overall agency focus, it changes. Safety is always priority one, but then again when you

start competing for resources, when you take resources away, by definition it looks like you're lowering the (safety) priority.”

And as a follow-up, from another interviewee:

“...if you get the sense that management only pays lip service to safety, then it flows downhill, and it may be the next day that somebody cuts a corner in a document, or a design, and they figure it's good enough and then it bites you. And management must continually reinforce this. It starts from the top, and if you don't have that, it won't happen from the bottom up.”

The loops summarized in this section were a small sample of the entire feedback structure of the model. As mentioned previously, complete documentation for the integrated model is provided in Appendix K.

6.4 MODEL RESULTS AND SCENARIO ANALYSIS

In many cases, the most important result from the use of a STAMP-based dynamic model is a deeper qualitative understanding of the overall system-level response to various scenarios and situations. The actual quantitative values generated in the simulations are sometimes of secondary importance in comparison to the qualitative learning opportunities presented by the model and the modeling process. For instance, the model can be used to identify unintuitive decision rules that outperform others in terms of cost, schedule, final system scope, and/or safety; to test policies that may yield desirable results at first yet be harmful over time; or to identify metrics that provide early warnings of undesirable trends.

Several model variables are meant to serve as relative performance indicators. One example in the ESMD Model is the variable Safety of the Operational System. The numeric value of this variable is not linearly related to the likelihood of accidents. Instead, it provides a reference variable for the comparison of multiple simulation runs (i.e., if the final value of “Safety of the Operational System” is higher at the end of one simulation than it was in another, then the product of the design process from that run is better in terms of the attributes that have been explicitly identified as important to the safety of the operational system). Such reference variables allow users to evaluate the relative importance of parameters and decision rules in the model.

Furthermore, the models are useful in the identification of tipping points or unstable equilibria in the system. Tipping points are points beyond which the stabilizing forces of the balancing loops in the system are overwhelmed by the destabilizing forces of the reinforcing loops of the system. Potential for tipping point dynamics may exist in a system but remain dormant until a specific combination of factors (or system states) activate the (often undesirable) tipping point behavior. Since the model equations are usually nonlinear, linear systems theory cannot be used to uncover potential tipping points and instabilities. Monte-Carlo simulations are usually the only way to identify tipping points. The model can then be used to design and evaluate policies to avoid or exploit tipping point dynamics.

In summary, the model is instructive in answering questions generally associated with risk management affecting the technical and organizational aspects of the system. These questions include, but are not limited to the following:

- Which actions will be relevant in addressing the issues at hand?
- Will decision-makers have the appropriate feedback to make sound decisions?
- Will this policy (control action) have an immediate or delayed effect?
- Will this policy (control action) have the intended effect, and if so, how long will this effect last?
- Will this policy (control action) have unintended, undesirable effects?
- What indications, if any, will be available if the policy (control action) does not work out?
- Where are the bottlenecks in the system that might or will prevent the intended effect of the policy (control action)?
- Over time, how will one policy perform relative to another?

The following sections provide examples of model results from risk analysis scenarios run on the model based on inputs from NASA experts. The risk scenarios investigated are: 1) Workforce Planning, 2) Management Reserves, 3) Schedule Pressure and Safety Priority, 4) Safety Influence, Power, and Leadership, and 5) Changes in Requirements and System Scope.

6.4.1 SCENARIO 1: WORKFORCE PLANNING

6.4.1.1 Scenario Motivation

The workforce planning scenario was inspired by concerns raised in a number of interviews about the current and future skills and capability of the workforce. The following quote from a recent report by the National Academies of Science Space Studies Board [NAOS, 2006] summarizes many of the long-term issues raised in the interviews:

“NASA is not currently experiencing a supply problem in terms of overall available personnel. But the agency is experiencing a more complex and subtle problem that will grow over time. Like other government agencies and aerospace contractors, NASA is experiencing difficulty finding experienced personnel in certain areas, such as systems engineers and project managers. NASA’s workforce also has a skewed age distribution arising from hiring policies first implemented in the 1990s. The agency did not experience a hiring freeze during that time, but it adopted policies whereby it filled specific positions but did not hire younger people and ‘grow’ them into positions. As a result the agency’s mean age has continued to rise over time, and it lacks younger employees with necessary skills. As the agency embarks on new human and robotic exploration programs, problems in fulfilling demand will likely increase because the agency has not been developing the necessary employees from within.” [NAOS, 2006]

Additionally, NASA is struggling with a short-term workforce problem referred to as “unfunded capacity”. This problem stems from programmatic changes across the agency in response to the Vision for Space Exploration (VSE). Whenever a NASA project or program ends or gets restructured (either through the planned conclusion or cancellation of the project/program) the employees working on that project/program must be reallocated to another project. As the newest NASA programs are almost exclusively oriented toward space exploration, many employees hired before the VSE cannot find a project related to their area of expertise. Additionally, employees sometimes do not transfer to new projects for which they are qualified because the projects are located at different centers. Consequently, there is bound to be a fraction of the NASA civil servant workforce not assigned to an Agency project/program at any given time. This contingent is referred to as the “unfunded capacity”. In the wake of programmatic changes made for the VSE, there are currently about 900 unfunded civil servant positions at NASA. With the impending retirement of the Space Shuttle in 2010, and workforce retirements due to the average age of Agency civil servants,

NASA may face a challenge in maintaining enough in-house skills to achieve its exploration objectives.

6.4.1.2 Scenario Description and Results

A simulation scenario was developed to explore workforce needs at ESMD between 2004 and 2016. The results are shown in Figure 103 and Figure 104. Figure 103 shows the results of an analysis of the ESMD Employee Gap for a case where transfers are accepted from the Space Shuttle Program and the unfunded capacity (see Figure 103A) and a case where transfers are not accepted from other directorates in the Agency (see Figure 103B). The contours in the figure represent different levels of hiring and transfer from the Space Shuttle Program (the blue/dark contour represents the largest monthly hiring rate of Science and Engineering personnel for the entire Agency since 1993 [NASA, 2006]). In both of these cases, the ESMD Employee Gap (the difference between the required and actual number of experienced ESMD civil servants) increases dramatically shortly before the Space Shuttle is retired and then tapers off towards the end of the simulation. This tapering occurs because funds from the Space Shuttle Program in the Space Operations Mission Directorate (SOMD) are transferred to ESMD, thus creating a need for ESMD to hire new employees and/or transfer employees in from unfunded capacity (towards the end of the simulation, the *ESMD Employee Gap* decreases mainly due to the fact that a portion of its budget gets transferred back to SOMD when the CEV/CLV is deployed for operations). Unfortunately, transferring civil servants in from SOMD is only partially effective because the Space Shuttle Program is comprised of a civil-servant to support-contractor ratio that is much lower than the Agency average. Figure 104 shows that when the ESMD Employee Gap is filled through the hiring of support contractors through a simple proportional control law (i.e. the bigger the gap, the more support contractors hired instead of civil servants) the ratio of civil-servants to support-contractors in ESMD trends down towards the ratio of civil-servants to support-contractors of the space shuttle program.




Figure 103: Increasing demand for ESMD technical civil servants (Fixed civil servant to contractor ratio)




Figure 104: Ratio of productive civil servants to support contractors

A conclusion from this scenario is that if ESMD inherits the budget and workforce of the space shuttle program it will also inherit its civil servant-to-support contractor ratio if hiring rates are not increased. Consequently, a follow-up question that must be raised based on this analysis is: *“Will a low civil servant-to-support contractor ratio in a development environment (ESMD) be equally effective as a low ratio in an operations environment (Space*

Shuttle Program)?” Further analysis will be required to answer this question, but it is very conceivable that a higher number of in-house civil servants will be necessary to oversee and integrate the work done by support and procurement contractors, especially since contractors working for one private company cannot readily oversee workers from another.

6.4.1.3 Recommendations

Based on the scenario analysis, and in the light of the “unfunded capacity” problem, it appears absolutely necessary to further investigate the impact of governmental hiring freezes on the future capability of the civil servant workforce. According to analysis results, hiring rates would have to be raised significantly to offset mass retirements and loss of knowledge and development capacity. In fact, in order to retain the current capacity, hiring rates should increase to a level unseen in 15 years at NASA. Additionally, more investigation is necessary to better understand the impact of the high contracting ratio predicted in the analysis on a new and complex development organization such as ESMD.

6.4.2 SCENARIO 2: INVESTIGATING MANAGEMENT RESERVES

6.4.2.1 Scenario Motivation

Many interviewees voiced the concern that system requirements may not be reasonably aligned with available resources. Another concern was that program management leadership may not be able to secure and manage sufficient reserves to mitigate uncertainty and complete system development on schedule. While discussing the program/project management model component, an interviewee discussed the interplay between program management, requirements and resources:

“In general, the higher up you go, the less they want you to take (requirements) out. Generally the pressure to cut cost is to cut cost without changing content [...]. That’s not always the case, clearly there (are) exceptions, but I would say in about 90% of the cases it’s cut cost, don’t cut content. And that goes back to that leadership thing where you’ve got to be willing to stand up there and say: ‘(We) can’t do it. You want it to cost X, I’ve got to drop Y or if you want us to do X, then it’s got to cost whatever we said.’ [sic] [...] At PDR (Preliminary Design Review) you go off and you look at all of those requirements and you say: ‘The requirements don’t stack up to the dollars that we signed up to before.’ Not surprisingly when do you replace project managers? Right around

those times! Now is it really the project manager's fault? [...] you knew going out that you gave them a task that they couldn't do for the dollars that they had, that goes back to the top part of your chart here (pointing at the PM component)."

6.4.2.2 Scenario Description and Results

A scenario was created to investigate the impact of system requirements planning and management reserves on development completion and the overall safety of the system (as defined by the relative system safety indicator). In the baseline scenario, the requirements are calibrated based on a planned system development time of 8 years (96 months) in order to hit the 2012 CEV launch deadline. The resources available are calibrated according to planned ESMD budgets and workforce data [NAOS, 2006].

A sensitivity analysis was performed based on these baseline parameters. The system requirements were randomly varied +/- 25% of the baseline value of 8 years worth of requirements; that is, requirements for 6 to 10 years of development. Given that the planned development deadline is fixed at 2012, the 6 years of requirements run is associated with more reserve as the system could be theoretically completed in 2010, and the 10 years of requirements is overly optimistic, since, given baseline resources and realistic pressure, the system could not be completed until 2014.

In a similar manner, the resources (budget and workforce) available are varied +/- 20% of the baseline value. 120% of the baseline value is the conservative management case with 20% management reserve, while 80% is the overoptimistic planning case with insufficient baseline resources to perform the work on time (the baseline does not include explicit reserves).

Figure 105 and Figure 106 provide some scenario results for the baseline and envelope case, namely the overoptimistic planning case (10 years requirements, 80% resources), and the sufficient reserves case (6 years requirements, 120% resources). As can be observed, the management reserves have a significant impact on completion time and relative system safety (the arrow in Figure 105 indicates increasing management reserves). In the worst case scenario, there is more work to do, and fewer resources to do it, resulting in longer completion time. Lower safety can be explained through the side effect loops shown in Figure 102. As there is more to do with fewer resources, managers apply additional schedule pressure,

resulting in lower safety priority, lower quality safety analyses, more unsatisfied safety requirements, and consequently lower overall system safety of the final system.




Figure 105: Fraction of work remaining for the baseline and envelope scenarios




Figure 106: Relative safety of the final system for the baseline and envelope scenarios

In addition to envelope results, the scenario sensitivity analysis allows the generation of outcome distributions. Using a random variation of the parameters documented above, the final system development outcomes were collected and arranged by frequency interval to obtain the distributions of Figure 107 for completion time (in years) and (final) system safety.

The advantages of this approach are multiple. Among others, the best (or worse) final outcomes can be traced back to the specific parameters that produced them, and the runs can be analyzed individually. In this scenario, the runs at the tail of the distribution correspond to envelope values of Figure 105 and Figure 106, but in other cases, extreme outcomes may be associated with surprising model parameters.




Figure 107: Outcome distributions (completion time and safety) for the sensitivity analysis

6.4.2.3 Recommendations

Based on this scenario, it appears that careful planning and reserve utilization, management, and monitoring are critical to dampen disturbances in workforce, budget, and technology availability. One of the indicators that planning or reserves may be inadequate is the workload of the ESMD workforce. Especially critical employment areas include system engineering and integration, safety engineering, and safety assurance, as these areas control the strength of the rework cycle, and as an interviewee mentioned: “*Doing things wrong the first time costs you more*”. Based on the scenario results, one could add that you also get a lower quality, less safe system in the end.

6.4.3 SCENARIO 3: EFFECT OF SCHEDULE PRESSURE AND SAFETY PRIORITY

6.4.3.1 Scenario Motivation

Schedule (and budget) pressure ended up being one of the most common themes discussed by interviewees. The Columbia Accident Investigation Board (CAIB) made it clear that the managers and engineers of the shuttle program were under tremendous pressure from the

NASA administration to meet the February 2004 deadline for the Space Station to reach “core complete” configuration [Gehman, 2003]. The CAIB recommended the implementation of measures (such as the ITA) to ensure that schedule pressure does not negatively affect safety-related decisions. Nevertheless, some of these recommendations (including the ITA as it was initially designed) were discarded for the development of the exploration system, only three years after the Columbia accident. As interviewees put it:

“Schedule is a major risk factor; [...] what we’re trying to do in exploration is pretty aggressive so it’s going to make it hard to not get caught in some of the same kind of [...] I mean everybody is really sensitized to not caving in to budget and schedule pressures (everybody knows) what happened on Columbia and Challenger, all of these things. Having a way to keep that from happening again, I think it’s going to be an issue.”

and, similarly:

“Then I think the other thing (primarily affecting safety risk in the system development) is having reasonable schedules and manpower and the right skills. The schedules now are pretty ridiculous to actually do the things that need to get done.”

6.4.3.2 Scenario Description and Results

A scenario was developed to investigate the impact of schedule pressure and enforcement in the exploration development program. As mentioned previously (see chapter 2), management pressures were implemented in the model as a simple PID-type controller. In a nutshell, a profile for the desired fraction of completed development was created based on actual and forecasted yearly budget allocations [NASA, 2004]. The schedule pressure applied at the program management and administration level is a function of the difference between the measured work completed and the desired work completed at any point in time. This simple controller framework is applied to the desired system development completion profile, as well as the desired safety analyses completion profile.

In the present scenario, the proportional gain responsible for the application of pressure at the program management level (when development falls behind schedule) is varied from a value of 0 to 10. Consequently, the pressure applied is simply equal to the gap in schedule completion times a proportional gain (K). The same variation (0 to 10) applies to the safety

pressure gain, that is, the pressure used to ensure that safety analyses are performed early enough to be used in design decisions

Figure 108 shows the estimated project outcomes for safety, schedule, and cost as a function of extreme values (0:Low, 10:High) of schedule pressure and safety priority gains. As can be observed, overly aggressive schedule enforcement has little effect on completion time (<2%) and cost, but has a large negative impact on safety. Inversely, priority of safety activities has a large positive impact, including a positive cost impact as less rework is required because high-quality safety analyses were used to influence design decisions in the first place.




Figure 108: Outcome (Safety, Schedule, Cost) as a function of schedule and safety priority (low, high)

Figure 109 shows the estimated cost result of a more continuous variation of the schedule and safety gains from 0 to 10. The improvement in cost observed when the schedule gain is low and the safety gain is high is achieved because less rework (associated with variable and fixed costs) is necessary as the safety work was done correctly and on time. The improvement in cost associated with high schedule pressure and low safety priority is achieved at the detriment of safety, which means development is finished earlier (lower fixed costs) but the final system is unsafe.

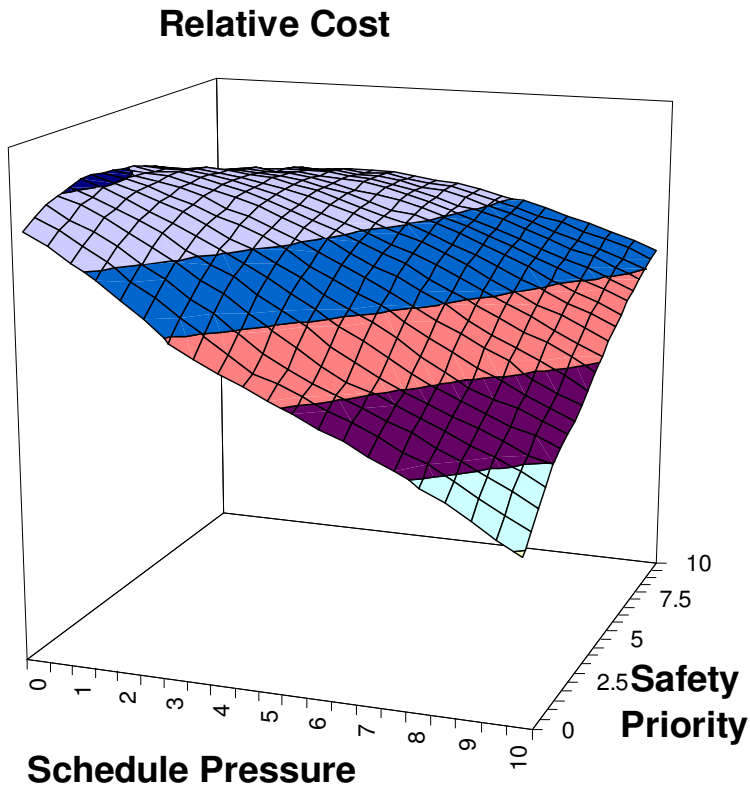


Figure 109: Estimated relative cost as a continuous function of schedule pressure and safety priority

6.4.3.3 Recommendations

Recommendations from this scenario include the monitoring of workforce workload, as extreme workload and employee burnout allow for more mistakes which necessitate more rework. In addition, ensuring that safety analyses are used in design decisions is a good way to verify the synchronization of the design and safety flows. Controlling safety requirements waivers and operational workarounds is another way of ensuring that schedule pressure does not take undue priority over safety concerns.

6.4.4 SCENARIO 4: EFFECT OF HIGH SAFETY INFLUENCE AND POWER ON DECISION-MAKING

6.4.4.1 Scenario Motivation

Many interviewees very well summarized some of the problems safety and mission assurance (SMA) people have in influencing decisions and being a powerful voice in the design process from the early lifecycle phases. Examples from interviewees include:

Interviewee 1:

“Historically, the Agency has leaned toward engineering and not respected SMA much. SMA was not been asked to do very much so it has not had the respect and staffing it needs. We kid around as SMA being a dumping program for people. Well if your products are irrelevant, the people that produce them can get away with being irrelevant. Some of the best people in the Agency are in SMA, there's just not enough of them. We have a real problem with technical competence in this area and that's how you get the respect, where you bring the value to the table.”

Interviewee 2:

“The model in SMA was, you get your job, then you go get a contractor who does the work and you just report the results. The reputation they've earned is that they come in and throw the problems around and walk away -- you don't try to help solve the problem. Who would want to go into an organization like that?”

Interviewee 3:

“I think in the old model that there was no credibility, that in the old model, when it was (a situation) we come in late, we see a problem, we tell you it's a problem and that you need to fix it and we disappear and (tell you to) call us when you do. I think that left very little credibility. There are those that will tell you that our organization has been fighting for credibility for a very long time and some of that is the police mentality, some of it is the feeling that our folks don't even know our own business, much less the space business -- and actually we've ended up in the past being viewed as (a place) when all else fails with Jim Smith over here, we can't get him to do good stuff where he is, let's just send him to S&MA, so we would end up with problem children who weren't particularly motivated in our process and the way we operated and that didn't do much for our credibility within the project.”

6.4.4.2 Scenario Description and Results

A scenario was created to investigate the impact of safety influence and power on decision-making. In the system safety component of the model, the variables *Power and Authority of System Safety Analysts*, as well as *Status and Credibility of System Safety Analysts* were varied +/- 50% about baseline values. The results shown in Figure 110 illustrate the potential impact of safety influence and power on system development. High safety influence and power has a large positive impact on safety (see bottom-left of Figure 110) because it

effectively removes the “relief valve” of accepting design problems, unsatisfied requirements, and operational workarounds (see bottom right of Figure 110). On the other hand, not accepting these problems may necessitate additional rework, which has slight negative impact on cost and schedule (see top of Figure 110). The negative impact on schedule and cost can be dampened by allocating more resources to system integration, and by carefully planning and anticipating safety analysis requirements. Indicators of safety influence and power on decision-making were identified in the model and include: 1) Safety-based design changes, 2) Overruling of safety decisions, 3) Adequacy & stability of safety resources, 4) Review time allocated to safety analyses, and 5) Unsatisfied safety requirements.

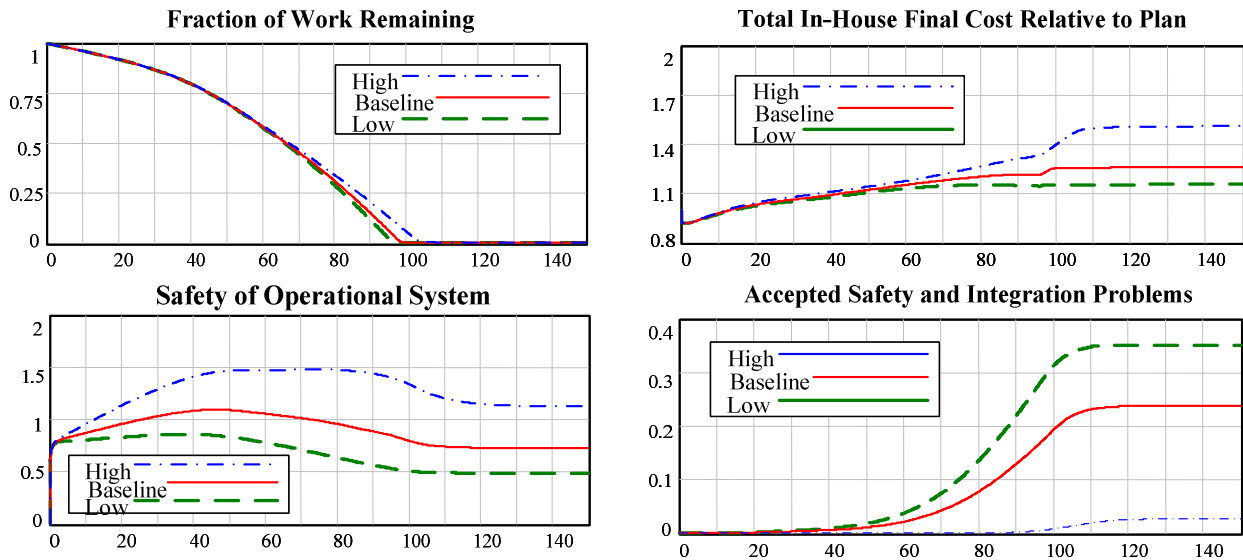


Figure 110: Impact of safety influence and power on project dynamics

In a similar way, the leadership and technical expertise of safety and mission assurance (SMA) personnel was a recurring theme during interviews. Leadership and technical expertise is closely related to safety influence and power on decision-making. Consequently, a second related scenario was created to investigate the impact of the assignment of technical leaders to safety analyses. The results (see Figure 111) show that assigning high-level technical leaders to safety analyses has a very high impact on safety, with a minimal impact on cost and schedule. One of the main reasons for this positive impact is that high-level technical leaders will have the capability to deliver high-quality safety analyses on time, and

the status and credibility necessary to influence design decisions (see right side of Figure 111). The effective result is a high quality product with safety “designed-in” that meets schedule and budget because of less rework. Indicators of the assignment of technical leaders to safety were identified in the model and include: 1) Attractiveness of safety positions, 2) Experience and skills of current and incoming workforce, and 3) Impact of safety analyses on design.

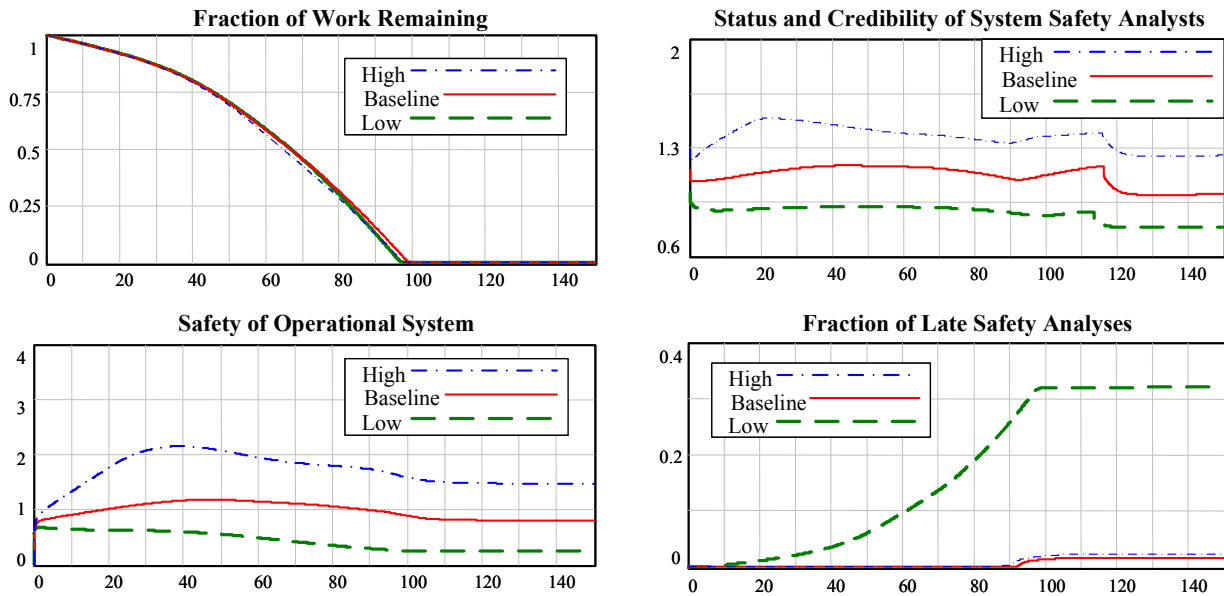


Figure 111: Impact of high-level technical leaders on project dynamics

6.4.4.3 Recommendations

Based on the analysis results of this scenario, ensuring that system safety analysts have high influence and power on decision-making, as well as high credibility and status is a very effective way to improve the safety of the system with minimal impact on cost and schedule. Allocating high priority to the assignment of high-level technical leaders to safety analyses (possibly through the rotation of technical leaders and “rising stars” into safety positions) can be an effortless way to positively impact the status, credibility, and impact of safety analysts.

6.4.5 SCENARIO 5: EFFECT OF CHANGES IN REQUIREMENTS AND SYSTEM SCOPE

6.4.5.1 Scenario Motivation

Having a fixed set of clear requirements from the beginning is arguably the most effective way to ensure the system will be developed on time and on schedule. As an interviewee puts it:

“You have to have a good set of requirements that can’t be over-burdensome, if they are, then safety is going to be compromised because you’ve provided too much detail at too high of a level. You could (also) provide too little detail at too high of a level and you’ve got the same problem. So you need clear requirements to get there.”

From a project management perspective, frozen requirements are a huge advantage. As a manager mentioned:

“I’m a longtime project guy so my bias is toward: ‘Give me my requirements and go away. You know once I have the requirements, I will deliver.’”

However, for a very large development environment such as the exploration system, freezing requirements early is very difficult because of the complexity of the system and the low maturity of the technology. This creates large difficulties and opens the door to changes in requirements and system scope later in the development cycle. As one interviewee puts it:

“Well one thing is - if you don’t know this already - at NASA, no decision is ever final. [...] we constantly strive to have the best technical thing that we can have, so that’s why no decision is ever final and then you end up impacting cost and schedule.”

6.4.5.2 Scenario Description and Results

A scenario was created to investigate the impact of requirements and scope change on project dynamics and system characteristics. In this scenario, the baseline simulation corresponds to requirements that are well-defined and frozen from the beginning of system development. A second simulation was run where very small requirements changes (<2% of the total requirements) are made at a 12-month regular interval. Finally, a third simulation was run where large requirements and scope changes (<20% of the total requirements) are made at a 60-month regular interval (see top-left of Figure 112). Upward changes in scope and

requirements (in top-left of Figure 112) indicate added requirements, while downward changes indicate abandoned requirements. Consequently, an entire cycle of downward/upward changes is associated with a replacement of requirements even if the total net change in the number of requirements and tasks is zero.

It is well understood and accepted that large requirements and scope changes will have significant negative consequences on project cost, schedule, and system characteristics. Not surprisingly, the scenario reproduces these expected results (see Figure 112) as large changes have a disastrous impact on schedule and system safety. Also, not surprisingly, the later the changes, the more negative impact they have. Another not so intuitive result, however, is that small but more frequent changes can have a similar negative impact on the system.

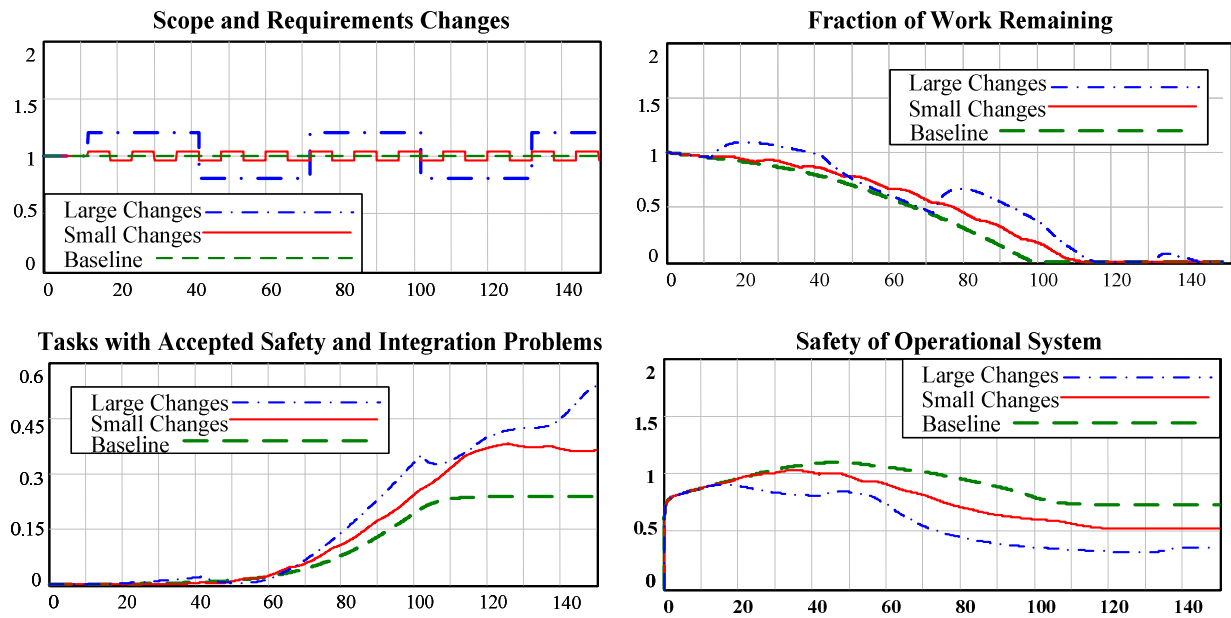


Figure 112: Effect of scope and requirements change on project dynamics and outcome

Detailed planning is necessary to limit requirement changes (even small ones) and their negative impact on system development. Many potential mitigation strategies were identified to lessen the impact of some amount of (inevitable) changes later in the development cycle. These mitigation strategies include: 1) the use of additional management reserves, 2) design and planning for operations, 3) on/off-ramps for technologies and design, and 4) improved system engineering and integration.

6.4.5.3 Recommendations

In a perfect world, requirements would be well-defined and frozen at the very beginning of a development project. However, this is rarely the case, in particular for new and complex engineering systems using unproven technologies that are developed and tested as the system is designed. The scenario presented in this section demonstrated that small, frequent changes can have a negative impact similar to larger, less frequent changes. Careful and complete requirements specification, combined with the mitigation strategies mentioned above can help lessen the negative impact of changes. The following metrics were identified to monitor the amount of smaller, more subtle requirements changes: 1) the number of discarded or otherwise performed but unused design tasks, safety analyses, and technology items, and 2) the number of accepted unsatisfied requirements and operational workarounds.

6.5 CASE STUDY SUMMARY

In this chapter, the model-building and analysis methodology was used in combination with a proposed interview-based research protocol to create a risk management model of the planned NASA space exploration system. The ESMD case study demonstrated that:

- 1) The model-building methodology introduced in this thesis facilitates the building and customization of dynamic risk management models to a point where tools could be implemented and used by domain experts with limited simulation knowledge.
- 2) The STAMP-based component approach combined with a well-designed research protocol focused on the expertise of interviewees facilitates model validation and confidence-building activities.
- 3) Augmenting STAMP with a dynamic simulation framework improves the ability to mitigate time-dependent risks such as those identified and addressed in the scenarios above.
- 4) The models can be used to create and test policies and control actions to improve risk mitigation in the development of a complex engineering system.

- 5) The models can be used to find indicators of risk increase for specifically identified and analyzed risks, as well as for general system migration toward states of higher risk.

CHAPTER 7: CONTRIBUTIONS, FUTURE WORK AND CONCLUSION

7.1 SUMMARY

STAMP [Leveson, 2004] improves over traditional accident models by using a control theory framework that allows the handling of system accidents and software-intensive systems, as well as complex decision-making influenced by managerial, organizational and social factors. In this thesis, the STAMP model is augmented with a new framework to create dynamic executable models used to analyze time-dependent risks, assist engineers and managers in safety-related decision-making, create and test risk mitigation actions and policies, and monitor the system for states of increasing risk. The usefulness of the STAMP-based model creation and analysis methodology was demonstrated using a risk analysis of the NASA Independent Technical Authority (ITA), an operation-centric system, and of the NASA Exploration Systems Mission Directorate, a development-centric system. The NASA models were calibrated and validated through extensive data gathering and interviews with domain experts at five NASA centers.

7.2 CONTRIBUTIONS

This section provides a short summary of the original contributions made in this thesis.

7.2.1 STATIC STAMP CONTROL STRUCTURES CREATION AND ANALYSIS GUIDELINES

A technique was introduced to support the creation and analysis of STAMP static safety control structures. The technique is based on the use of generic static connectors that define relationships between components of the socio-technical system. Example connectors include: direct report (authority), direct oversight, performance appraisals, resource allocation, and personnel appointment. The use of generic relationships to create safety control structures allows automated analysis of the static structure to be performed. Completeness and consistency criteria for the analysis of static control structures were proposed. Completeness criteria provide guidelines to ensure that the static control structures are well

defined. Consistency criteria provide warnings and raise flags when the combination of generic connectors used in the structure may lead to unsafe and dysfunctional interactions between components. Preliminary testing was performed on the ITA system that allowed the identification of potential problems in the ITA reporting structure and responsibility allocation

7.2.2 MODEL-BUILDING METHODOLOGY

The main contribution in this thesis is a methodology to facilitate the creation and validation of dynamic executable safety-risk models based on the STAMP accident model. The structure of the dynamic models created using this methodology mirrors that of the static STAMP safety control structure, which greatly facilitate model building and analysis. The methodology is based on the customization of generic dynamic components created using accident reports, relevant system safety and organization theory literature, and interviews with domain experts and system stakeholders in the organizations studied.

The component-based approach has multiple advantages. The initial components provide a structured basis for discussion and greatly facilitate the creation of final customized components. The approach allows interviewees to focus on their area of expertise and facilitates the refinement and piecewise-testing of model components. During the ESMD project, the complexity of models was reduced to a point where significant validation progress could be made in a single-hour interview. In some cases, insights were collected and scenarios could be created and run at the component level to focus on a single issue of concern. In most cases, interview participants felt confident enough to take charge of the component discussion, scratching out variables and adding others directly on the model printouts.

Instead of starting from a clean sheet, the generic components combined with the STAMP control structure offer useful guidelines, making model creation available and usable to managers, scientists and engineers with limited knowledge of simulation models. The creation of software packages including a repository of generic components, as well as model assembly and analysis features can further improve the accessibility and usability of the models.

The ESMD case study presented in this thesis illustrated the usefulness of the model-building methodology to facilitate the creation and validation of complex risk analysis models.

7.2.3 MODEL-BASED RISK ANALYSIS TECHNIQUES

Many techniques were presented to facilitate model-based risk analysis. Some of these techniques were originally created for use in other fields, but were customized in this thesis and applied to the type of safety risk analysis performed using the STAMP framework. The model analysis techniques include:

- 1) Preliminary tool designs for the visualization and understanding of model structure and behavior. The tools can be used over time to improve the mental models of decision-makers through interactive scenario-based learning, but more importantly, they can be used immediately to assist in daily safety decision-making.
- 2) Scenario-based risk analysis can be performed on risks identified by system stakeholders, or in a more comprehensive and structured way, using the STAMP risk identification process. The scenario analyses usually include Monte-Carlo type sensitivity analyses to identify interesting behavior patterns such as tipping points and to obtain model outcome distributions.
- 3) The creation and testing of policies and control actions to mitigate time-dependent risk increases.
- 4) The identification of early indicators of risk increase to better monitor safety erosion and states of risk increasing risk.

7.2.4 INTERVIEW-BASED MODEL VALIDATION PROTOCOL

The interview-based research protocol developed for the ESMD project was very effective in ensuring the accuracy of the STAMP safety control structure and facilitated model creation and validation. The protocol can be reused and customized for future risk analysis projects.

7.2.5 PROJECT-SPECIFIC INSIGHTS

The modeling projects performed for the NASA ITA and ESMD generated insights that may help mitigate risks in real complex systems. For example, the hiring limitations at NASA may have to be revisited in order to maintain in-house knowledge and capabilities. The projects also identified areas of concerns where future research efforts should be focused. These areas of concern include the impact of contracting ratios on complex system development and the dynamic patterns of requirements waiver and operational workarounds accumulation.

7.3 FUTURE WORK

This section discusses areas of future research that were identified as having high potential to further improve risk management in complex systems.

7.3.1 MODEL-BUILDING AND ANALYSIS SOFTWARE TOOLS

The methodology introduced in this thesis greatly facilitated the creation of the ESMD model. However, model creation would have been much more difficult without extensive experience and knowledge of the non-specific simulation package used for model-building. As effective as the methodology may be, it will not be extensively used in real systems until a commercial quality software package is available to automate the mundane tasks required to assemble and analyze the models. The software tools should include a repository of generic components that can be easily augmented as new components are developed. The tools should allow for easy customization of the components without extensive simulation knowledge. The components should be easily connected through generic interfaces and connectors as defined in chapter 4. The entire model building and testing criteria introduced in this thesis should be integrated in the software package.

The model analysis techniques introduced in chapter 5 should also be integrated into the software package. The tool should facilitate the creation and documentation of risk analysis scenarios to help managers and engineers in making informed safety-related decisions.

7.3.2 MODEL VISUALIZATION TOOLS

Prototype model structure and behavior visualization tools were developed based on the ITA model [Friedenthal, 2006]. Further research should focus on the evaluation of design principles for the creation of effective interactive visualization. Interactive visualization tools are necessary to improve the domain expert's understanding of the model and to ensure a good correlation between the model and the real system.

7.3.3 VALIDATION IN OTHER DOMAINS

The generic components were created mostly based on NASA examples. The components created during the ITA analyses greatly facilitated the ESMD model creation. It is believed that many of the generic components created will be useful for the creation of future models. However, systems are developed and operated differently in the private sector than in large government agencies. Consequently, as more risk analysis models are created and validated, the repository of generic components should be augmented with generalized components found to be useful in various other systems.

7.3.4 MICRO-THEORIES OF RISK INCREASE IN COMPLEX SYSTEMS

As part of the interview-based component validation process, potential "micro-theories" of risk increase (or dynamic archetypes) were identified. Further validation, data gathering and empirical testing will be required to improve our understanding of how these archetypes influence risk dynamics in complex systems. Once better understood, the dynamic archetypes can be integrated in the repository of generic components and software packages to further facilitate model-building. Potential micro-theories and dynamic archetypes identified during interviews include:

7.3.4.1 Pressure to push safety analyses further downstream the development cycle

Some interviewees discussed the pressures felt to push safety analyses further down the development cycle in order to ensure that safety analyses do not impact development schedule:

“There’s a pressure early on to say the safety info is good enough in functional analysis. [...] Then when you get to design, there’s a pressure to say we can work it as part of verification, or maybe we don’t need everything right now, and especially when you get to manufacturing: ‘yeah, we can work that out later in operations’. This is where you compromise and there’s a lot of pressure at this point because now you’re starting to commit big money in manufacturing [sic]. And you want to get that hardware moving. So until manufacturing, there pressure to say: ‘Well, that’s enough work, we’ll catch up later’, after that, then: ‘maybe we can find workarounds in operation.’”

7.3.4.2 Pressure to reduce the Design-Analysis-Cycle (DAC) time

Many interviewees discussed the tradeoff dynamics associated with the length of design cycles:

“We design in cycles. The Design-Analysis Cycles (DAC) usually run on 6-week intervals. This creates pressure, and there’s a tendency to shorten the design cycle even more. But the fidelity of analysis goes down with shorter cycles, and fatigue kicks in, people make subtle mistakes. You need some schedule pressure, but not so much that you burn out people or make mistakes.”

“We shorten cycles to produce outputs used as someone else’s input. Shorter cycles help you make the schedule, but you lose fidelity. You can’t drag the cycle so long that you get better answers but waste other people’s time. So people sitting around and waiting creates a pressure to work faster to provide required info/work... and then when some work comes back with unexpected answers, it ripples into schedule.”

7.3.4.3 Balance in Personnel Rotation

Some interviewees discussed a problem where the regular rotation in program managers hinders the timely handling of critical issues:

“Each program manager knows he’s there four or five years at the most. If it’s not going to happen in those four or five years, if it’s -- let’s say it’s a life issue or something like that -- he’s going to push it on to the next guy. He says, ‘Well okay, it’s on my list of things to do, but it’s on the bottom.’ He knows he’s not going to be there by the time that thing is going to happen so it comes to the next guy. Now whoever brings that problem to him, is it going to be brought up with the same intensity that it was four or five years ago? I don’t know... If it’s a near term consequence that’s going to affect his four or five-year tenure on the program, he will probably do something about it. If it’s not within his window of when he’s going to be there, chances are he’s going to try to push that downstream. That’s just a fact of life, and it’s a sad fact of life.”

7.3.4.4 Requirements Waiver Accumulation

The accumulation of critical requirements waivers was discussed and documented extensively in the CAIB report. Requirement waivers are an ongoing topic of concern for managers and engineers. Parallels to the development and operation of complex systems in the private sector should also be investigated. An interviewee summarized the issue:

“The problem with pushing back on rules is once you push back and get one waiver, it opens a Pandora’s Box and then you get waivers on almost everything. It becomes the path of least resistance, which is: ‘I don’t have to do it because I can get a waiver’, because that’s two weeks worth of paperwork to get a waiver and not do the analysis. Really, a waiver should come with a lot of analysis, and if you do the analysis and find it is not an issue, then why do you need a waiver?”

7.4 CONCLUSION

The methodology and techniques presented in this thesis provide the foundation for STAMP-based dynamic model building and analysis. The framework further enhances the STAMP accident model by providing more powerful tools and techniques to mitigate time-dependent risks in complex systems. The methodology allows for seamless integration with the STAMP process, thus facilitating the analysis of identified risks. Throughout the NASA projects, the models created showed great potential to identify patterns of risk increase and the factors that cause them. Many of these risks increase patterns would not have been identified using existing methods. The models also allowed the design and testing of non-intuitive policies and processes to provide more effective and lasting mitigation of time-dependent risks in complex systems. Finally, it was shown that effective “warning systems” can be created by correlating the early indicators identified during model analysis with data collected in the real system. Monitoring and warning systems are used to detect states of increasing risk before the situation deteriorates to a point where an accident occurs.

APPENDIX A:

ACRONYMS

AA	Associate Administrator
ARC	Ames Research Center
CAIB	<i>Columbia</i> Accident Investigation Board
CaLV	Cargo Launch Vehicle
CD	Center Director
CDL	Center Discipline Lead
CEV	Crew Exploration Vehicle
CIL	Critical Items List
CLV	Crew Launch Vehicle
CSRL	Complex System Research Laboratory
DFRC	Dryden Flight Research Center
DTWH	Discipline Technical Warrant Holder
DTrA	Discipline trusted Agents
ESD	Engineering Systems Division
ESMD	Exploration Systems Mission Directorate
FMEA	Failure Modes and Effects Analysis
FTE	Full-Time Equivalent
FY	Fiscal Year
GRC	Glenn Research Center
GSFC	Goddard Spaceflight Center
HA	Hazard Analysis
HQ	Headquarters
ISS	International Space Station
ITA	Independent Technical Authority
JPL	Jet Propulsion Laboratory
JSC	Johnson Space Center
KSC	Kennedy Space Center
LaRC	Langley Research Center
LPRP	Lunar Precursor and Robotics Program
MDAA	Mission Directorate Associate Administrator
MIT	Massachusetts Institute of Technology
MSFC	Marshall Space Flight Center
NASA	National Aeronautics and Space Administration
NESC	NASA Engineering and Safety Center
NSF	National Science Foundation
OCE	Office of the Chief Engineer
OSMA	Office of Safety and Mission Assurance
PA&E	Program Analysis and Evaluation
P/P	Program/Project
P/PM	Program/Project Management

P/PM CE	Program/Project Chief Engineer
PRA	Probabilistic Risk Assessment/Analysis
RLEP	Robotic Lunar Exploration Program
S&MA	Safety and Mission Assurance
S&E	Science and Engineering
S.M.	Master of Science
Sc.D.	Doctor of Science
SE&I	System Engineering and Integration
SA	Safety Analysis
SOMD	Space Operations Mission Directorate
SOW	Statement of Work
SR&QA	Safety, Reliability, and Quality Assurance
SSC	Stennis Space Center
STAMP	Systems-Theoretic Accident Model and Process
STPA	STAMP Analysis
STWH	System Technical Warrant Holder
STrA	System Trusted Agents
USRA CPMR	Universities Space Research Association Center for Program/Project Management Research
USRA/NASA	Universities Space Research Association/NASA Academy for Program/Project Leadership
APPL	Program/Project Leadership
VSE	Vision for Space Exploration
WYE	Work Year Equivalent

APPENDIX B: TYPES OF UNCERTAINTY

This appendix complements the uncertainty section of chapter 1 by providing more detailed descriptions and examples of different types of uncertainties associated with complex systems.

B.1 AMBIGUITY

Ambiguity (sometimes called vagueness) arises from the imprecision associated with terms and expressions used for human communication. In practice, ambiguity can be reduced through the use of more precise linguistics and definitions, that is, by moving toward more formal syntax and semantics, but it is likely that ambiguity will remain an unavoidable part of human discourse. Ambiguity in a parameter stems from an inability to empirically measure it. Clarity tests have been proposed to verify that a statement is well-specified or non-ambiguous [Howard, 1984], despite some debate about whether ambiguity really is a type of uncertainty [Bedford, 2001]. In theory, through the use of a precise vocabulary using formally defined syntax and semantics, it would be possible to reduce ambiguity to an arbitrary low level. In practice, this reduction is rarely done because of the large effort involved.

B.2 ALEATORY UNCERTAINTY

Aleatory uncertainty (also called stochastic, probabilistic, random, inherent, intrinsic or physical uncertainty) represents inherent variations associated with a physical system. Aleatory uncertainty can usually be differentiated from other types of uncertainty by its representation as a distributed quantity whose range is known but whose exact value from time to time will depend on chance. The most frequently used representation for aleatory uncertainty is a probability distribution [Oberkampf, 1999]. Some events may appear aleatory only because we do not fully understand the underlying mechanisms that produce them. In this case, aleatory uncertainty would not be due to intrinsic randomness, but rather to a lack of


information about the event-generation process. A certain quantity may appear random to one observer while appearing deterministic to another observer who is aware of the mechanisms by which the quantity is generated. For example, the outcome of throwing a dice may appear completely aleatory, but in theory, if a sophisticated model was available to perfectly understand the dynamics of the dice as it hits the table, it would be possible to predict the outcome of the dice throw. In the absence of such a model, it is convenient to approximate the outcome as being random. Because of this theoretical possibility of using highly sophisticated models, there is some debate as to whether aleatory uncertainty is a separate type of uncertainty or merely another case of epistemic uncertainty.

B.3 EPISTEMIC UNCERTAINTY

Epistemic uncertainty stems from a lack of information in any phase of the modeling process. This definition stresses that the primary cause of epistemic uncertainty is incomplete information or knowledge about some characteristics of the system or its environment. Epistemic uncertainty can be further classified into model, phenomenological, and behavioral uncertainty.

B.3.1 MODEL

Model uncertainty (also called structural uncertainty) is related to the accuracy of the model used to represent a system. The use of various simplified relationships between variables used to represent the “real” relationships in a system is a form of model uncertainty [Melchers, 1999]. For example, the use of the equation in Figure 113 as a model of the dynamics of a mass-spring-damper mechanical system (see Figure 113) is a linear approximation of the highly nonlinear dynamics taking place in the real system including various friction and heat losses, material damping and hysteresis, among other nonlinear dynamic phenomena. Model uncertainty can be further decomposed into approximation error, associated with simplified relationships as in Figure 113, numerical errors associated with finite precision arithmetic, and model programming errors.



$$m\ddot{x} + c\dot{x} + kx = 0$$

Figure 113: Mass-Spring-Damper Simplified Model

B.3.2 PHENOMENOLOGICAL

Phenomenological uncertainty is associated with what engineers call “unk unks” or “unknown unknowns”. This type of uncertainty mostly arises during the development of novel envelope-pushing technologies or systems. During system development, at the time when critical decision-making must be made, some information will remain unknown. For example, for novel systems with little operational experience, some “failure modes” will not be identified, and if the system is unsuccessful, it may be due to a seemingly unimaginable scenario, or “unknown unknown”. This type of uncertainty is relevant to new complex systems, where unidentified hazards could cause accidents. However, many of the most recent aerospace accidents (and most accidents in general) were not caused by unknown risks, but rather by a failure to adequately manage and mitigate well known and documented risks [Leveson, 1995; Leveson, 2004].

B.3.3 BEHAVIORAL

Behavioral uncertainty is associated with uncertainty in the behavior of individuals and organizations. In the context of complex engineering systems, it can be further classified into: (1) Requirements uncertainty arising from incomplete or conflicting system requirements and uncertainty in the boundary between requirements and design decisions. (2) Design uncertainty associated with design choices that have not yet been decided upon, (3) Volitional uncertainty associated with difficulties in predicting the future behavior and decisions of other system stakeholders, and (4) Human uncertainty or error associated with possible mistakes or blunders during the development and operation of a system [Hastings, 2004].

B.4 INTERACTION

Interaction uncertainty is associated with difficulties to predict interactions between different components of a system or between different disciplinary areas. Interaction uncertainty often results from a lack of system integration activities, where the behavior of each component is well understood, but the interactions between components or discipline areas are not. Interactive complexity can greatly increase the difficulty of reducing interaction uncertainty.

This appendix complements the organizational risk theories section of chapter 1 by providing a more detailed analysis of the strengths and limitations of the Normal Accident Theory (NAT) and the High-Reliability Organization research.

C.1 NORMAL ACCIDENT THEORY (NAT)

Charles Perrow's initial formulation of what has come to be known as Normal Accident Theory (NAT) was developed in the aftermath of the accident at the Three Mile Island nuclear power plant in 1979 [Perrow, 1982]. Perrow introduced the idea that in some technological systems, accidents are inevitable or "normal" [Perrow, 1999]. He defines two related dimensions: interactive complexity and tight coupling, which determine a system's susceptibility to accidents.

Interactive complexity refers to *the presence of unfamiliar or unplanned and unexpected sequences of events in a system that are either not visible or not immediately comprehensible*. A tightly coupled system is one that is highly interdependent: *Each part of the system is tightly linked to many other parts and therefore a change in one part can rapidly affect the status of other parts*. Tightly coupled systems respond quickly to perturbations, but the response may be disastrous. Loosely coupled or decoupled systems have fewer or less tight links between parts and therefore have more capacity to absorb failures or unplanned behavior without major destabilization.

According to NAT, systems that are interactively complex and tightly coupled will experience accidents that cannot be foreseen or prevented. Perrow calls these system accidents. When the system is interactively complex, independent failure events can interact in ways that cannot be predicted by the designers and operators of the system. If the system is also tightly coupled, the cascading of effects can quickly spiral out of control before operators are able to understand the situation and perform appropriate corrective actions. In such systems, apparently trivial incidents can cascade in unpredictable ways and with possibly severe consequences.

C.1.1 NAT LIMITATIONS

Perrow made a significant contribution by identifying interactive complexity and tight coupling as risk-increasing system characteristics. However, his initial conclusion that nothing can be done to prevent accidents in complex interactive systems is very pessimistic and based upon the assumption that redundancy is the only engineering solution to improve safety. At the same time, he argues that redundancy may not help because it introduces additional complexity and encourages risk taking. In fact, he provides many examples of redundant safety devices or human procedures that were the direct cause of accidents. Sagan [Sagan, 1993] further documented how redundancy was the cause of many close calls in the handling of nuclear weapons and may effectively reduce the efficacy of nuclear security policies [Sagan, 2004].

Perrow's arguments and pessimism are based on the belief that reducing complexity and coupling will always be against the interest of the most powerful stakeholders, thus will not occur. This may not always be the case. In fact, Leveson [Leveson, 1995] argues that redundancy and the use of protection systems are among the least effective and the most costly approaches to designing for safety and describes many non-redundancy approaches to system design for safety. The most effective approaches involve eliminating hazards or significantly reducing their likelihood by means other than redundancy, for example, substituting non-hazardous materials for hazardous ones, reducing unnecessary complexity, decoupling, designing for controllability, monitoring, interlocks of various kinds, etc. Operations can also be made safer by eliminating and reducing the potential for human error. A simple example is the use of color coding and male/female adapters to reduce wiring errors.

Nevertheless, it is a fact that complexity and coupling are introduced in a system because they often allow greater functionality and efficiency to be achieved, usually at the cost of higher risk. However, simpler, decoupled designs can usually achieve the same goals. The problem boils down to minimizing tradeoffs and determining how much risk is acceptable. However, as the risk perception of various stakeholders will be very different, this problem does not have a single, and especially not a simple solution [Slovic, 1999].

C.2 HIGH RELIABILITY ORGANIZATIONS (HROS)

High Reliability HROs are defined by Roberts [Roberts, 1990] as the subset of hazardous organizations that enjoy a record of high safety over long periods of time:

“One can identify this subset by answering the question, ‘how many times could this organization have failed resulting in catastrophic consequences that it did not?’ If the answer is on the order of tens of thousands of times, the organization is ‘high’ reliability. [Roberts, 1990].”

The field of High Reliability Organizations research is based on observations made during the study of two aircraft carriers, U.S. air traffic control, utility grid management, and fire fighting teams [La Porte, 1991]. These observations seem to counter Perrow’s hypothesis by suggesting that some interactively complex and tightly coupled systems operate for long periods of time with very few accidents.

The literature associated with the HRO field is large and growing. Nevertheless, most HRO researchers agree on four primary organizational characteristics that they claim substantially limit accidents and simultaneously result in high levels of performance: (1) prioritization of both safety and performance and consensus about the goals across the organization [La Porte, 1991]; (2) promotion of a “culture of reliability” in simultaneously decentralized and centralized operations [Weick, 1987]; (3) use of organizational learning that maximizes learning from accidents, incidents, and near misses [La Porte, 1991]; and (4) extensive use of redundancy [Rochlin, 1987].

C.2.1 HRO LIMITATIONS

There are problems associated with the HRO definition used by researchers, namely: *“an organization where tens of thousands of potentially catastrophic events did not result in catastrophic consequences”*. In fact, this type of safety record would probably not be acceptable, except in cultures or countries where frequent catastrophic consequences and death is tolerable. A “catastrophic failure” can be hypothesized to constantly be within the realm of possibilities for most high-hazard activities.

Apart from definition problems, HRO researchers claim that the systems they studied fall within the type of interactively complex and tightly coupled systems described by Perrow. There is some controversy associated with this claim. For example, one could argue that air traffic control (ATC), for example, is as safe as it is precisely because the system design is deliberately decoupled in order to increase safety. The ATC system is carefully divided into non-interacting sectors and flight phases (enroute, arrival, and takeoff and landing) with the interfaces between the sectors and phases (for example, handoff of an aircraft between two air traffic control sectors) carefully limited and controlled. Loose coupling is also ensured by maintaining ample separation between aircraft so that mistakes by controllers can be remedied before they impact safety. Different parts of the airspace are reserved for different types of aircraft or aircraft operation (e.g., visual flight rules vs. instrument flight rules). Extra warning devices, such as collision avoidance systems, also exist to prevent accidents.

The fact that these loosely coupled systems are safe seems to support Perrow's arguments rather than contradict them. The High Reliability Organization researchers themselves emphasize the low level of complexity in the example systems they studied:

“HROs struggle with decisions in a context of nearly full knowledge of the technical aspects of operations in the face of recognized great hazard ... The people in these organizations know almost everything technical about what they are doing—and fear being lulled into supposing they have prepared for every contingency ... This drive for technical predictability has resulted in relatively stable technical processes that have become quite well understood within each HRO. [La Porte, 1991] (emphasis added).”

The fact that these systems allow perfect knowledge contradicts the definition of interactive complexity, which Perrow defined as system designs for which the interactions between components could not be thoroughly planned, understood, predicted, or guarded against.

Another potential limitation is the difficulty of applying the four HRO best practice principles outside the systems where they were observed, namely:

C.2.1.1 Prioritization of both safety and performance and consensus about the goals across the organization.

For aircraft carriers during peace time, the primary goal is to get aircraft landed and launched safely or, if that goal is not successful, to safely eject and recover the pilots: There are no goal conflicts with safety. If conditions are risky, for example, during bad weather, flight operations can be delayed or canceled without major consequences. In a different context, it may be much more difficult to prioritize safety as much as performance. The grounding of the *USS Enterprise* and the accidental shooting down of an Iranian commercial aircraft by the *USS Vincennes* indicate that combat conditions have a strong effect on high reliability performance [Rochlin, 1991].

C.2.1.2 Promotion of a “culture of reliability” in simultaneously decentralized and centralized operations.

The second characteristic of High Reliability Organizations is that organization members are socialized and trained to provide uniform and appropriate responses to crisis situations [Weick, 1987]. This field-level response to crises is the “decentralized response” that forms such a large part of HRO philosophy. The other side, “simultaneous centralization,” refers to the maintenance of clear chains of command in crisis situations. For example, La Porte and Consolini [La Porte, 1991] argue that while the operation of aircraft carriers is subject to the Navy’s chain of command, even the lowest-level seaman can abort landings. Clearly, this local authority is necessary in the case of aborted landings because decisions must be made too quickly to go up a chain of command. Overtraining of emergency responses is a standard practice in the training of operational personnel working in potentially dangerous, time-critical conditions. Note also that low-level personnel on aircraft carriers may only make decisions in one direction, that is, they may only abort landings. The actions governed by these decisions and the conditions for making them are relatively simple.

More interesting cases arise when decision-making is not time critical. La Porte and Consolini [La Porte, 1991] state that all personnel, regardless of rank, are trained to own a problem when they see it until it is solved or until someone who can solve the problem takes responsibility for it. This approach works only because the systems they studied were loosely coupled. In systems that are interactively complex and tightly coupled, taking individual

action and acting alone may lead to accidents when local decisions are uncoordinated with other local or global decisions. Figure 114 shows an analysis by Rasmussen of the Zeebrugge ferry accident [Rasmussen, 1997]. Those making decisions about vessel design, harbor design, cargo management, passenger management, traffic scheduling, and vessel operation were unaware of the impact of their decisions on the others and the overall impact on the process leading to the ferry accident. The type of bottom-up decentralized decision-making advocated for HROs can lead to major accidents in complex socio-technical systems.

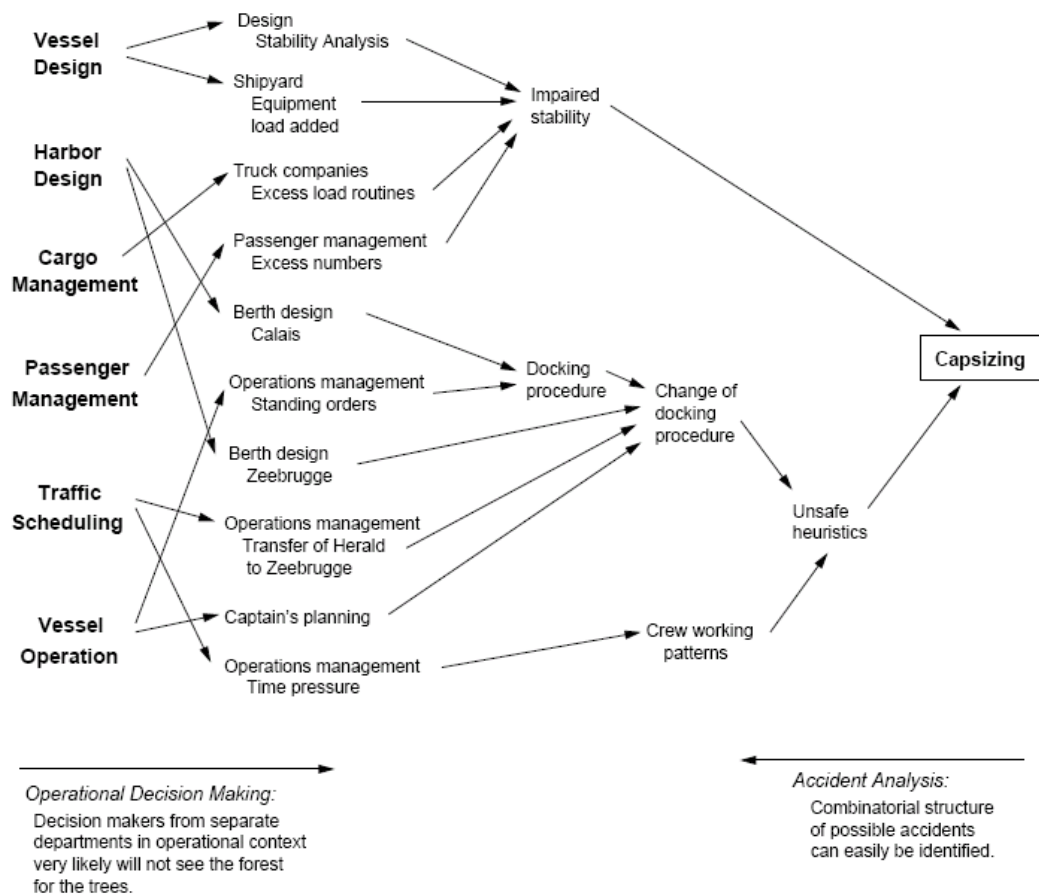


Figure 114: The Interactions in the Zeebrugge Ferry Accident (from [Rasmussen, 1997])

C.2.1.3 Use of organizational learning that maximizes learning from accidents

A third characteristic of HROs claimed by some proponents of this theory is that they use sophisticated forms of organizational learning. The argument is made that limiting learning to trial and error is not practical in these organizations. Instead, HROs use “imagination,

vicarious experiences, stories, simulations, and other symbolic representations of technology and its effects” as substitutes for trial-and-error learning [Weick, 1987]. This process resembles what engineers do in a more rigorous way when performing hazard analysis. More interesting is the claim by some that HROs try to maximize learning from accidents, incidents, and near misses [La Porte, 1991]. It is difficult to argue against learning from mistakes, but for complex safety-critical systems where mistakes can have disproportionate consequences, the difficulties of learning should not be underestimated. Many difficulties have been identified that limit the effectiveness of learning from a small number of problems and near misses. March [March, 1991] provides a summary of the benefits and limitations of learning from limited experience. Among these difficulties is the fact that learning is limited to acceptable interpretations of the causes of accidents and incidents. Political, financial and legal concerns may limit the set of acceptable interpretations from which learning can occur. Moreover, it was shown that individuals identify different accidents “causes” depending on their own position and status [Leplat, 1987]. Other researchers suggest that accidents are often attributed to factors in which the individuals are less directly involved, consistent with classic work in attribution theory [Jones, 1967; Ross, 1977].

A difficulty of learning from near-misses or hypothetical incidents is that it is very expensive to construct hypothetical histories from which to learn. Consequently, the costs of implementing effective organizational learning are high and the problems of competition for resources arise again. Moreover, the impact of hypothetical histories ordinarily cannot compare with the dramatic power of realized history [Fischhoff, 1975]. It is almost impossible to match the power of actual dramatic events on beliefs and potential learning. Think about the profound impact of the 1986 Challenger explosion on the beliefs about the safety of space travel.

March [March, 1991] also describes another difficulty related to the use of hypothetical histories for learning, especially for organizations dealing with complex safety-critical systems:

“[...] hypothetical histories may be ambiguous and thus unpersuasive. Where organizations face possible events of great consequence but small likelihood, the use of near-histories to augment simple experience is sometimes controversial. If

the probability of disaster is very low, near-histories will tend to picture greater risk than will be experienced directly by most organizations or individuals in a reasonable length of time. In such case, near-histories are likely to be treated as generating too pessimistic a picture. For example, long before the fatal Challenger flight, the spacecraft flew a series of successful missions despite its faulty O-rings. Some engineers interpreted the indications of O-ring problems during these early flights as symptoms that past successes had been relatively lucky draws from a distribution in which the probability of disaster was relatively high [Boisjoly, 1987]. Others, including some key personnel in NASA, considered these estimates of danger as exaggerated because, in the realized history, the system had been robust enough to tolerate such problems [Starbuck, 1988].”

Trial and error is not a very efficient way to learn, particularly for complex systems where the factors that can be involved in accidents may be very large. A billion dollar Milstar satellite was lost when only past errors that had led to accidents were considered; the human error that led to the loss (a decimal-point mistake in manually copying a roll rate filter constant in the Centaur launcher attitude control system) had never been identified as the cause of an accident before and no safeguards were implemented. The most important point here is that learning from accidents is not the only or even the most effective way to lower risk in high-tech systems. The organizations studied by HRO researchers are characterized by unchanging or very slowly changing designs and technology, which makes learning from accidents and incidents more effective. Organizations like NASA and the military that operate at the edges of technological innovation do not always have past experience from which to learn. Lessons learned on old technologies are also often inapplicable to newer ones. For example, digital systems are changing the nature of accidents and even changing the types of errors made by operators [Sarter, 1996; Leveson, 2006]. Experience with older, electro-mechanical systems does not apply to these new system designs and technology.

In addition, the difficulty of implementing effective organizational learning should not be underestimated. Among others, Senge [Senge, 1990] provides a good summary of behavioral patterns and archetypes responsible for resistance to learning and change in organizations.

C.2.1.4 Extensive use of redundancy

A fourth characteristic often cited about HROs is the extensive use of redundancy. HROs are “characterized especially by flexibility and redundancy in pursuit of safety and performance,”

[La Porte, 1996] where redundancy is defined as “the ability to provide for the execution of a task if the primary unit fails or falters” [La Porte, 1991]. According to Roberts, HROs use technical redundancy, where parts are duplicated (e.g., backup computers) and personnel redundancy, where personnel functions are duplicated (e.g., more than one person is assigned to perform a given safety check) [Roberts, 1990]. On aircraft carriers, for example, control for setting the arresting gear ultimately rests in the hands of at least three people, with oversight from the carrier’s airboss.

The role of redundancy in increasing the safety of socio-technical systems is a major point of disagreement between Normal Accident Theory (NAT) and HRO. Once again, the problem seems to be that the proponents of each are arguing about different types of systems. Interactive complexity, tight coupling, and working in environments of uncertainty and imperfect knowledge limit the effectiveness of redundancy and, in fact, redundancy under these circumstances can actually increase the risk of an accident, as Perrow and Sagan [Sagan, 1993; Perrow, 1999; Sagan, 2004] so persuasively argued.

It is sometimes argued [Marais, 2004; Perrow, 2004] that some of the systems studied by HRO researchers do not exhibit the characteristics and dynamic patterns of interactively complex, tightly-coupled systems. Moreover, HRO researchers admit that some of these systems are well-understood and characterized by low levels of uncertainty [La Porte, 1991]. In these systems, redundancy can be effective in preventing a single component failure (or sometimes multiple component failures) from leading to an accident. Even in these cases, however, there are limitations. For example, common-mode failures, where supposedly independent redundant components fail due to the same cause, limit the effectiveness of redundancy in protecting against component failure. A Lockheed L-1011 flying out of Miami in 1983 [NTSB, 1983] lost oil pressure in all three engines simultaneously because both mechanics did not put O-rings on three newly installed engine oil plugs. Inadequate preventive maintenance is only one type of common error that will affect all components, including the supposedly redundant backups.

Redundancy depends on an assumption of random failure of the components for its effectiveness. But many, if not most, causes of accidents in interactively complex and tightly-

coupled systems do not involve random component failure. In fact, engineers designing complex systems have long known about and tried to protect designs against accidents where no components fail, i.e., the accident is caused by dysfunctional interactions among components. The same is true for human components and human decision making. As illustrated by the Zeebrugge accident analysis described in the previous section, individual decisions may be reasonable in their context but not when combined at the system level.

The use of redundancy can, in fact, lead to dangerous decision making when false reliance is based on it and the need for additional safety measures is discounted. The decision to launch the Challenger Space Shuttle on its fatal flight was partly based on over-reliance on redundant O-rings. The failure of the primary O-ring led to the failure of the secondary O-ring [Leveson, 1995]. Redundancy does not protect against underlying design errors, only random failures. Worse, the overconfidence provided by the redundancy convinced decision-makers that the Shuttle would survive a cold-weather launch even if the primary O-ring failed.

When systems contain software, redundancy is not useful in protecting against commands that can lead to accidents. In fact, most software-related accidents can be traced back to errors in the software requirements, i.e., a misunderstanding about what the software is supposed to do under some circumstances. In these accidents, the software did not fail in the same way as hardware fails. Unless there was an implementation error, the software did exactly what the programmers intended it to do, which is different from a hardware failure where the system did not do what it was intended to do. In addition, software redundancy management systems are so complex that they often introduce errors and can lead to system failures themselves.

Redundancy is only one limited way to increase reliability (but not necessarily safety) in some special cases; under other circumstances it can be the cause of or contributor to accidents.

C.3 DEBATE AND LIMITATIONS

For years, there has been a standing, raging debate in the organizational safety literature between advocates of the Normal Accident Theory and that of HRO proponents. The debate has been left without a clear winner and has died down in the past years, arguably after a recognition of the impossibility to resolve it [Rijpma, 1997; Rijpma, 2003]. Much has been

said about the debate, and the purpose of this section is not to provide an exhaustive discussion about the pros and cons of each approach. A more thorough analysis can be found in [Marais, 2004]. It is useful to note however, that from the author's point of view, both approaches and visions oversimplify the cause of accidents. HRO underestimates the problems of uncertainty. NAT recognizes the difficulty of dealing with uncertainty but underestimates and oversimplifies the potential ways to cope with uncertainty. Both seem to believe that redundancy is the only way to handle risk. The contribution of Perrow to understanding accidents in complex systems by identifying interactive complexity and tight coupling as critical factors should not be discounted. His top-down system view of accidents versus the bottom-up, component reliability view of the HRO theorists is critical in understanding and preventing future accidents. But the theory is incomplete and leads to more pessimism than required with respect to designing and operating complex high-risk systems. While the HRO theorists do offer more suggestions, most of them are either prohibitively costly or even inapplicable to complex systems. In other cases, the suggestions oversimplify the problems involved and will not be effective.

Another popular theory of organizational risk perception and acceptance was developed by Diane Vaughn based on data gathering and observations related to the 1986 Challenger accident. The Normalization of Deviance theory (NoD) [Vaughan, 1996] that resulted is a five-step "risk normalization" process responsible for an increase over time of the risk considered acceptable by certain organizations. The theory has received considerable attention from sociologists, but has had little impact on the design and operation of complex engineering systems. The Challenger accident did not result simply because of a knowingly accepted increase in the risk level. Rather, some aspects of system behavior had come to be acceptable to managers and engineers based on insufficient supporting data. Moreover, there was a loss of sight of the priority of hazard mitigation strategies, combined with an over-reliance on the seemingly independent redundancy provided by the second field joint O-ring [Rogers, 1986]. Instead of focusing on hazard mitigation, the focus switched to a negotiation into how much requirements violation was acceptable. This change in focus was the result of many factors including overly powerful schedule and budget pressures, which were adequately portrayed by Vaughan. Even though the "Normalization of Deviance" framework itself may not accurately portray the way risk increases in complex systems, the Vaughan

study correctly emphasized the time-dependent processes of risk increase in socio-technical systems. A more thorough analysis and critique of the NoD theory is provided in [Marais, 2005] and [Leveson, 2004].

D.1 ITA CONTEXT AND MODELING

During the study of the impact of the ITA on the safe operation of NASA's space shuttle fleet [Leveson, 2005], we discovered that the ITA impact was highly reinforcing, but that this reinforcing effect could be positive or negative, depending on model parameters. A model was created based on the Technical Authority Implementation Guidance document to capture the effect of ITA implementation on the dynamics of the system (see [Leveson, 2005]). From a system-level perspective, the credibility and effectiveness of ITA directly affects Launch Rate, System Safety Efforts and Efficacy, and the way Incident Learning and Corrective Actions are performed, including the strength of the safety program and the handling of requirements waivers. In the other direction, the Credibility and Effectiveness of ITA is directly affected by the availability of employees with high levels of System Safety Knowledge and Skills.

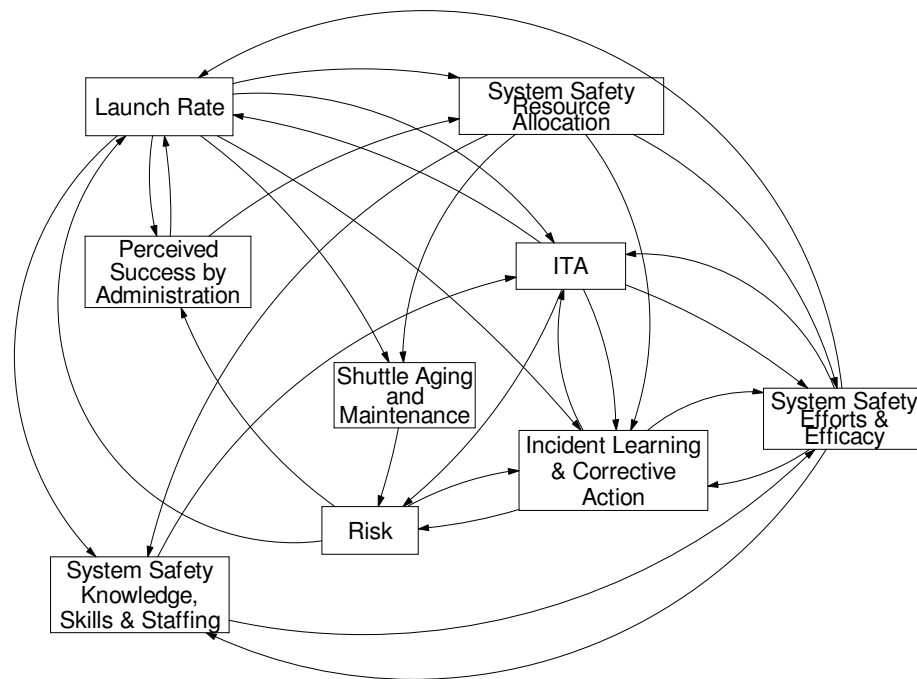


Figure 115: The Nine Subsystem Models and their Interactions

According to the ITA Implementation Plan, Technical Warrant Holders (TWHs) are supposed to be unaffected by schedule pressure. Trusted Agents, however, still have obligations to the project manager, so schedule pressure and Launch Rate will still affect their work. The

effectiveness of ITA personnel is highly dependent on the quality, thoroughness and timely availability of safety analysis performed by safety experts and, therefore, it is directly affected by the System Safety Efforts and Efficacy, modeled separately. The number of open incident investigations and waiver resolutions may also affect the workload and effectiveness of the ITA designees. Finally, as the “guardian” of NASA’s technical conscience, ITA promotes the raising of safety issues and concerns that can result in proactive changes in the system that will decrease system Risk.

Figure 116 provides an overview of the internal feedback structure of the ITA model. The internal dynamics of this model are highly reinforcing. Four internal reinforcing loops create these dynamics: Attractiveness of being a TWH, TWH Resources and Training, Ability to Attract Knowledgeable Trusted Agents, and Trusted Agent Training Adequacy. If the external influences from outside parts of the model were removed, the Effectiveness and Credibility of the ITA would either grow rapidly (if left unbounded) or would collapse. The reinforcing polarity depends on the gain of each loop at every instant in time. A highly effective and credible ITA will have high Influence and Prestige, resulting in a great ability to attract highly skilled and well-respected technical leaders, ensuring the TWHs have enough power and authority to perform their functions. In addition, an effective and credible ITA will be able to obtain and manage the resources necessary for their functioning and to ensure that TWHs and Trusted Agents are provided with the resources and training necessary to remain highly knowledgeable and effective over time. On the flip side, these interactions can create a downward spiral that will act in the opposite direction.

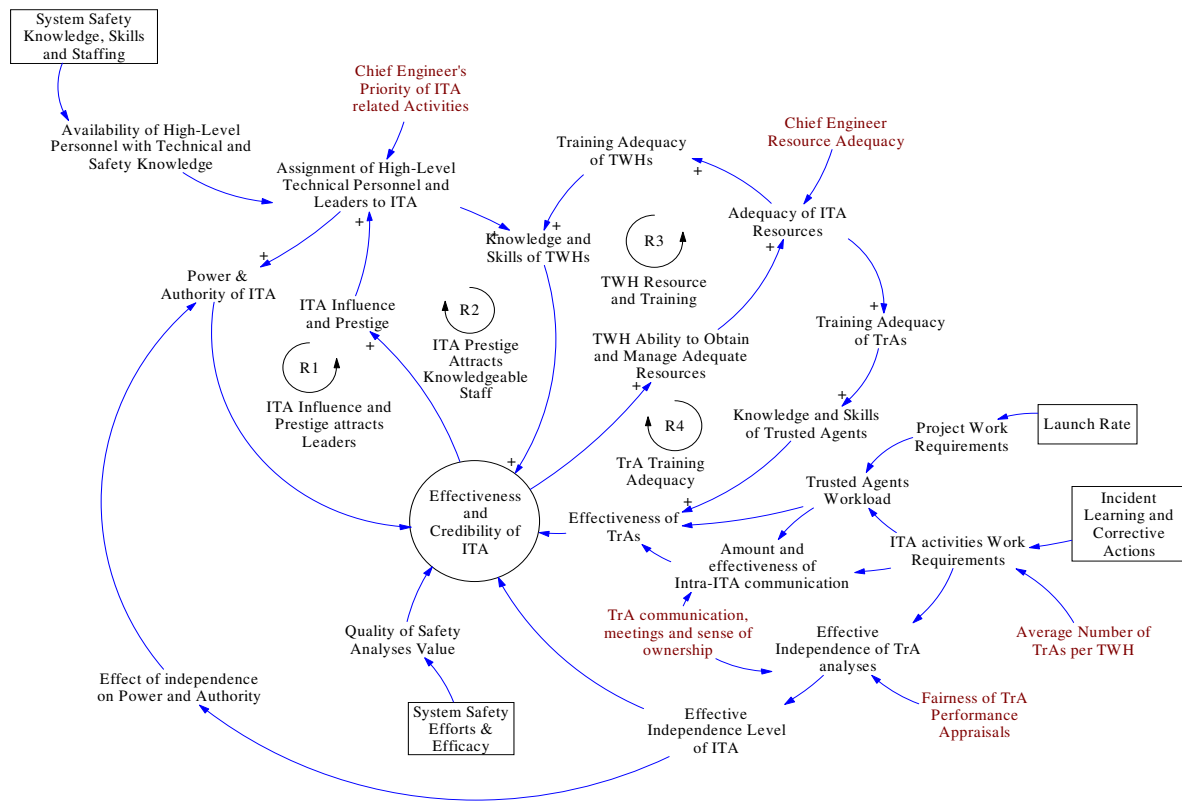


Figure 116: ITA Model Structure

D.2 INITIAL ITA MODEL ANALYSIS

While testing the model, we realized that the reinforcing polarity depends on initial values and on the value of exogenous parameters of the ITA model, mainly: Chief Engineer's Priority of ITA-related activities, Chief Engineer Resource Adequacy, Average number of Trusted Agents per Warrant Holder, Fairness of Trusted Agents Performance Appraisals, and Trusted Agents Communication, Meetings and Sense of Ownership. However, when the loops are closed and the ITA model is integrated within the system, many other balancing loops affect the behavior of the system and the dynamics become more complex.

In order to investigate the effect of ITA parameters on the system-level dynamics, an initial 200-run Monte-Carlo sensitivity analysis was performed. Random variations representing +/-

30% of the baseline ITA exogenous parameter values were used in the analysis. Figure 117 shows the results of the 200 individual traces.




Figure 117: ITA Sensitivity Analysis Trace Results

The initial sensitivity analysis shows that at least two qualitatively different system behavior modes can occur. The first behavior mode (behavior mode #1, Figure 117) is representative of a successful ITA program implementation where risk is adequately mitigated for a relatively long period of time. More than 75% of the runs fall in that category. The second behavior mode (behavior mode #2 in Figure 117) is representative of a rapid rise and then collapse in ITA effectiveness associated with an unsuccessful ITA program implementation. In this mode, risk increases rapidly, resulting in frequent hazardous events (serious incidents) and accidents.

D.3 ITA BEHAVIOR MODE ANALYSIS

Because the results of the initial ITA sensitivity analysis showed two qualitatively different behavior modes, we performed detailed analysis of each to better understand the behavior exhibited. Using this information, we were able to identify some potential metrics and indicators of increasing risk as well as possible risk mitigation strategies. The ITA support structure is self-sustaining in both behavior modes for a short period of time if the conditions are in place for its early acceptance. This short-term reinforcing loop provides the foundation

for a solid, sustainable ITA program implementation. The conditions under which this initial success continues or fails is important in identifying early warning metrics.

D.3.1 BEHAVIOR MODE #1: SUCCESSFUL ITA IMPLEMENTATION

Behavior mode 1, successful ITA program implementation, includes a short-term initial transient where all runs quickly reach the maximum Effectiveness and Credibility of ITA. This behavior is representative of the initial excitement phase, where the ITA is implemented and shows great promise to reduce the level of risk. After a period of very high success, the Effectiveness and Credibility of ITA slowly starts to decline. This decline is mainly due to the effects of complacency: the quality of safety analyses starts to erode as the program is highly successful and safety is increasingly seen as a solved problem. When this decline occurs, resources are reallocated to more urgent performance-related matters and safety efforts start to suffer. The decrease in Effectiveness and Credibility of ITA is not due to intrinsic ITA program problems, but rather to a decrease in the quality of safety analysis upon which the TA and TWHs rely.

In this behavior mode, the Effectiveness and Credibility of ITA declines, then stabilizes and follows the Quality of Safety Analyses coming from the System Safety Efforts and Efficacy model. A discontinuity occurs around time 850 (denoted by the arrow on the x-axis of Figure 117), when a serious incident or accident shocks the system despite sustained efforts by the TA and TWHs. At this point of the system lifecycle, time-related parameters such as vehicle and infrastructure aging and deterioration create problems that are difficult to eliminate.

Figure 118 shows normalized key variables of a sample simulation representative of behavior mode #1, where the ITA program implementation is successful in providing effective risk management throughout the system lifecycle. As previously mentioned, although the Effectiveness and Credibility of ITA starts to decline after a while, due to eroding System Safety Efforts and Efficacy, ITA remains effective at mitigating risk and is able to avoid accidents for a long period of time. This behavior mode is characterized by an extended period of nearly steady-state equilibrium where risk remains at very low levels.

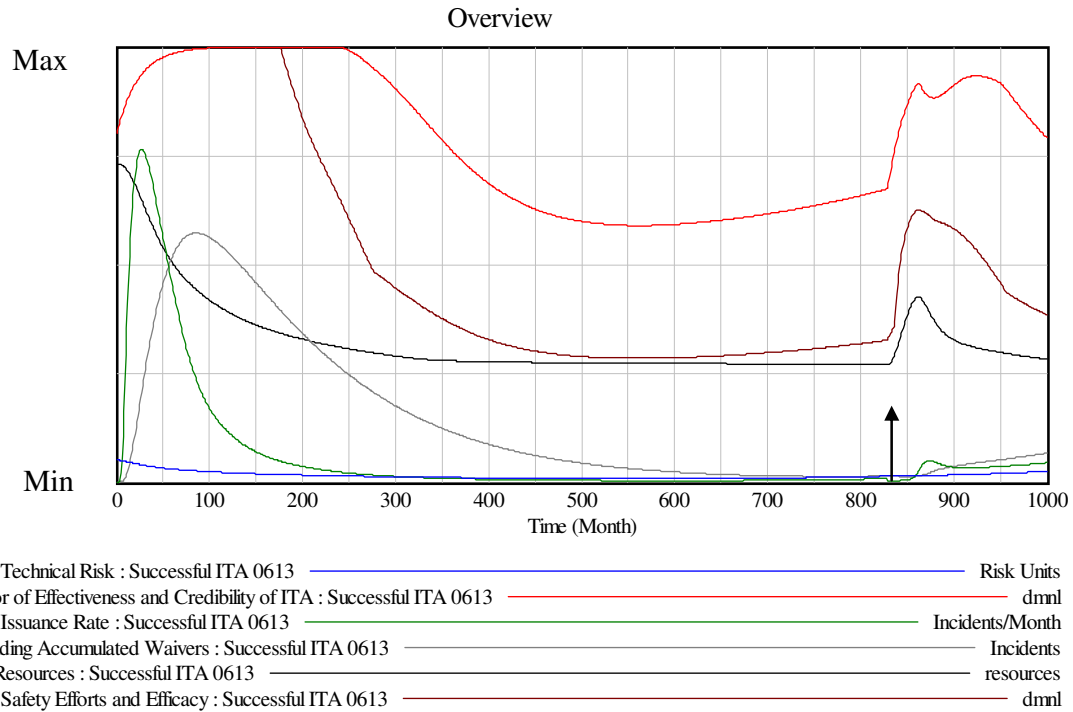


Figure 118: Behavior Mode #1 Representing a Successful ITA Program Implementation

D.3.2 BEHAVIOR MODE #2: UNSUCCESSFUL ITA IMPLEMENTATION

In the second behavior mode (behavior mode #2 in Figure 117), Effectiveness and Credibility of ITA increases in the initial transient, then quickly starts to decline and eventually reaches bottom. This behavior mode represents cases where a combination of parameters (insufficient resources, support, staff...) creates conditions where the ITA structure is unable to have a sustained effect on the system. As ITA decline reaches a tipping point, the reinforcing loops mentioned previously act in the negative direction (they have a <1 loop gain) and the system migrates toward a high-risk state where accidents and serious incidents occur frequently (at the arrows on the x-axis in Figure 119).




Figure 119: Behavior Mode #2 Representing an Unsuccessful ITA Program Implementation

The key normalized variables for a sample simulation run representative of the second behavior mode are shown in Figure 119. These variables are: System Technical Risk, Indicator of Effectiveness and Credibility of ITA, Waiver Issuance Rate, Outstanding Accumulated Waivers, Safety Resources, and System Safety Efforts and Efficacy. This behavior mode represents an unsuccessful implementation of the ITA program. As risk increases, accidents start to occur and create shock changes in the system. Safety is increasingly perceived as an urgent problem and more resources are allocated for safety analyses, which increase System Safety Efforts and Efficacy, but, at this point, the TA and TWHs have lost so much credibility that they are not able to significantly contribute to risk mitigation anymore. As a result, risk increases dramatically, the ITA personnel and safety staff become overwhelmed with safety problems and starts to issue a large number of waivers in order to continue flying. This behavior mode includes many discontinuities created by the frequent hazardous events and provides much useful information for selection of metrics to measure the effectiveness of ITA and to provide early indication of the system migrating toward a state of increased risk.

This appendix provides a repository of generic components to be used along with the component-based model creation methodology presented in chapter 4, the model and policy analysis techniques of chapter 5 and the ESMD case study presented in chapter 6. Before the components are presented, the research approach, data sources and projects used to create and validate the components are documented.

E.1 RESEARCH APPROACH FOR COMPONENT CREATION

The research approach used to create the dynamic components introduced in this Appendix combines the use of grounded theory building methodology [Creswell, 1994; Strauss, 1994], as well as causal loop model building exercises and formal simulation models [Sterman, 2000] to produce an inductively-derived understanding of the dynamics associated with the migration of complex safety-critical socio-technical systems toward a state of high-risk. The research process includes multiple iterations between observations, expert interviews, data-collection, causal-loop inferences and diagramming, as well as links to relevant literature and formal model building and analysis to converge toward components that provide grounded qualitative micro-theories of safety-critical systems migration toward high-risk that can be used as building blocks for the larger custom model creation and analysis methodology introduced in this thesis.

In summary, the component creation approach combines the following elements:

- Data and information collection from interviews, observations, documents and reports
- Relevant literature as a source of understanding, insight and context-creation for quasi-validation
- Causal relationships inductively or deductively derived and reviewed by safety and domain experts
- Formal model building, simulations, and analysis
- Generalization of findings and re-structuring into reusable micro-theories of risk increase used as the basis for the creation of generic customizable risk management model components

The research progressed along the following steps:

- 1- Formulation of the problem space and boundary based on experience, observations and intuition.
- 2- Secondary data gathering based on interactions with expert groups.
- 3- Coding and sorting of the data collected into emerging categories through the use of traditional qualitative research methods
- 4- Proposition of causal links between categories and mapping through the use of causal loop diagramming
- 5- Review of causal loop formulations by system safety and domain experts
- 6- Creation of formal dynamic simulation models based on the causal loop formulations, as well as data gathered from experts, documents, and literature.
- 7- Performing multiple scenario analyses on the formal model and compare the results to available documents, data and literature.
- 8- Presenting the results to system safety and domain experts in order to elicit feedback to provide a first-order validation of model behavior
- 9- Based on feedback results, performing further data collection using documents, literature and interviews with experts
- 10- Refining the model formulations and continue to perform formal simulations and analyses while collecting insights and lessons learned along the way.
- 11- Converging to micro-theories of migration of complex safety-critical socio-technical systems toward a state of high-risk
- 12- Using the formulated micro-theories, as well as insights and lessons learned along the way, generalize the dynamic micro-theories obtained and create a toolset and methodology to facilitate the creation of customized models used to detect, monitor, and prevent the migration of specific systems to a state of higher risk.

While these steps appear to be well-ordered and sequential, the actual process is highly iterative and frequent backtracking to previous steps is the norm rather than the exception. Various sources of data were used in the creation of micro-theories of risk increase to form the basis of the component-based model creation process.

E.2 PROJECTS AND CASE STUDIES

The components presented in this appendix are based on the generalization of micro-theories of migration toward high risk that were initially developed through their use in real projects and case studies. Two main projects were used as the basis of the extraction of generic micro-theories of migration toward high-risk. One project was operation-centric, and another was development-centric. For the operations part of the system, the example used is the model created for the space shuttle risk analysis and the NASA Independent Technical Authority

(ITA) risk and vulnerability analysis. For the development part of the system, the example used is the model created for the risk analysis performed for the NASA Exploration Systems Mission Directorate (ESMD).

E.3 DATA SOURCES

This section describes various sources of data that were used in the creation of micro-theories and generic model components to facilitate the creation of dynamic risk management models.

E.3.1 GROUP MODEL BUILDING AND REVIEW FROM SAFETY AND ORGANIZATION THEORY EXPERTS

Theory-building activities include causal loop mapping and review by multiple faculty members in a multidisciplinary research group on system safety created at MIT in the wake of the Columbia accident. Members of the group included graduate students, research staff and faculty members from the Aeronautics and Astronautics Department, the Department of Engineering Systems Design, and the Sloan School of Management. Some members of the group have extensive knowledge of the internal functioning of NASA (and other complex systems) through a long-term association with the Agency.

E.3.2 BOOKS, ACCIDENT REPORTS, AND RISK LITERATURE

Several books and sources were used to capture and understand the causal structure, including books on NASA's safety culture such as McCurdy [McCurdy, 1994; McCurdy, 2001], books on the Challenger [Vaughan, 1996] and Columbia accidents [Hollnagel, 2005], NASA mishap reports (CAIB [Gehman, 2003], Mars Polar Lander [Young, 2000], Mars Climate Orbiter [Stephenson, 1999], WIRE [Branscome, 1999], SOHO [NASA, 1998], Huygens [Link, 2000], etc.), other NASA reports on the manned space program (SIAT [MacDonald, 2000] and others) as well as many of the better researched magazine and newspaper articles, some of which have been reviewed in the first two chapters of this thesis.

E.3.3 INTERVIEWS WITH DOMAIN EXPERTS

While much data can be obtained from the literature and accident reports, specific questions about causal links and decision-rules used during system development and operation can only be answered by domain experts involved in daily system activities. In order to get answers to these specific questions, NASA employees were interviewed at five different NASA centers. Informal interviews were conducted at the Kennedy Space Center (Cape Canaveral, FL) before the ITA risk and vulnerability analysis was performed. During the ESMD project, more formal interview sessions were conducted using a semi-structured format. In all, 41 interview sessions were conducted, with a total of 44 interviewees at the following NASA centers: NASA Headquarters (Washington, D.C), Marshall Space Flight Center (Huntsville, AL), Johnson Space Center (Houston, TX) and Langley Research Center (Langley, VA). Over 200 pages of interview transcripts were used as the basis of the component and model creation and validation activities. Specific details of the interview logistics and protocol are presented in Chapter 6.

E.3.4 DOMAIN-SPECIFIC QUANTITATIVE DATA SOURCES

The primary source of qualitative data used in the NASA projects comes from interviews with NASA employees and domain experts. In addition, various sources of quantitative data were used. Detailed quantitative data about resource allocation and development progression were content-sensitive and unavailable for our studies. However, various sources of NASA-specific quantitative data were found in the public domain and used to calibrate the assumptions underlying our theory and model-building activities. The sources of quantitative data are discussed in Chapter 6, but include multi-year data on NASA civil servant and support contractor headcounts and characteristics (technical area, age, experience, retirement eligibility, etc.), NASA budget requests and procurement reports, as well as space shuttle program data provided in the CAIB report [Gehman, 2003].

E.4 DATA ANALYSIS

Data analysis began at the very beginning of the observation phase and continued throughout the projects following a highly iterative trajectory. Traditional data coding and sorting approaches [Creswell, 1994; Strauss, 1994] were used to create categories that were further linked through causal relationships inductively derived from the available data. The newly formulated causal relationships suggested further data gathering and reviews that were used to refine the causal links and eventually the formal model.

Diagramming of candidate causal loops and relationships was used throughout the research to elicit information and feedback from experts during group activities and interviews, as suggested by Strauss and Corbin [Strauss, 1994]. The diagramming technique was used during and after review sessions and interviews to capture the most important points and candidate causal relations. This practice is also in line with Sterman's [Sterman, 2000] argument that causal-loop diagramming is a particularly useful method to capture and understand causal relationships in complex systems.

The qualitative analysis was followed by formal model building. As causal loops by themselves do not carry the richness associated with the dynamics of the interactions they represent, formal models were created and simulations were performed. The objective of the simulation phase was twofold. First, simulation allowed us to better understand the contribution of individual components to the dynamic processes associated with the system-level migration of socio-technical systems toward high risk. Consequently, simulations were critical in the development of micro-theories of migration toward high-risk used in generic dynamic components. Second, simulations allowed further validation and improved confidence in the causal relationships defined earlier by ensuring that the model outputs was adequately correlated to data and to experts' mental models of the expected reference modes associated with various simulation scenarios.

As suggested by Mass [Mass, 1991], surprise model behavior was used throughout the model-building and simulation phase to elicit possible problems with the formulations of causal structures and decision-rules, as well as to investigate potential insights derived from observations of simulation outputs. As mentioned previously, surprise behavior should be investigated at any point of the model lifecycle to ensure that it does not result from a problem

in the formulations of the causal relations or the decision-rules embedded in the model. Surprise behavior ranged from completely unexpected model behavior to unexpected discontinuities or oscillations in some variables that should be investigated.

The theory and model building activities underwent multiple refinement cycles where a tentative model was created, documented and presented to various audiences including academics and experts in aerospace engineering, system safety, organizational learning, labor relations, history of technology and political science both at NASA and at MIT. New insight, reviews and data collection provided additional information for further model development and analysis. This ongoing cycle of model refinement and insight collection continues until the model provides acceptable overlap and agreement with existing theory and literature, experience, observations and the mental model of academic and industry experts.

E.5 REPOSITORY OF DYNAMIC GENERIC COMPONENTS FOR SYSTEM OPERATION

This section provides a selection of operations-based dynamic generic components that can be used as the basis for the creation of dynamic STAMP-based risk management models by following the methodology described in this thesis. Critical variables for each component and their normal units, default values, equilibrium values and range are also provided in associated tables.

E.5.1 CONGRESS AND EXECUTIVE COMPONENT

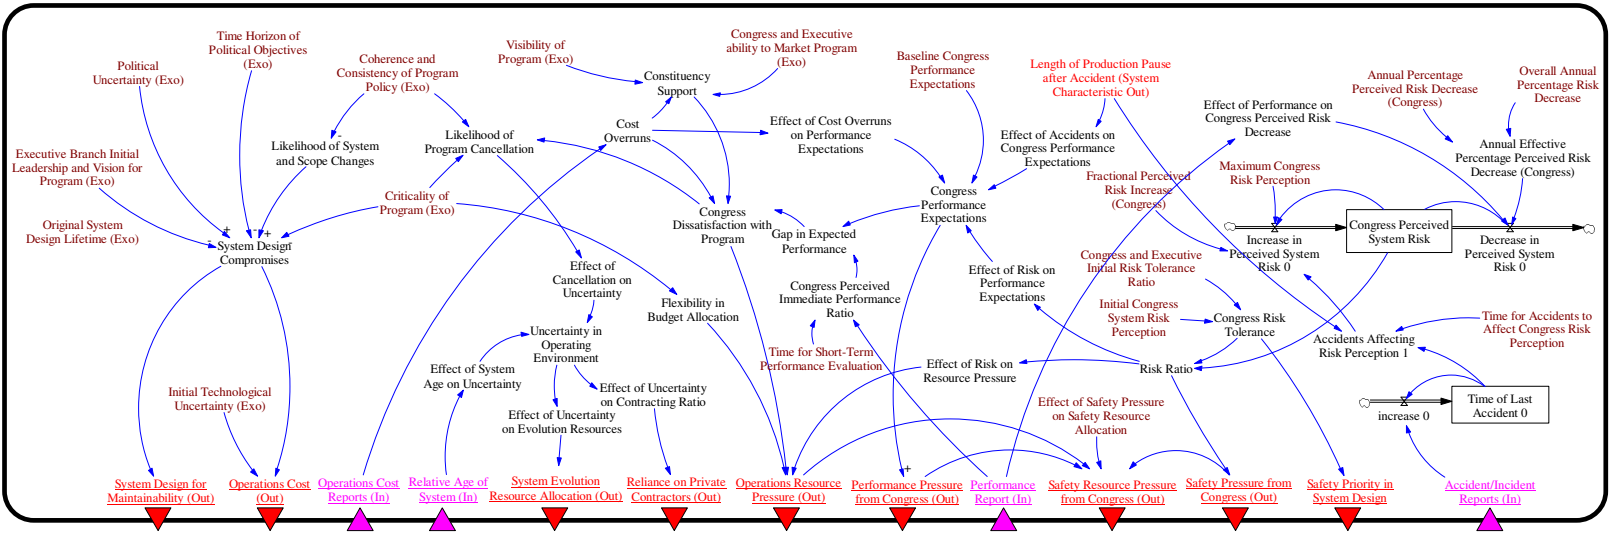


Figure 120: Congress and Executive Component

Congress and Executive						
Component Inputs	Units	Default	Equilibrium	Min	Max	User-Defined
Operations Cost Report	dmnl	1	1	0	2	
Relative System Age	dmnl	1	1	0	5	
Performance Report	dmnl	1	1	0	4	
Accident and Incident Report	dmnl	0	0	0	1	
Component Outputs	Units	Default	Equilibrium	Min	Max	User-Defined
Operations Cost	dmnl	1	1	0	4	
Operations Resource Pressure	dmnl	1	1	0	2	
Performance Pressure from Congress	dmnl	1	1	0	4	
Reliance on Private Contractors	dmnl	1	1	0	2	
Safety Pressure from Congress	dmnl	1	1	0	4	
Safety Priority in System Design	dmnl	1	1	0	2	
Safety Resource Pressure from Congress	dmnl	1	1	0	2	
System Design for Maintainability	dmnl	1	1	0	2	
System Evolution Resource Allocation	dmnl	1	1	0	2	
Exogenous Variables	Units	Default	Equilibrium	Min	Max	User-Defined
Annual Percentage Perceived Risk Decrease	dmnl/year	0.05	0	0	0.5	
Baseline Congress Performance Expectations	dmnl	1	1	0	2	
Coherence and Consistency of Program Policy	dmnl	1	1	0	2	
Congress and Executive Ability to Market Program	dmnl	1	1	0	2	
Congress and Executive Initial Risked Tolerance	dmnl	1	1	0	2	
Criticality of Program	dmnl	1	1	0	2	
Executive Branch Initial Leadership and Vision for Program	dmnl	1	1	0	2	
Initial Congress System Risk Perception	dmnl	1	1	0	2	
Initial Technological Uncertainty	dmnl	1	1	0	2	
Length of Production Pause after Accident	month	6	n/a	0	48	
Original System Design Lifetime	years	20	n/a	0	100	
Political Uncertainty	dmnl	1	1	0	2	
Time for Accidents to Affect Congress Risk Perception	month	3	n/a	0	24	
Time for Short-Term Performance Evaluation	month	3	n/a	0	24	
Time Horizon of Political Objectives	dmnl	1	1	0	2	
Visibility of Program	dmnl	1	1	0	2	
Intermediate Variables	Units	Default	Equilibrium	Min	Max	User-Defined
System Design Compromises	dmnl	1	1	0	2	
Likelihood of System and Scope Changes	dmnl	1	1	0	2	
Likelihood of Program Cancellation	dmnl	1	1	0	2	
Cost Overruns	dmnl	1	1	0	4	
Constituency Support	dmnl	1	1	0	2	
Congress Satisfaction with Program	dmnl	1	1	0	2	
Congress Performance Expectations	dmnl	1	1	0	2	
Congress Perceived System Risk	dmnl	1	1	0	2	
Congress Risk Tolerance	dmnl	1	1	0	2	
Uncertainty in Operating Environment	dmnl	1	1	0	2	

Table 5: Sample Table for Congress and Executive Component Variables

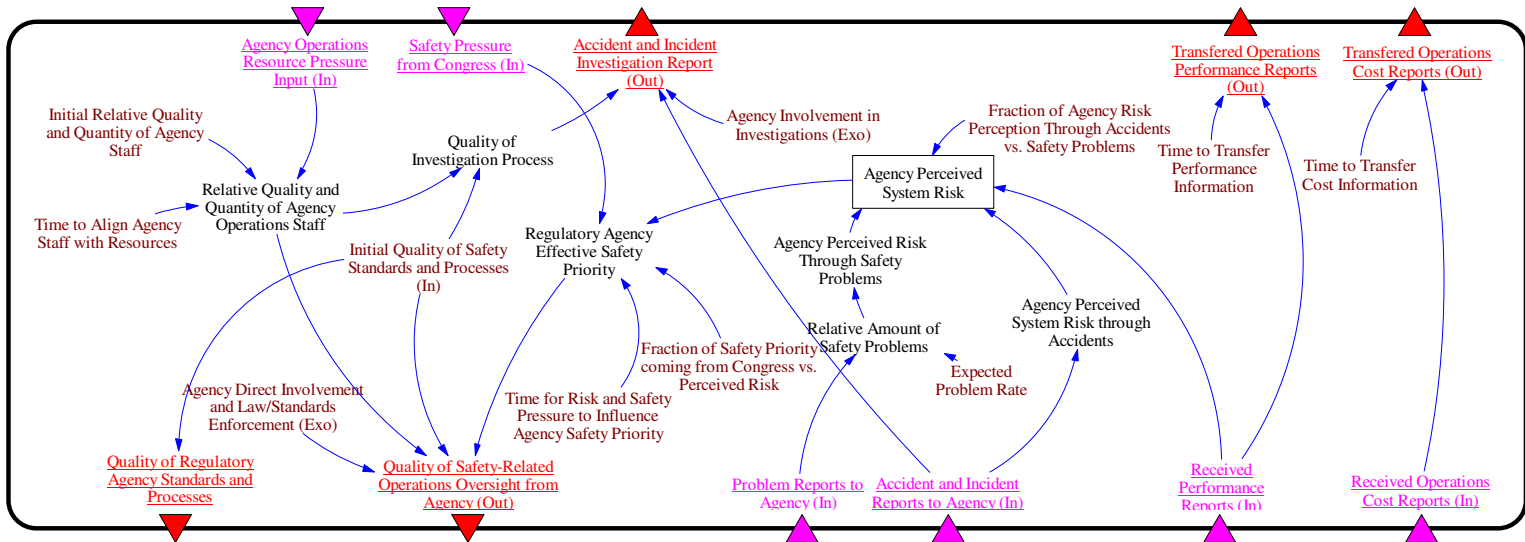


Figure 121: Regulatory Agency Component

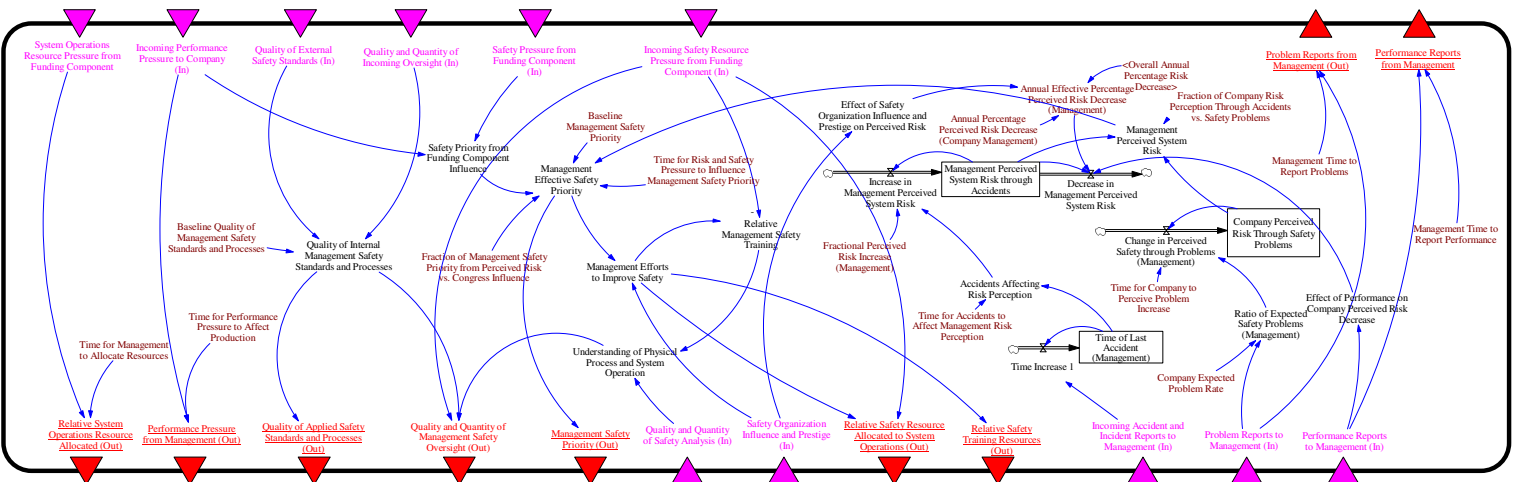


Figure 122: Company Management Component

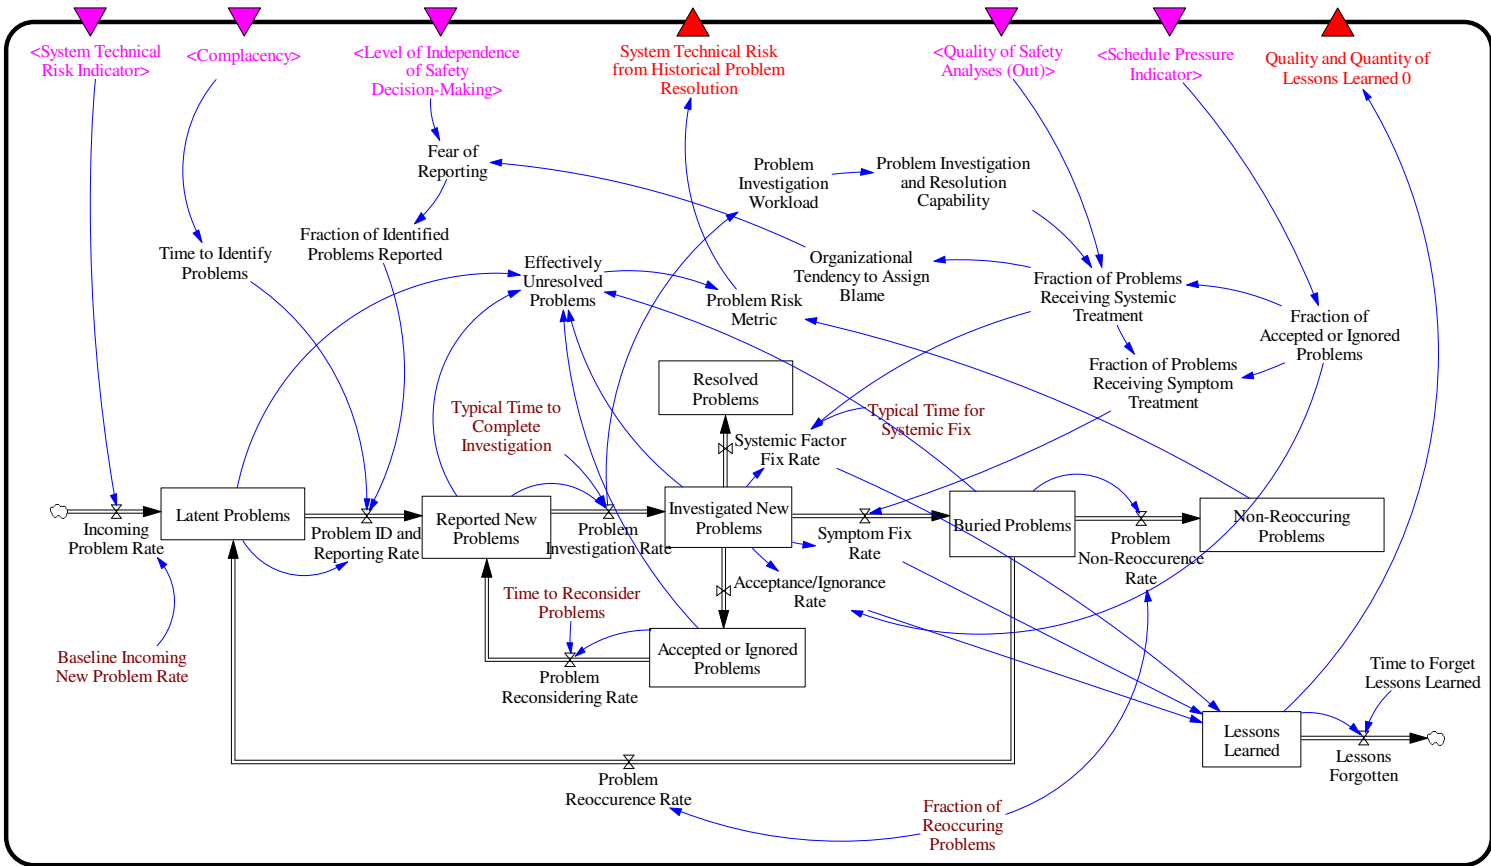


Figure 123: Operations Management Component - Problem Resolution and Learning

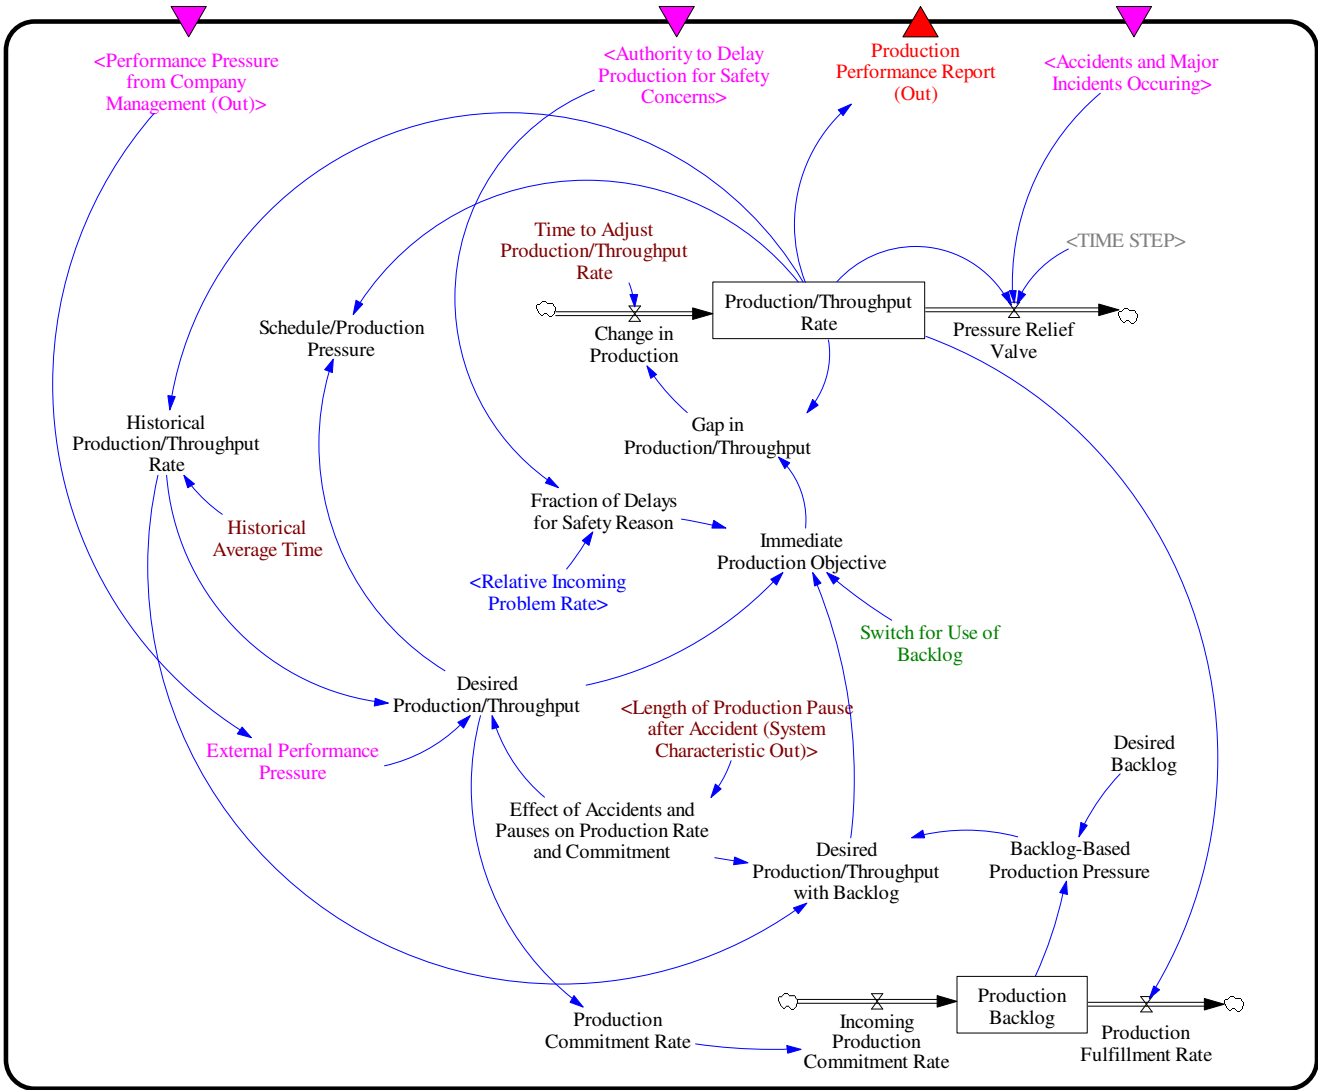


Figure 124: Operations Management Component - Production System

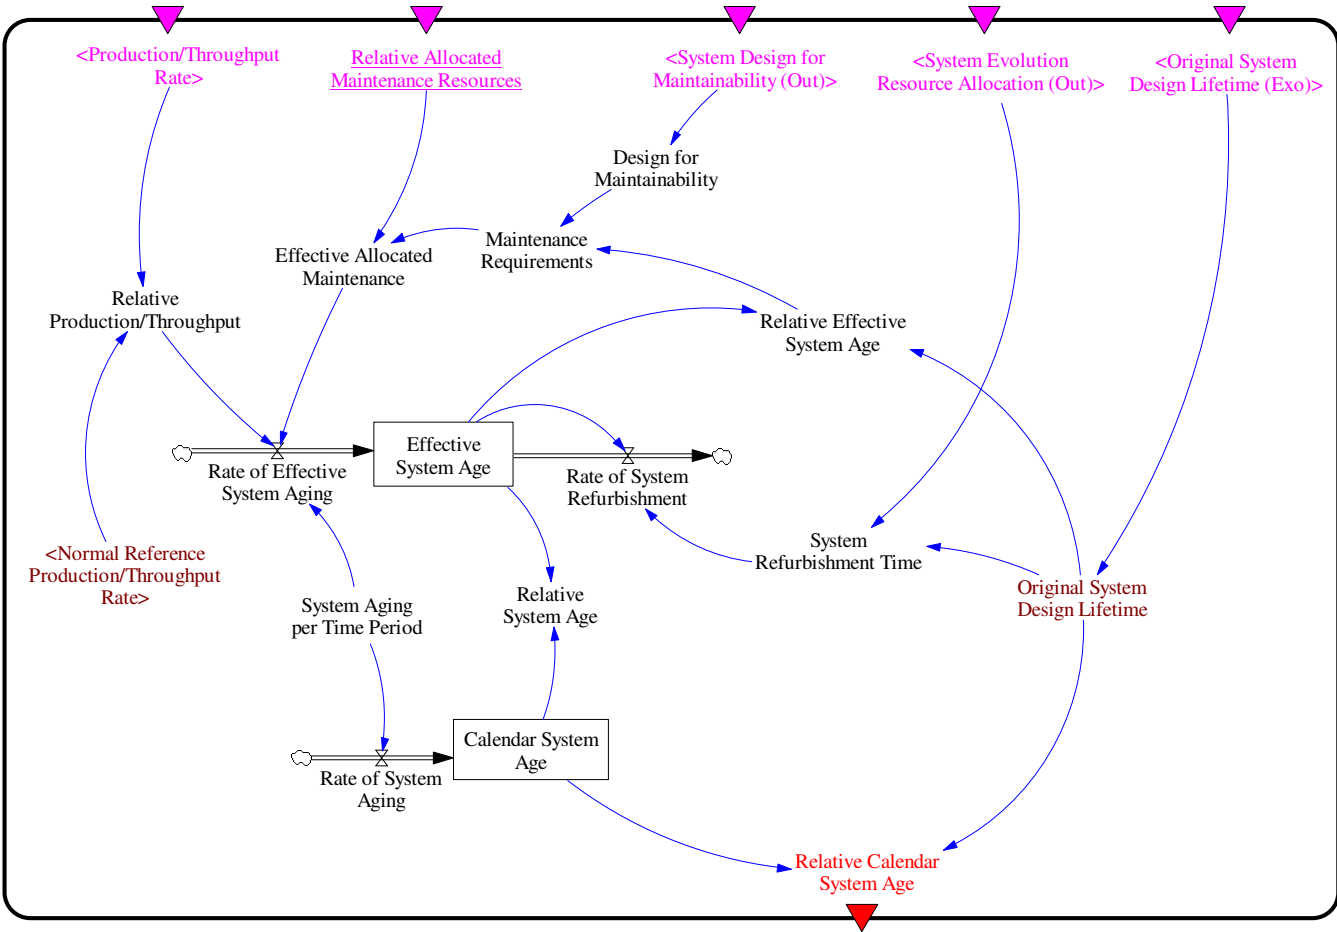


Figure 125: System Maintenance and Evolution Component

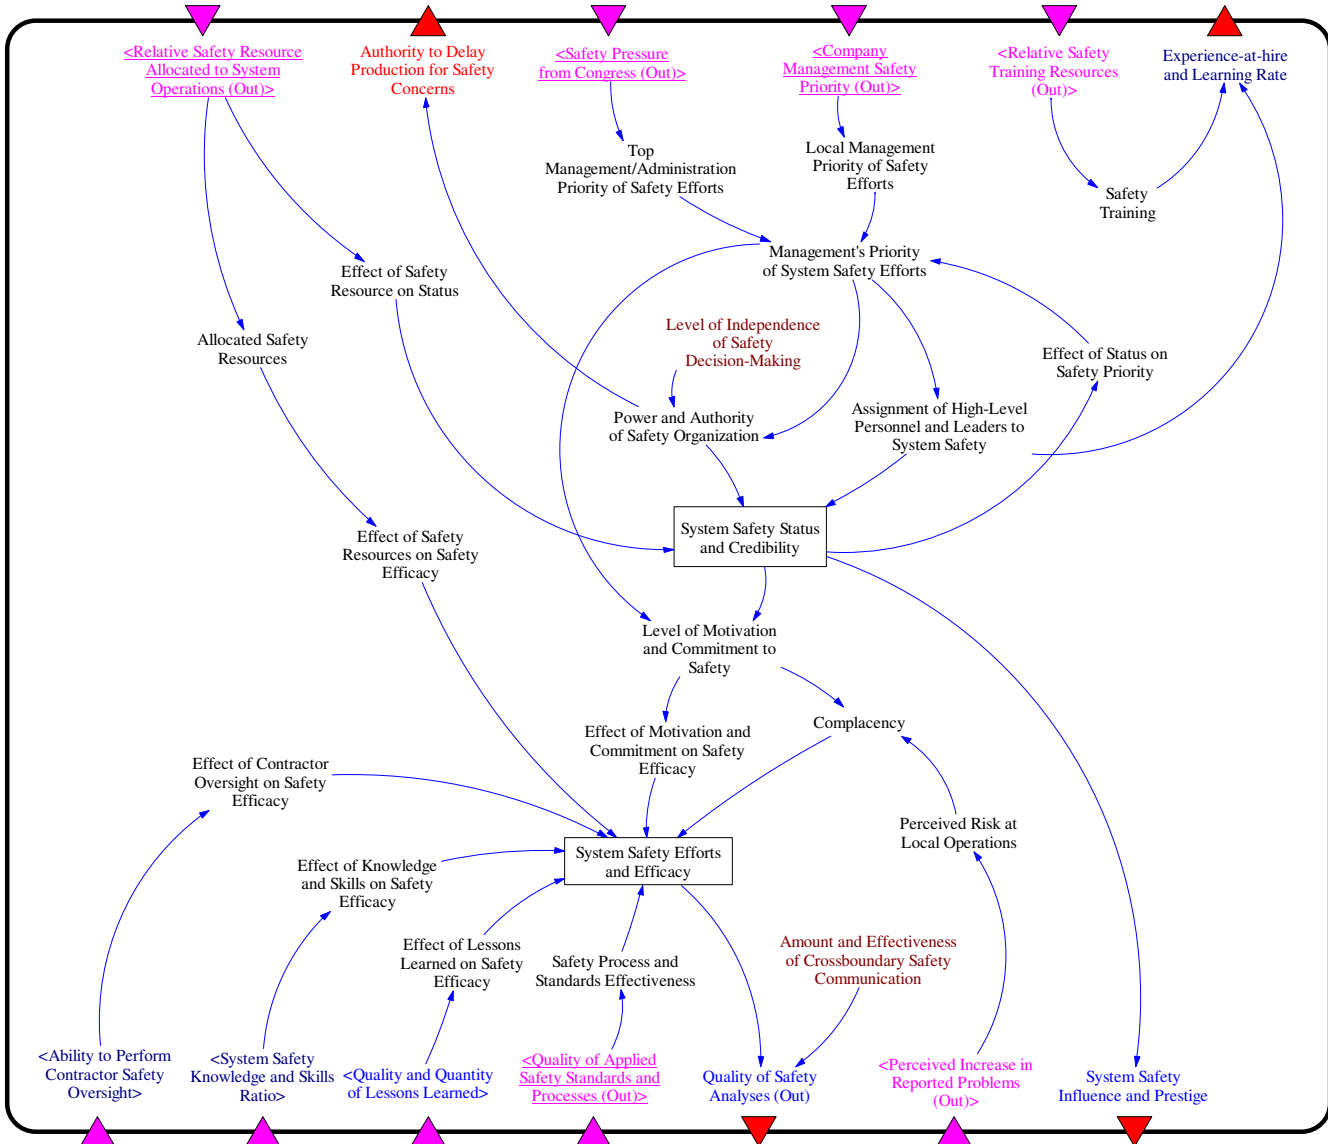


Figure 126: System Safety Component – Status, Effort and Efficacy

E.5.8 SYSTEM SAFETY COMPONENT - KNOWLEDGE, SKILLS AND STAFFING

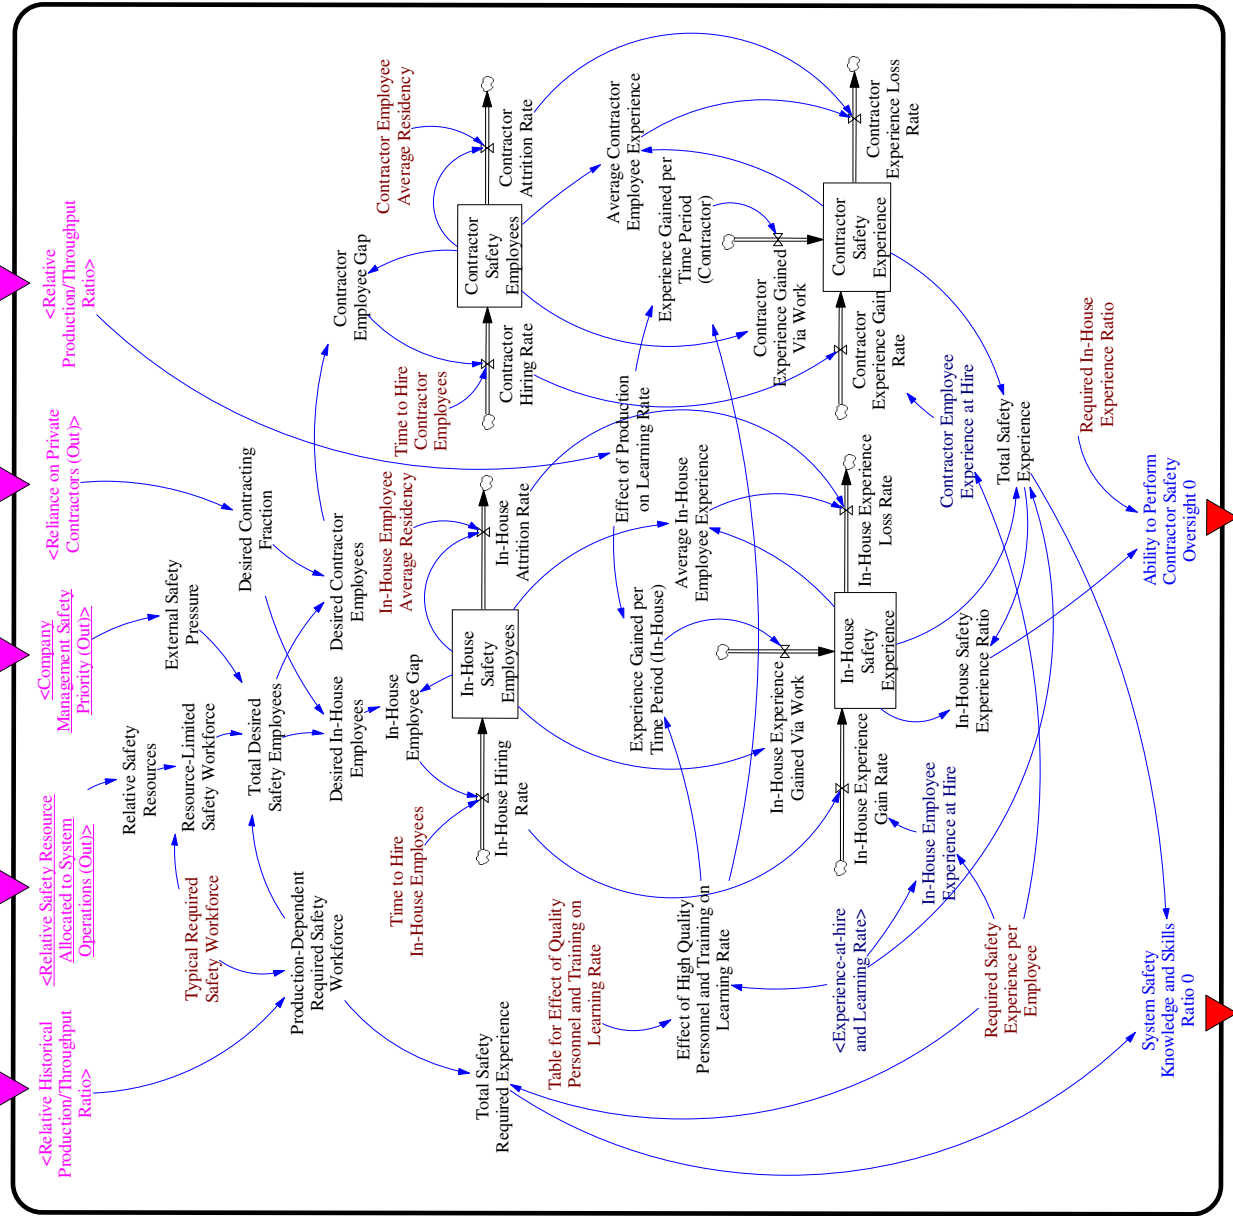


Figure 127: System Safety Component – Knowledge, Skills and Staffing

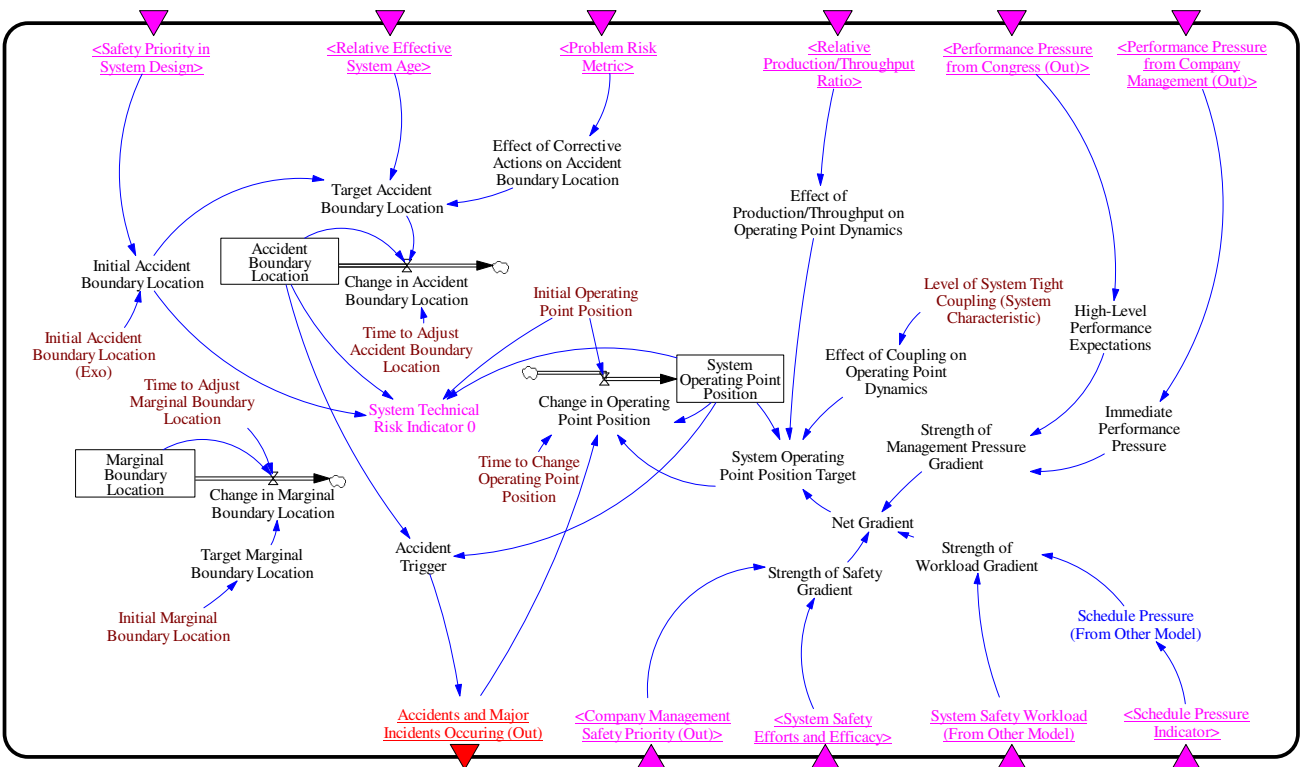


Figure 128: Operating Process Component - System Technical Risk and Safety Dynamics

E.6 REPOSITORY OF GENERIC DYNAMIC COMPONENTS FOR SYSTEM DEVELOPMENT

E.6.1 CONGRESS AND EXECUTIVE (DEVELOPMENT)

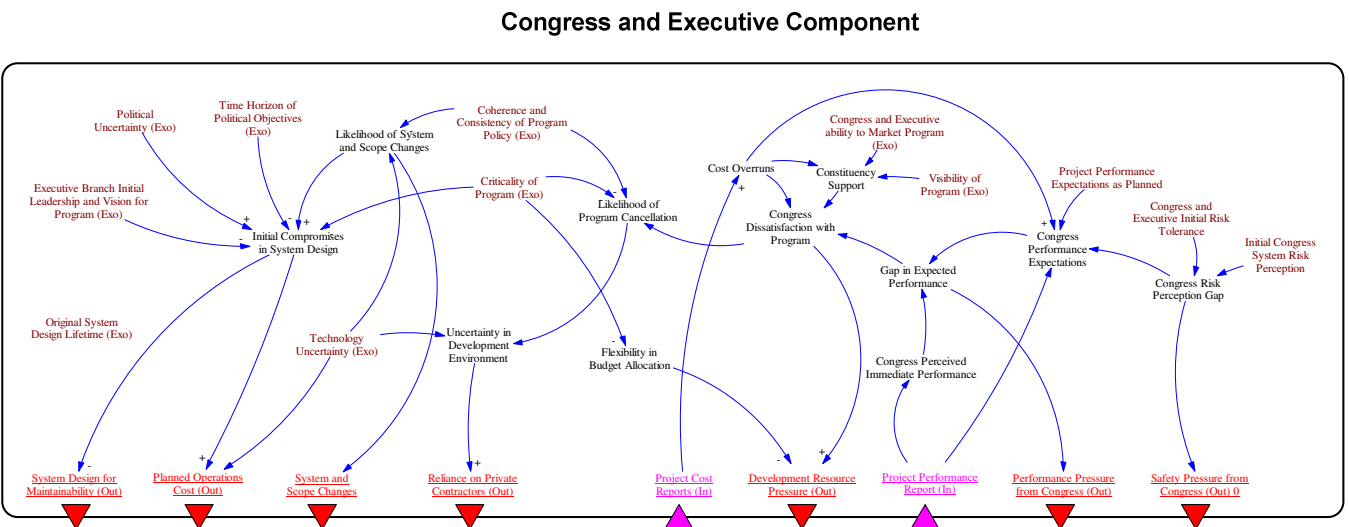


Figure 129: Congress and Executive Component (Development)

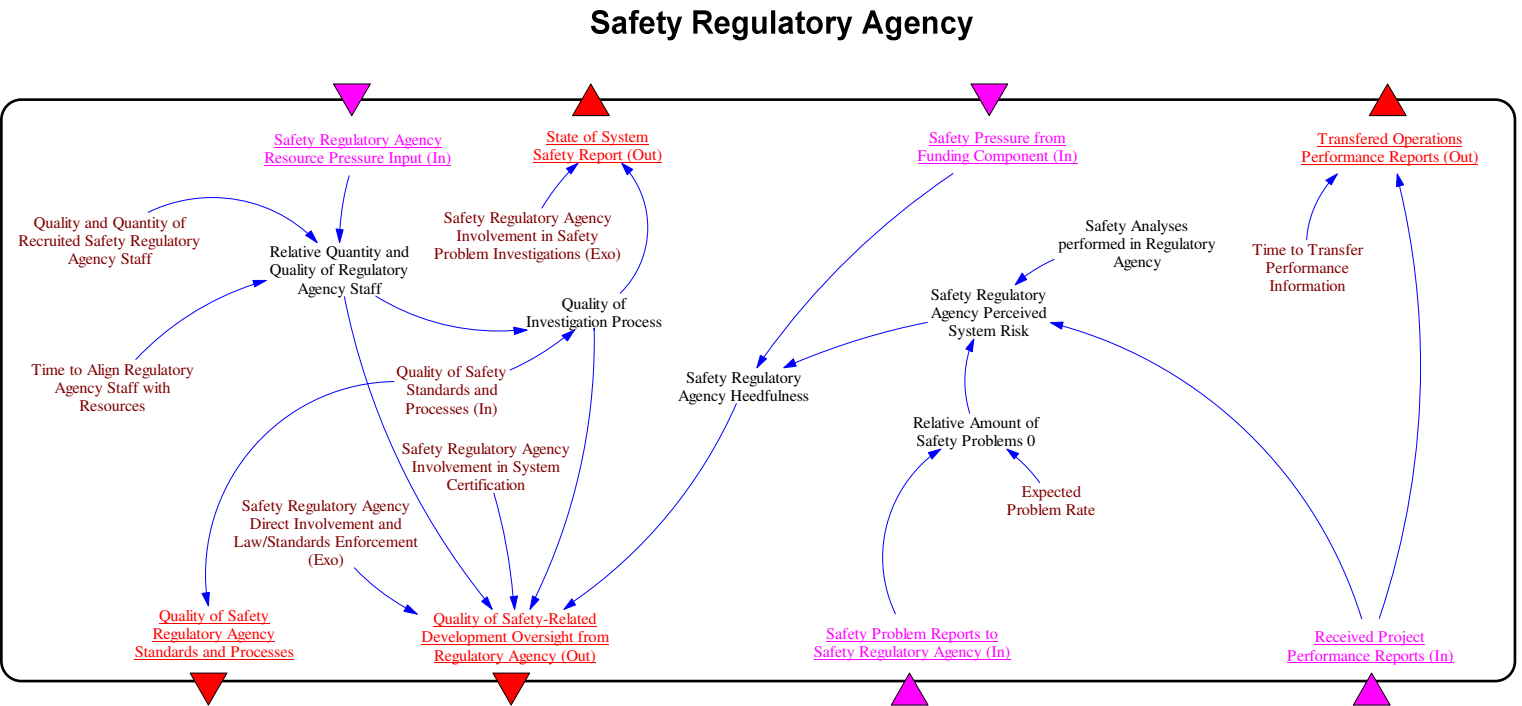


Figure 130: Safety Regulatory Agency (Development)

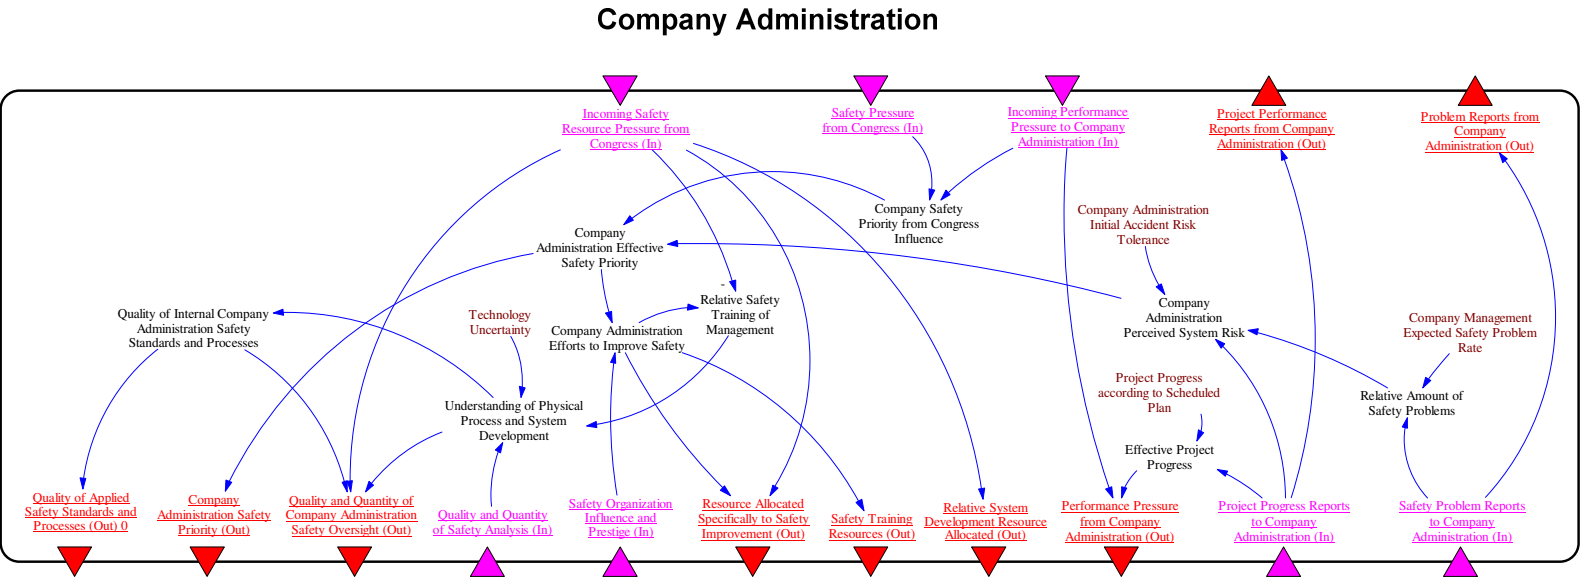


Figure I31: Company Administration (Development)

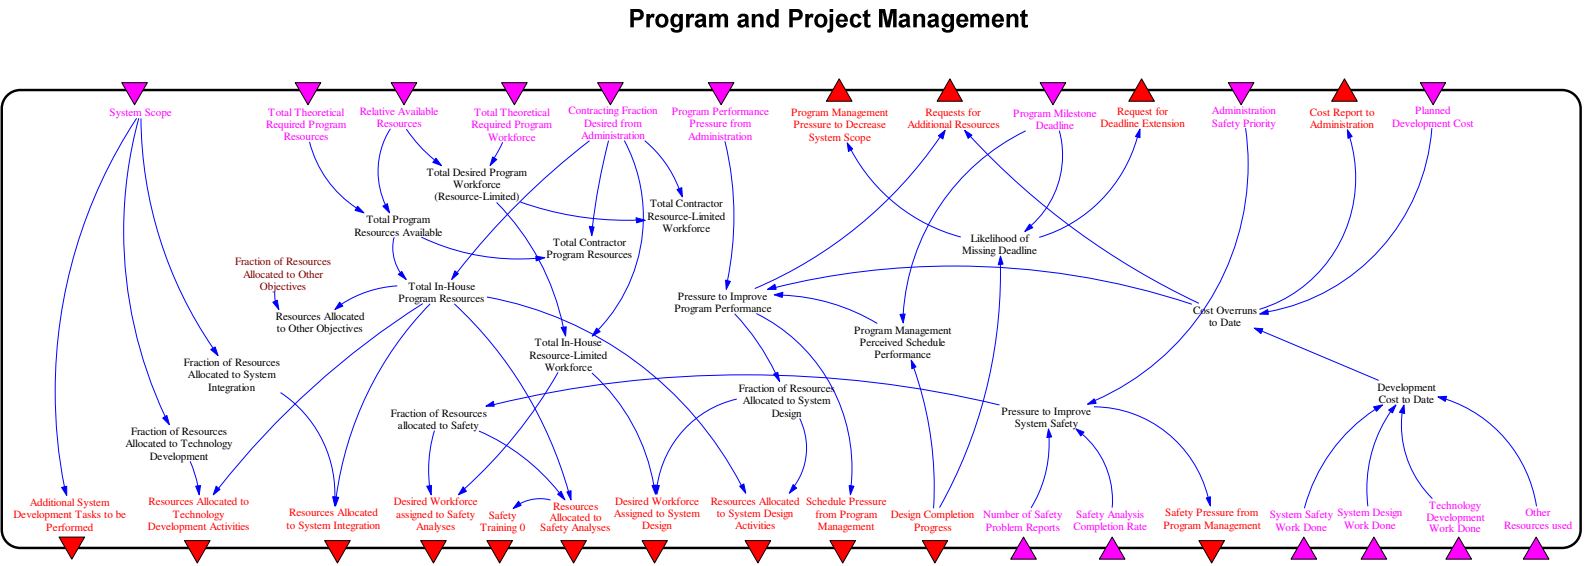


Figure 132: Project and Program Management (Development)

Engineering - System Development Completion and Safety Analyses

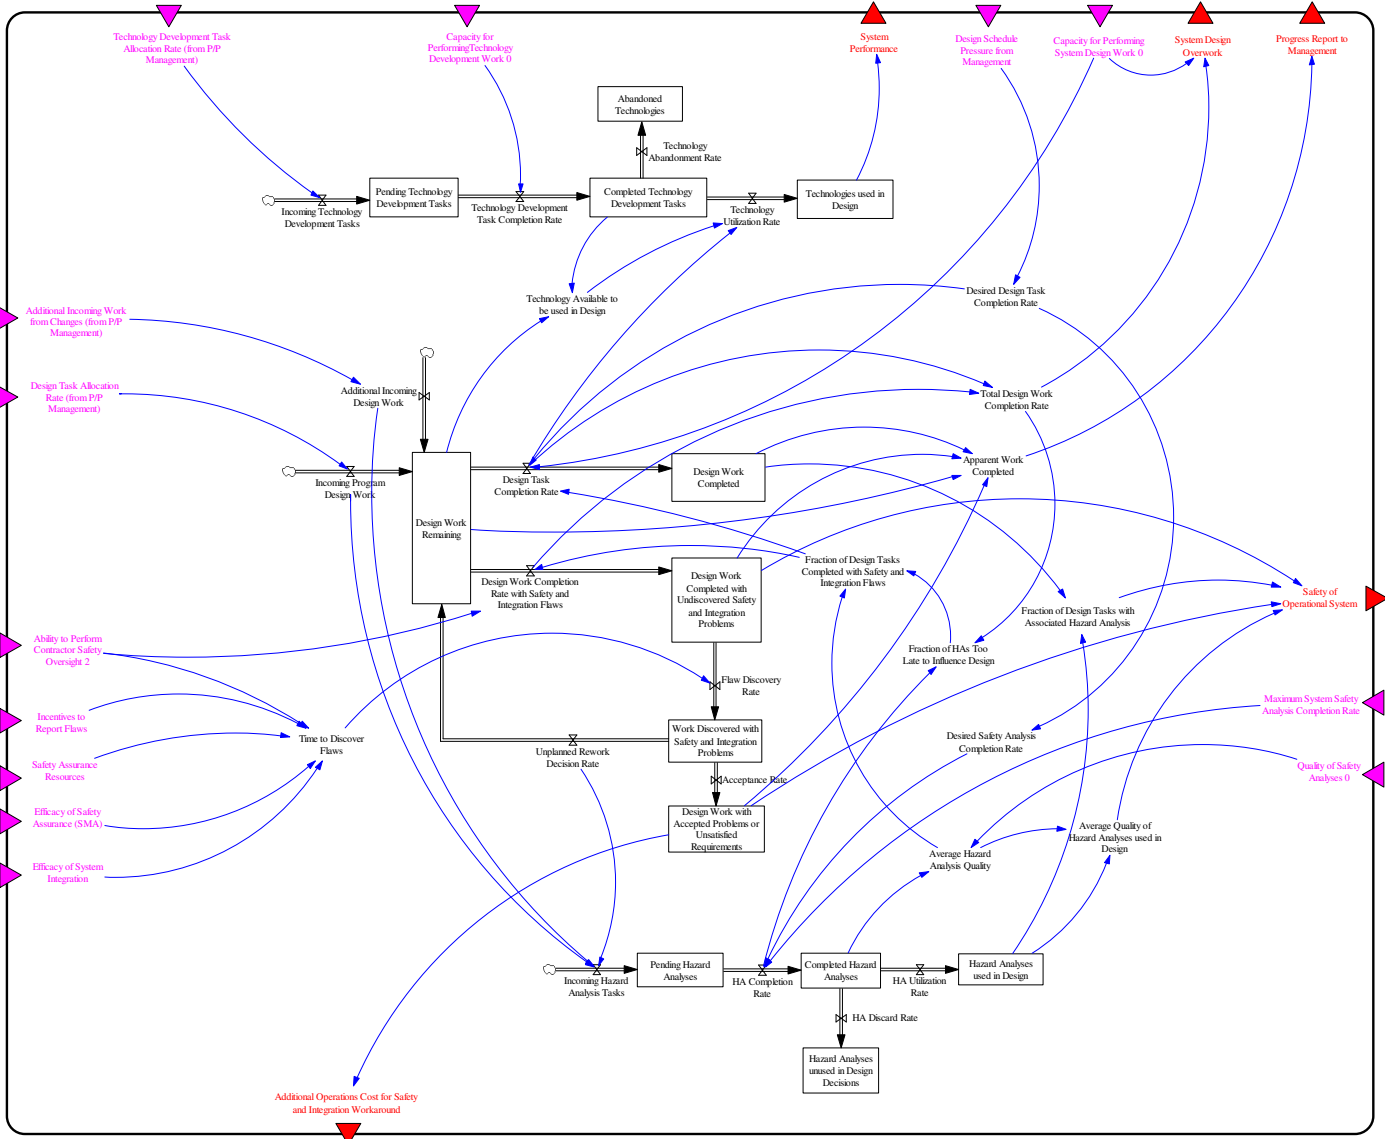


Figure 133: Engineering - System Development Completion and Safety Analyses

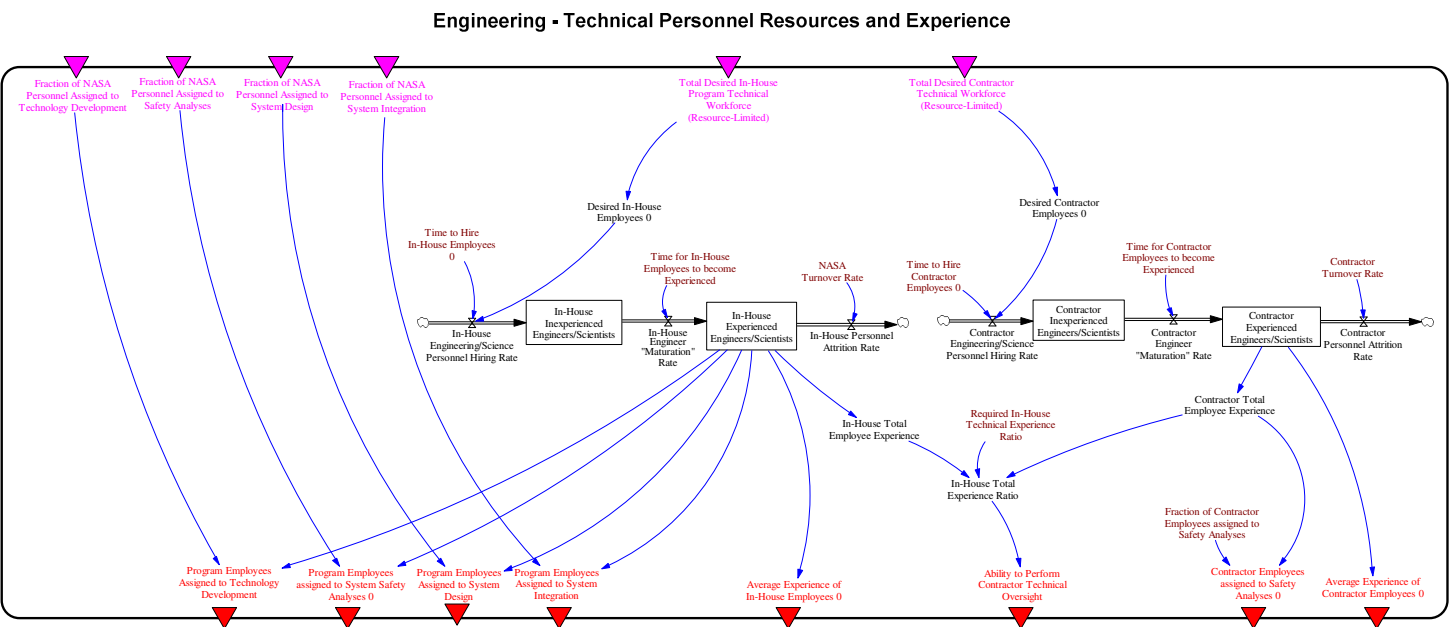


Figure 134: Technical Personnel Resources and Experience (Development)

Engineering - Effort and Efficacy of System Safety Analysts

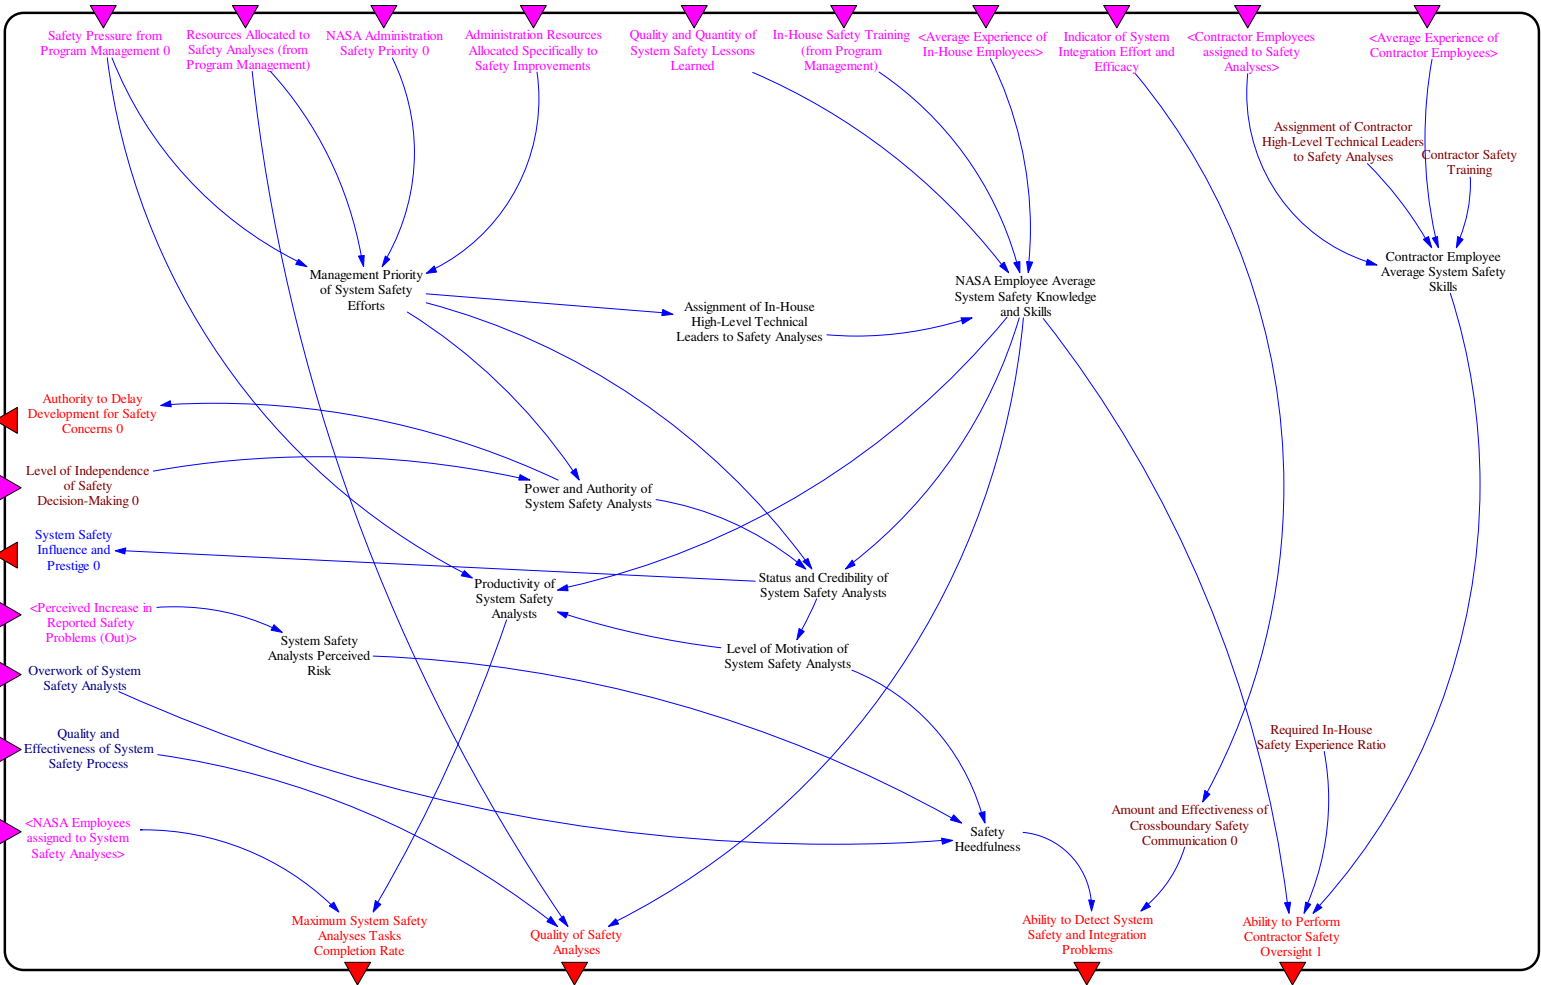


Figure 135: Engineering - Effort and Efficacy of System Safety Analysts

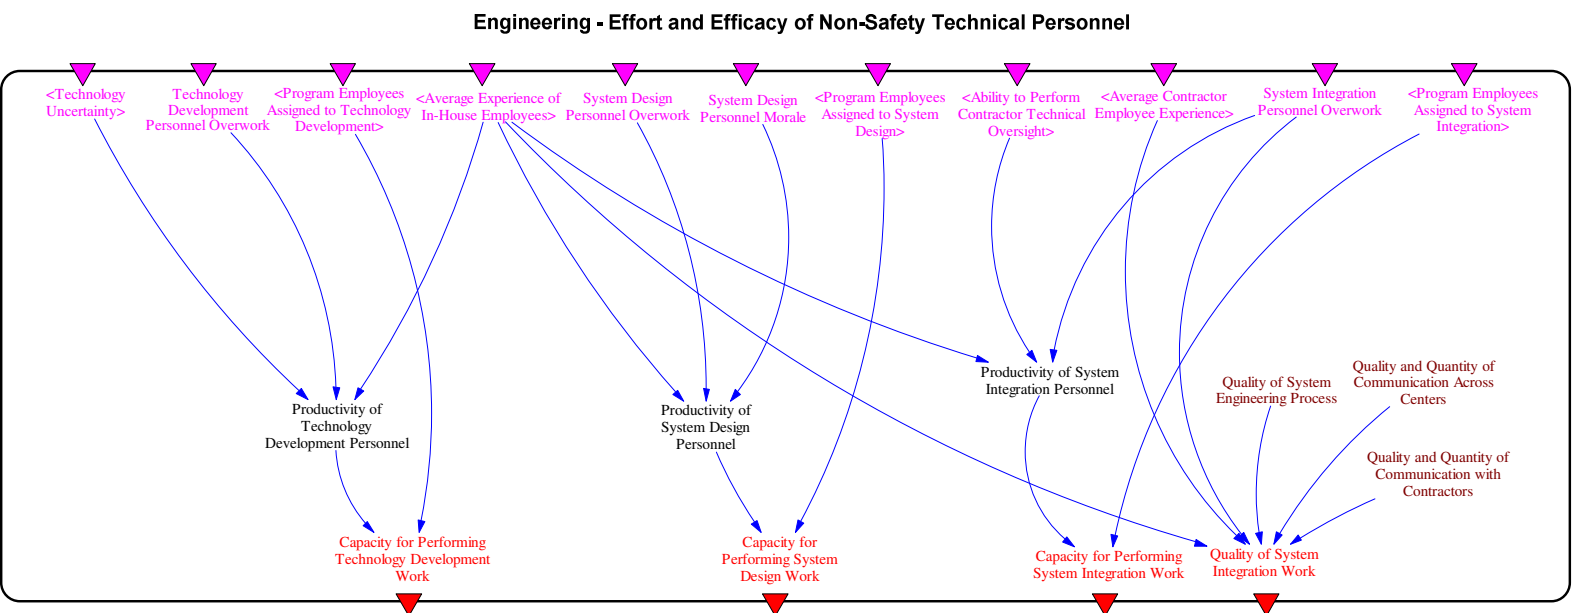


Figure 136: Engineering - Efforts and Efficacy of Non-Safety Technical Personnel (Development)

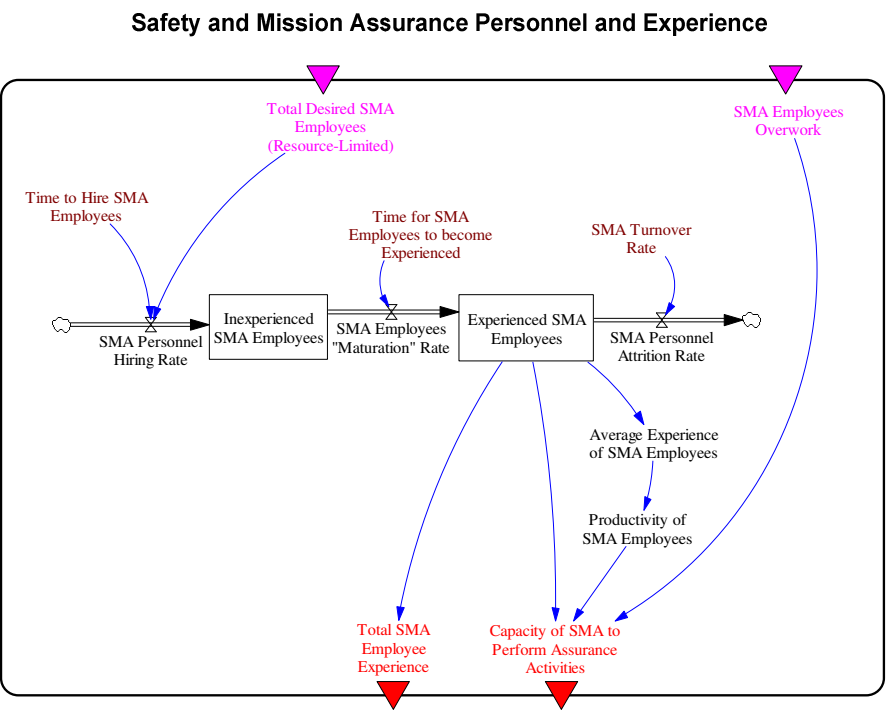


Figure 137 : Safety and Mission Assurance Personnel and Experience

Introduction to Study*(Consent Form)*

Thank you for participating in our study.

Our objective is to create and validate a description of how safety-related decision-making is performed during the development of NASA's Space Exploration System. Once we have a better understanding of the mechanisms behind safety decision-making, a simulation model will be created to capture this dynamics. This model will be used to test new safety/risk management policies and improve safety monitoring during the development of the new space exploration system.

You were selected because you have the expertise necessary to help us improve, modify and validate the causal structure of safety decision-making. There is absolutely no right or wrong answer. We are interested in your perspective on the factors you think contribute to safety-related decision-making. Your participation, while greatly appreciated, is entirely voluntary and the information you share with us will remain confidential.

If you agree to it, we would like to record this conversation for further data analysis. This interview will be more of an interactive review of the model rather than a Q&A session and the recording will allow us to concentrate on the interactive process while ensuring we do not miss important bits of information. However, the recording is voluntary and you may choose to opt-out completely or at any point.

Initial Orientation Question*(NASA Chart)
(Safety Responsibilities)*

Please review this chart of the NASA structure for the development of the space exploration system.

1. Where does your position fit in this structure and what kind of expertise do you bring?
2. Focusing on your area of expertise, how would you describe the flow of resources and information across participants (in boxes) in the system (examples)?
3. How do you describe the role you play in safety during the development of the space exploration system?

Please review the list of responsibilities associated with ensuring the development of a safe exploration system.

4. If you were to design the system, how would you distribute the responsibilities?

Safety-Risk Question

(Safety Risk Chart)

5. What do you think are the most important **factors** (3-5) that affect safety in the development of NASA's space exploration system?

Can you characterize or discuss the relationship between these factors and safety?

Model Analysis Question

*(Large Chart of Model Component)
(Large Chart of Second Model Component)*

6. Please review the large chart that corresponds to a component of the model. This represents our current understanding of the safety-related decision-making taking place in this area. We will step through it and ask for your comments and ideas on how to better represent and capture the causal structure that leads to safety decision-making in the development of the space exploration system. We hope you will use the chart to record comments, make modifications, and provide explanations. Please do not hesitate to let us know if you think the proposed structure is wrong or inaccurate.
7. We will now use the remaining time to go over a component of the model that describes the system development progression.
8. Are there any other things you would like to tell us that we did not discuss or cover in the diagrams?

Again, thank you for your help.

Following the Columbia accident in 2003, an interdisciplinary complex systems safety working group was created at MIT. One of the objectives of the group is to further our understanding of the dynamic processes that led to the flawed safety-related decision-making at the origin of the Columbia disaster. Our first attempt at integrating technical, organizational, and cultural factors into a comprehensive STAMP-based modeling framework were promising. The NASA Office of the Chief Engineer (OCE) asked MIT CSRL (Complex Systems Research Lab) to perform a risk analysis of NASA's planned implementation of an Independent Technical Authority, as mandated by the CAIB [Gehman, 2003] as part of the Space Shuttle return-to-flight effort following the Columbia Accident. The goal of this study, conducted in the summer of 2005, was to identify and evaluate the risks associated with this implementation of the ITA within the existing NASA structure and processes. Some of the results of the ITA study were integrated in previous chapters of this thesis to illustrate the piecewise application of the methodology.

Following the CSRL ITA analysis report, interest grew in the models and methodology. In the spring of 2006, NASA ESMD contracted with MIT CSRL to demonstrate the application of this new risk analysis and management approach to the ESMD space exploration enterprise. The statement of work provided by ESMD required the development of an ESMD model based on previous work with the NASA Office of the Chief Engineer model for the Independent Technical Authority. This model shall consider the work done by the CAIB regarding influences and effects on Systems Risks and Safety.” The model(s) shall include, as necessary, organizational entities including Congress, regulatory agencies, industry associations, unions, courts, NASA HQ, ESMD, SOMD, Office of Safety and Mission Assurance, Constellation Program, Shuttle and Station Programs, Johnson/Marshall/Kennedy Space Centers, CEV, CLV, RLEP, and Advanced Technology. The model(s) shall be iterated during development with NASA personnel to ensure their accuracy. Final deliverables include a formal report and a presentation for the Associate Administrator for ESMD.

One major difference between the ITA Project and the ESMD project is that the shuttle program is basically an operational program while ESMD will be focused for many years on development rather than operations. The models deal with completely different lifecycle phases and the risks and concerns for each phase are vastly different. Nonetheless, the methodology introduced in this thesis is mostly independent from the lifecycle phase.

H.1 INTERVIEWS WITH NASA OFFICIALS

Interviews with NASA officials were used as a primary source of information throughout the study for a many reasons. First, each organization is unique and many elements of its functions are unwritten or even guarded closely at times. The amount of information that can be gathered through a review of internal documentation and the broader literature in the public domain is limited even under normal circumstances. At the time of this study the scarcity of such documentation was exacerbated by the fact that NASA was in the midst of a dramatic reorganization in response to the Columbia Accident, the appointment of a new Administrator, and the implementation of the new Vision for Space Exploration.

Moreover, the model-building methodology described in this thesis relies heavily upon the customization of generic structures through interactions with domain experts at various levels and areas of a complex system. The objective of these interactions is to customize the dynamic structure that influence decision-making in various components of the organization. Typically, information about the structure, heuristics and criteria used for daily decision-making can only be obtained through discussions with employees and stakeholders highly familiar with the organization under study.

In all, 44 people were interviewed during 41 interviews conducted over a three-month period at NASA Headquarters, the Marshall Space Flight Center, the Johnson Space Center, and the Langley Research Center. Many of the interviewees worked in the Exploration Systems Mission Directorate, but a number worked above the directorate-level and still others worked in other mission directorates, particularly the Space Operations Mission Directorate (SOMD). Overall, interviewees included representatives from the Office of the Administrator, the Office of the Chief Engineer, the Office of Safety and Mission Assurance, the NASA Engineering and Safety Center, the Office of Program Analysis and Evaluation, Mission Directorate Offices, ESMD Directorate Offices, Program Offices, Project Offices, the Office

of Institutions and Management, Center Safety and Mission Assurance Directorates, Center Engineering Directorates, Center Mission Operations Directorates, and the Astronaut Office.

Each interview was recorded and over 200 pages of interview transcripts were accumulated. These transcripts were used extensively for model creation and validation, but would not be released as is because of confidentiality reasons.

A research process was developed that allowed the collection of interview data, the development of the model, and the calibration and debugging of the model in a concurrent fashion. The model-building process used in the ESMD study follows the methodology of this thesis. It uses a combination of causal loop diagramming, formal simulation model building supplemented by the use of NASA documents including quantitative data sources, literature and accident reports. The interview protocol selected is semi-structured, using a set of questions for consistency, while leaving interviewees with the freedom to elaborate on any topic. This technique allows the interviewee to provide information that the interviewer might have unintentionally suppressed with a more rigid set of questions.

The standard duration for an interview session was one hour, though some were as short as approximately twenty minutes or as long as approximately one hundred minutes. At the beginning of each session, the interviewee was briefed on system dynamics and STAMP from a set of introductory slides and asked to both sign a consent form and give verbal consent to the voice recording of the session (The interview protocol is provided in Appendix F). The sessions consisted of four major sections based on time constraints and interviewee expertise: 1) Discussion of the NASA ESMD structure, 2) Responsibilities survey, 3) Safety risk survey, and 4) Review of a specific model component. Each of these portions of the interviews is described in the subsections below.

H.1.1 Discussion of the NASA ESMD Structure

The structure of the NASA organization (with a particular focus on ESMD) was briefly discussed with each interviewee in order to answer structural questions, identify the domain expertise of interviewees to better focus the interview, and to identify organizational issues and discrepancies the roles of the organization elements. The chart in Figure 138 was presented to each interviewee to identify which component better fits their expertise. The

interviewees were then asked to describe how resources and information flowed across the elements of the organizational structure. Finally, interviewees were asked to describe their role in safety during the development of the space exploration system.

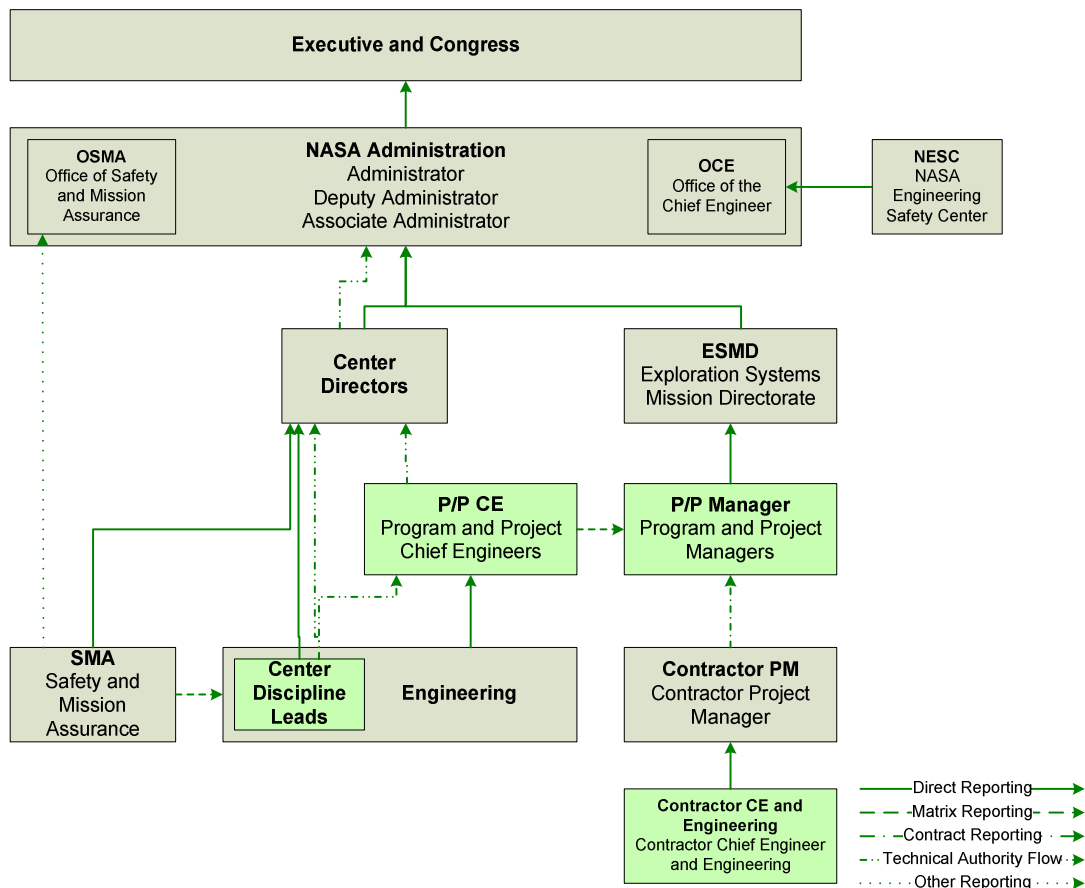


Figure 138: Exploration-Centric NASA Structure

H.1.2 Safety Responsibilities

Whenever time and interviewee expertise permitted, interviewees were asked their opinion on the allocation of selected safety-related responsibilities. This interview section was derived from previous work on the ITA model. One of the objectives was to further improve confidence in the safety-related processes included within organization components. Furthermore, the data collection allows the identification of disagreements and potential areas of inadequate control due to overlap or gaps in control responsibilities. Ultimately, it

appeared that the discussions generated by the questions were more informative to the modeling process than the responses themselves.

H.1.3 Safety Risk Identification

Schedule and budget restrictions for the ESMD study limited the detailed completion of the static STAMP control structure analysis including component responsibilities and thus the generation of a complete list of project risks using the technique summarized in section 5.2.1. As the usefulness of the technique had been previously demonstrated during the ITA project (see [Leveson, 2005] for complete results), the ESMD project used a more classic (and less complete) identification process involving interviews of experts. In this interview section, participants were asked to list what they feel to be the three to five most important factors affecting safety-related risk in the development of the new space exploration system. The questions were used to identify risks deemed to be most detrimental to safety if not addressed properly. The responses were used as the basis for the creation of risk analysis scenarios, as documented further in this chapter. Additionally, the questions were important to the modeling effort in order to ensure that some of the model structure addresses these risk factors.

H.1.4 Review of Model Components

The last interview section was used to review the causal structure of a component of the model that was closely related to the interviewer's area of expertise. The components were printed on large paper sheets (poster size) and placed on a table in front of interviewees. During the review, interviewees were guided through the draft component structure, first through a high-level component description, then one variable and relationship at a time. Questions were asked on specific relationships that we felt were more uncertain. As much as possible, interviewees were asked to try to quantify relationships between variables. The review process was highly interactive, with modifications and notes written directly on the large component sheet. Usually, after the first few questions were asked, interviewees were comfortable enough to review other relationships and variables in the model, provide comments on them without further prompting, and suggest changes or additions to the model based on their knowledge and experience. In all, the interviewees in 38 of the 41 interviews

reviewed model components and both the qualitative and quantitative inputs that they provided led to a number of changes in the model’s causal structure and underlying equations. Figure 139 provides a summary of the number of interviewees who participated in the review of each component.

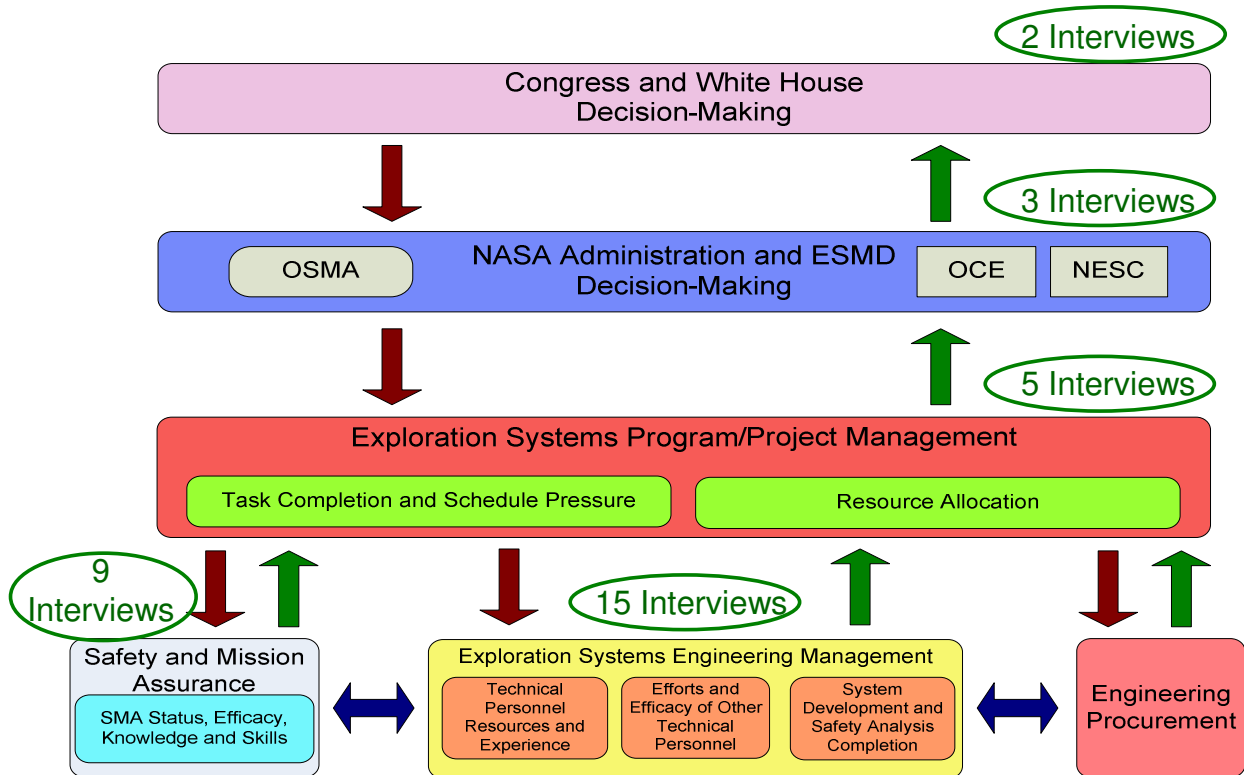


Figure 139: Model components and interviews

H.2 ADDITIONAL DATA SOURCES

Interviews served as the primary source of qualitative data. Some interviewees provided quantitative data, but as the vast majority of detailed budget and scheduling data was content-sensitive, most of the quantitative data was obtained from a number of sources in the public domain. Of these, the most noteworthy source of data was the workforce data cubes on the NASA People Website [NASA, 2006] (see Figure 140). This data source is highly interactive and allows for the customization of data reports and the extraction of very specific data about the NASA civil servant workforce (e.g. the resignation rate of NASA civil servants with less

than nine years of experience). Table 6 provides a summary of the quantitative data used as a foundation for the workforce planning and budget estimation algorithms in the model.

SOURCE	TYPES OF DATA
Workforce Data Cubes on the NASA People Website [NASA, 2006]	<ul style="list-style-type: none"> • Center support contractor headcounts for FY 2002 • Headcounts of civil servant workforce in Science and Engineering (S&E) positions • S&E civil servant workforce age, experience, hiring counts, attrition counts, retirement eligibility • Age of civil servant new hires • Etc.
FY 2004 to FY 2007 NASA Budget Requests [NASA, 2004]	<ul style="list-style-type: none"> • Budget breakdowns to the program level (historical and forecast) for FY 2002 to FY 2011 • Estimates of civil servant unfunded capacity
FY 2002 to FY 2004 NASA Procurement Reports [MSFC, 2002]	<ul style="list-style-type: none"> • Total procurement dollars for FY 2002 to FY 2004 • Procurement Awards by type of effort for FY 2002 to FY 2004
Columbia Accident Investigation Board Report [Gehman, 2003]	<ul style="list-style-type: none"> • Space Shuttle Program civil servant and support contractor workforce for FY 1993 to FY 2002 • Space Shuttle Program budget for FY 1993 to FY 2002

Table 6. Budget and personnel data sources and types of data used in the model

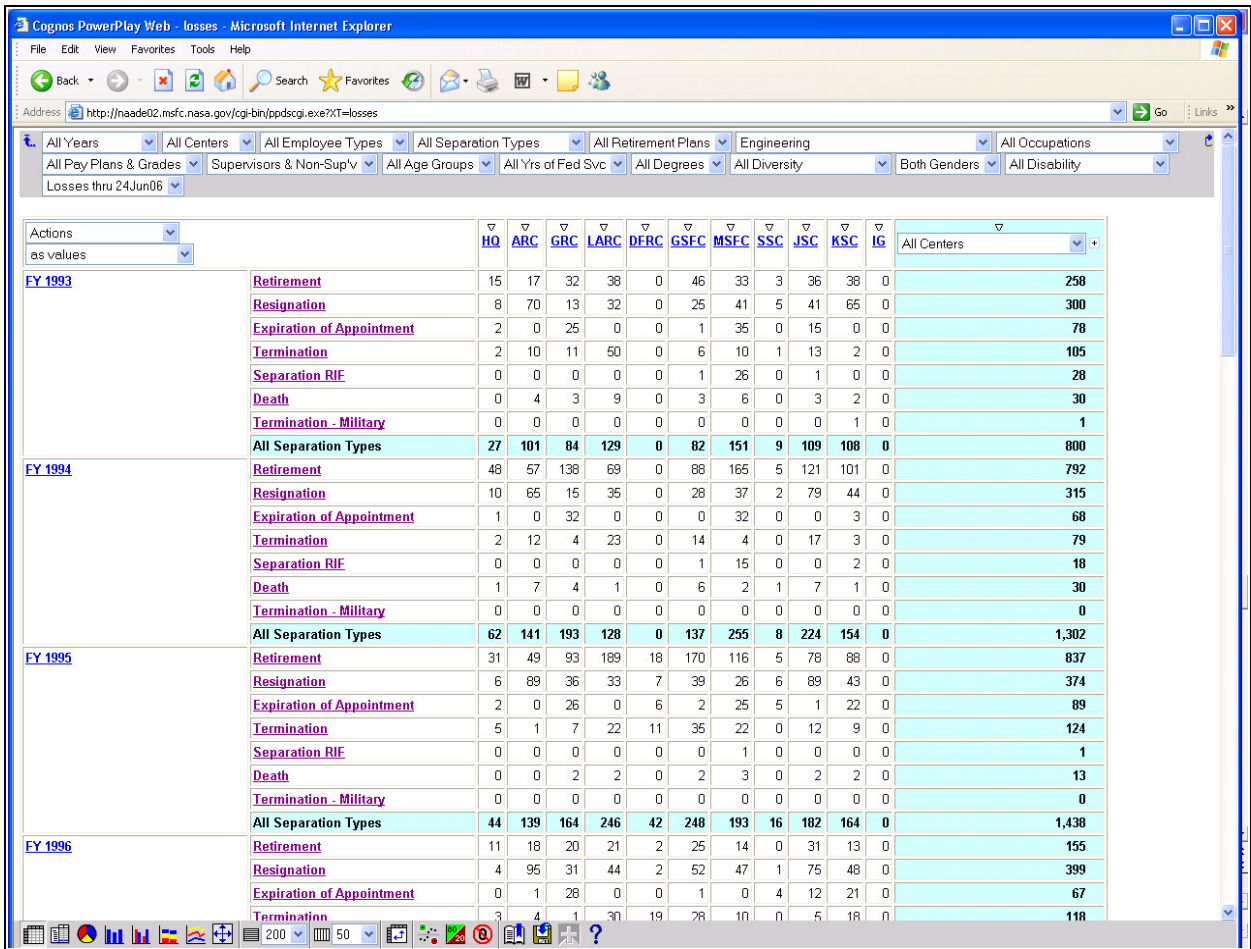


Figure 140: Datacubes on the NASA people website

APPENDIX I: MODEL CONVENTIONS AND ASSUMPTIONS

The assumptions and conventions used in the model are as follows:

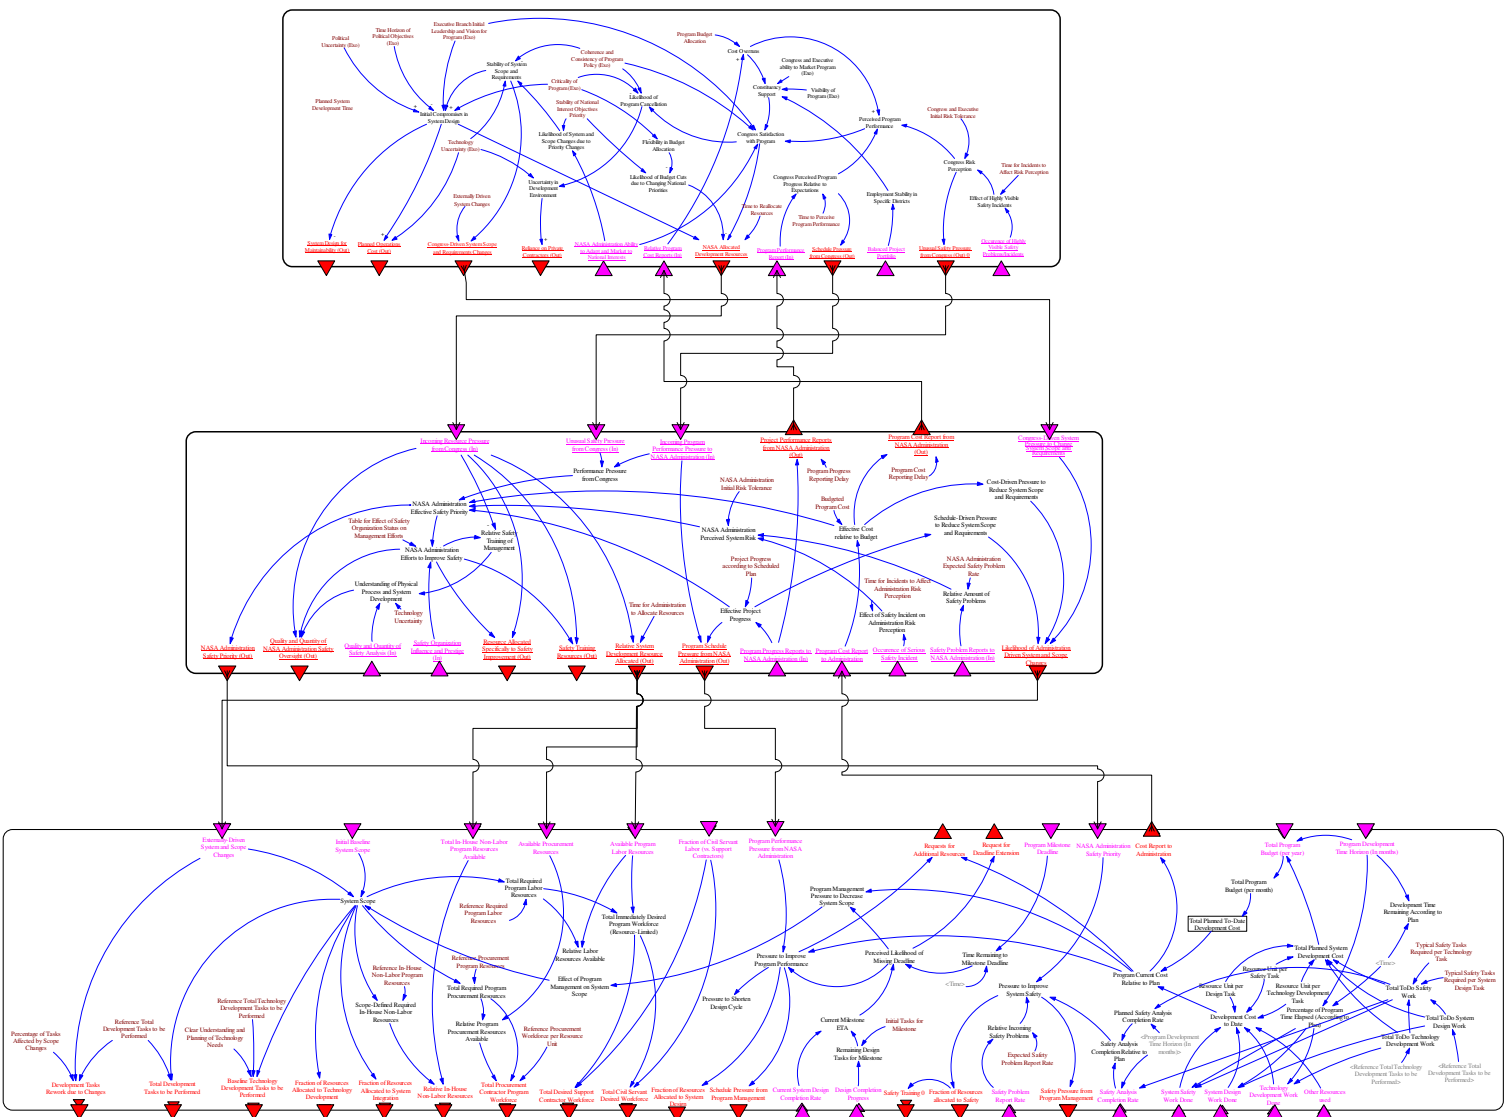
- The start date of the simulation is January 1, 2004
- The end date of the simulation is July 1, 2016
- The workforce and budget in the model is that of the entire Exploration Systems Mission Directorate
- Inflation is not accounted for in the model
- Civil servants do not get fired from ESMD, they get transferred out when there is a surplus of ESMD civil servants
- The planned initial system development time (CEV, CLV, LPRP) is 8 years
 - Only the development of the major projects at the start of the simulation are tracked
 - Projects that start up after the start of the simulation can affect the progress of the tracked projects (e.g. CaLV)
- The explicit program management schedule and budget reserves are set to zero in the baseline model. The impact of non-zero reserves is discussed further in specific scenarios
- The budget for ESMD is exogenous in the baseline model
 - There is a switch in the model for turning this assumption off
 - Between 2004 and 2011, the budget is based on a curve fit of budget request forecasts
 - Beyond 2011, the budget peaks in 2012 and decreases back to \$6 Billion in order to approximate the effect of CEV/CLV ops deployment
 - Between 2010 and 2011 roughly \$4 Billion will be transferred to ESMD from SOMD for Shuttle retirement, raising the ESMD budget to roughly \$8.8 Billion. The assumption is that more than half of the \$4 Billion ESMD received from SOMD would go back to SOMD for CEV/CLV deployment and thus the budget was rounded down to \$6 Billion
- The scope of the projects tracked in the model is fixed in the baseline scenario
 - There is a switch in the model for turning this assumption off
- Whenever the term procurement is used in the context of the model, it refers to hardware/software acquired from prime contractors. It does not apply to support contractors who work on-site at the NASA centers
 - Support contractors are considered to be a part of the in-house workforce and are modeled as people

- Prime contractors are not considered to be a part of the in-house workforce and are modeled as money that must be managed by the in-house workforce.
- The amount of reliance on contracting is fixed in the baseline model. A switch has been implemented to make this parameter endogenous and dependent on various uncertainty parameters, but it is deactivated in the baseline model.
- Roughly 42% of NASA's budget is allocated to external procurement
 - External procurement does not include funds given to on-site support contractors
 - The assumed value of 42% is derived from procurement statistics by effort for FY 2002 to FY 2004
 - The figure below shows how we broke down procurement efforts
 - The values were estimated at 40%, 43%, and 44% in FY 2002, FY 2003, and FY 2004 respectively

	<u>Category</u>	<u>Total (Millions)</u>	<u>Number of Awards</u>
	<u>Total</u>	<u>\$ 9,085.9</u>	<u>11,650</u>
50% Support Contractors	<u>Research & Development</u>	<u>\$ 1,751.3</u>	<u>1,978</u>
	Space Station	508.2	10
50% Procurement Contractors	Aeronautics & Space Technology	435.9	984
	Space Flight	309.5	67
	Space Science & Applications	273.5	285
	Space Operations	10.2	23
	Commercial Programs	3.2	17
	Other Space R&D	77.6	146
	Other R&D	133.2	446
		<u>Services</u>	<u>\$ 5,882.2</u>
Support Contractors	Professional, Admin. & Mgmt. Support	3,212.9	660
	ADP & Telecommunications	767.5	253
	Operation of Gov't-owned Facilities	571.6	52
	Special Studies & Analyses-Not R&D	319.3	130
	Transportation, Travel & Relocation Svc.	286.0	31
	Quality Control, Testing & Inspection	147.0	33
	Maint., Repair or Alteration Real Property	102.6	215
	Architect & Engineering Services	86.6	175
	Other Services	388.7	1,905
Procurement Contractors	<u>Supplies & Equipment</u>	<u>\$ 1,452.4</u>	<u>6,218</u>
	Space Vehicles	1,122.2	57
	ADP Equipment, Software, Supplies & Support Equipment	113.0	1,809
	Instruments & Laboratory Equipment	27.9	424
	Fuels, Lubricants, Oils & Waxes	24.1	76
	Chemicals & Chemical Products	24.3	46
	Electrical & Electronic Equip. Component	20.9	119
	Furniture	7.4	212
	Aircraft Launch, Landing & Ground Equip.	4.3	2
	Other Supplies & Equipment	108.3	3,473

Figure 141. Assumed Breakdown of Procurement Efforts [MSFC, 2002].

- Roughly 65% of the portion of the budget that does not go to external procurement goes to technical employees (both civil servants and support contractors)
 - The assumed value of 65% is estimated from workforce statistics between FY 1994 and FY 2006
 - In FY 2002 the percentage of total workforce (including support contractors) in science and engineering was 69%
 - Between FY 1994 and FY 2006 the percentage of the civil servant workforce in engineering ranged from 58.75% to 60.5% while the percentage in science ranged from 4.98% to 5.94%.



This appendix provides detailed documentation of the ESMD model components used in the integrated model. For each component, the structure is shown and detailed documentation is provided. The entire software-generated model documentation is available from the author upon request.

Congress and Executive Component

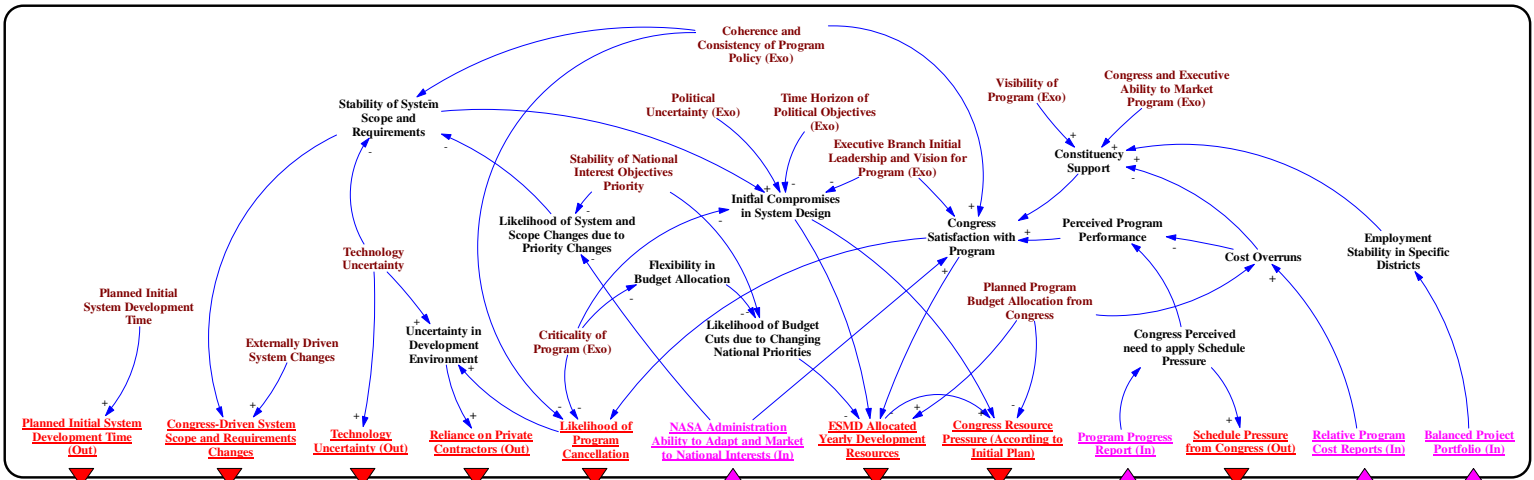


Figure 142: Congress and Executive Component Structure

The Congress and Executive component is responsible for defining the vision for the US space exploration enterprise, as well as providing the level of funding necessary to develop and operate a safe exploration system. Many external factors affect the ability and willingness of the Congress and Executive to define and implement a realistic and safe system. Some of these external factors include: *Political Uncertainty*, *Time Horizon of Political Objectives*, and the *Executive Branch Initial Leadership and Vision for Program*. Those three external factors influence the variable *Initial Compromises in System Design* along with the *Criticality of the Program*. In turn, the variable *Initial Compromises in System Design* will have an impact on the life cycle cost of the system. This variable, when combined with *Technology Uncertainty* will increase the operational costs of the system.

In addition, *Technology Uncertainty* affects the *Stability of System Scope and Requirements*. The *Coherence and Consistency of Program Policy* can counterbalance this effect. The ability of NASA to market the space exploration program will influence the effective criticality of the program. This ability is very important because a program perceived as being critical is less likely to be canceled, compromised, or subjected to budget cuts. *Uncertainty in the Development Environment* is a factor that affects the amount of outside contracting desired by the Agency. *Uncertainty in the Development Environment* is caused by a combination of *Technology Uncertainty* and *Likelihood of Program Cancellation*. While we were not able to find hard quantitative evidence that *Uncertainty in the Development Environment* affects the amount of *Reliance on Private Contractors*, many of the interviewees confirmed that a relation exists. The Congress and Executive model component receives project cost and performance reports from NASA. *Cost Overruns* have a negative impact on *Congress Satisfaction with the Program*. Similarly, *Constituency Support* can also be influenced by *Cost Overruns*. However, other factors also play a role, including the *Visibility of the Program* and the *Congress and Executive Ability to Market the Program* to constituents. In addition to cost, the *Perceived Program Performance* also has an impact on *Congress Satisfaction with the Program*. A highly visible project that is perceived to be on schedule and within budget, and that provides quality jobs in selected congressional districts is likely to be perceived as highly successful and be strongly supported by members of Congress. On the other hand, if Congress becomes dissatisfied with a project that is not

perceived to be critical to national objectives, inevitable budget cuts are likely to be directed toward this particular project.

Another theme often mentioned by interviewees is that of Congressional risk tolerance and risk perception. Space flight is a risky business. While every effort must be put in place to mitigate hazards and improve safety as much as possible, many interviewees mentioned a mismatch between risk perception at the Congress and Agency levels. Interviewees mentioned that this mismatch could lead to unrealistic safety, cost, and schedule expectations from Congress.

Congressional and Executive dynamics are extremely complex. In this model component, we did not attempt to precisely quantify the relationships between different variables. Instead, we merely tried to improve our confidence in the existence of these relationships. In the baseline model, the variables in this component are in equilibrium, that is, unless the values of external variables in this component are modified, the component will have negligible effect on the dynamics of the system. Nevertheless, all the relationships have been implemented in the model, thus allowing us to test Congress and Executive-related policies as well as scenarios where external events affect national priorities and Agency funding. Similarly, in the baseline model, NASA's budget is exogenous, that is, based on existing predictions and not affected by national priorities and future Congressional satisfaction. The model is equipped to relax this assumption by making part of NASA's budget allocation dependent on *Congress Satisfaction with the Program*.

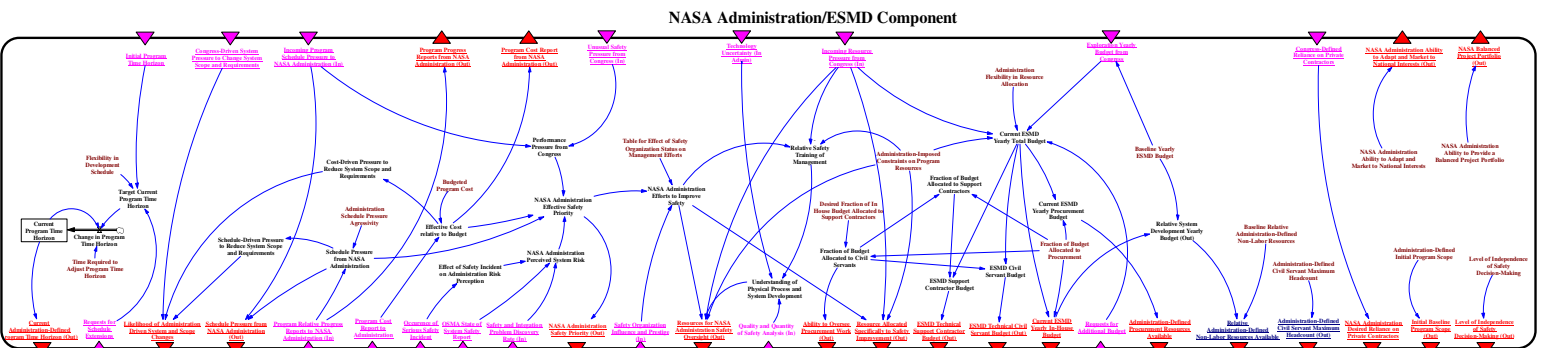


Figure 143: NASA Administration and ESMID Component Structure

The purpose of the NASA administration and ESMD component is to define the agency level policies, requirements, and guidelines that will enable the development of a safe and successful exploration system. The Agency receives directives and funding from Congress, and allocates resources according to program needs. However, NASA has limited flexibility in resource allocation because some of the budgets associated with larger programs are dictated at the Congressional level. For example, by the mandate of the President, the US space shuttle program has to be supported until it is retired around 2010. This and other constraints on resource allocation flexibility take up a significant portion of NASA's approximately \$17 billion annual budget. The primary function of the NASA Administration and ESMD component is to allocate resources (human and material) to different programs while respecting the constraints imposed by Congress and presidential administrations.

In addition to dealing with budget and financial matters, the Agency and directorates are also responsible for providing programs and projects with a highly qualified and trained civil servant workforce. This workforce is drawn mostly from the individual NASA centers, which provide the institutional technical knowledge and skills necessary for a safe and successful exploration enterprise. The technical workforce available at the centers is then matrixed to the individual programs and projects based on technical needs and the availability of funding.

ESMD monitors the progress and cost of the individual programs under its responsibility. The *Incoming Resource Pressure from Congress* affects NASA and ESMD by reducing resources available for safety training and improvements, as well as potentially compromising the quality of safety oversight provided to individual programs and projects by NASA/OSMA and ESMD. In addition to these effects on safety, budget pressure flows downstream and affects the resources available to develop the exploration system.

The NASA Administration and ESMD component has very few exogenous inputs. *Technology Uncertainty* is one of those external inputs. This exogenous input influences the management and technical personnel *Understanding of the Physical Process and System Development Environment*. The *Technology Uncertainty* variable is not meant to have a precise numerical value. Instead, it is used to investigate scenarios where the chosen system architecture includes completely proven or field-tested technologies versus scenarios with

new, undeveloped technologies. Another external input is the planned profile of work completion. Projects are usually not completed in a linear way. There are project phases that require a larger workforce and project managers can have varying levels of flexibility in hiring support and procurement contractors to supplement their workforce during critical phases. For the ESMD baseline model, we made the work completion profile proportional to the total ESMD budget allocation (confirmed and projected) at any point in time. This budget-weighted profile is a first order estimate that may not be entirely accurate, but it is believed to be a much better approximation of work completion than a linear profile.

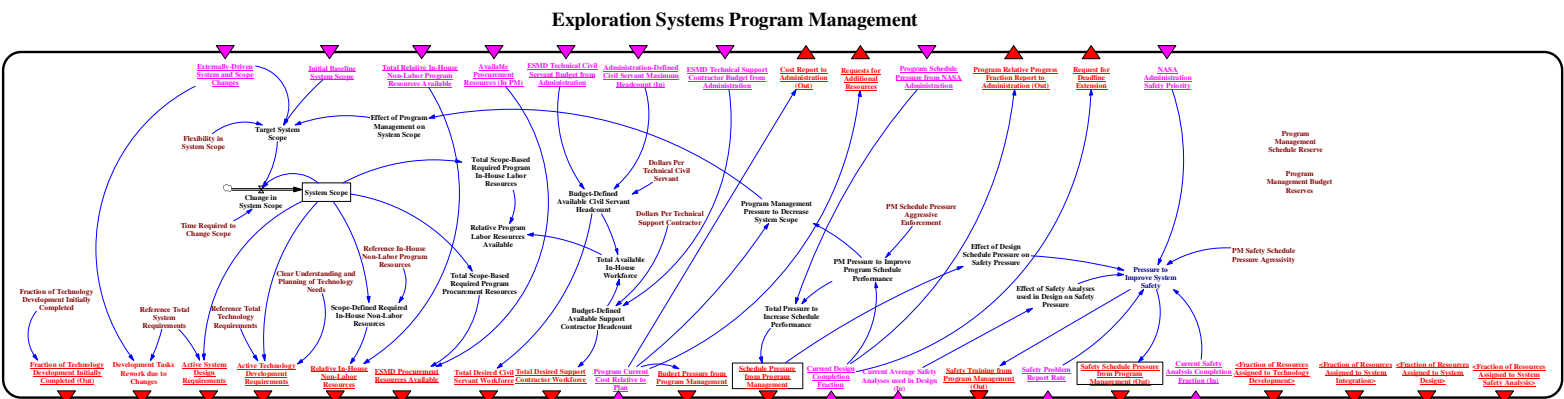


Figure 144: Exploration Program and Project Management Component Structure

The purpose of the Exploration Systems Program Management component is to reproduce the behavior of program and project managers during real system development. Program managers have to ensure that the system under development meets technical requirements (including both safety and performance requirements) while remaining within budget and on schedule. Program managers use multiple levers to achieve these objectives, including reshuffling schedules, reallocating resources (human and material), and applying various pressures to lower-level managers, engineers and other technical workers.

The program management component is essentially a control system trying to regulate system development. In general, the task assigned to program management is a multi-objective control problem that may require trade-offs between different system qualities. In most cases, the overall objectives will be dictated by higher-level elements of the control structure such as the NASA Administration, the Executive Branch, or Congress, but the implementation details will be the responsibility of the programs and projects. For example, the *System Scope* and high level technical requirements as well as the budget and workforce available are usually constraints applied by the Agency or the Directorate, which are based on desires, requirements, and constraints at the Congressional and Executive level. Program managers have to report to the Administration at various system development milestones. Another potential constraint is the amount of contracting desired by Congress or the agency. In this model, resources are allocated to five different bins: 1) Technology development, 2) System integration, 3) Safety analyses, 4) System design, and 5) Other ESMD projects. The total amount of in-house resources available has a large impact on system development. It allows project managers to allocate sufficient resources into the five different bins while keeping management reserves to account for uncertainty, disturbances, and unplanned events. In case of a high likelihood of missing a development deadline, program management can place *Requests for Deadline Extension* or in some cases request a decrease in system scope, which translates into either dropped or waived requirements.

Another way to alleviate the likelihood of missing schedule is to improve the rate of design completion. This improvement in design completion rate can be achieved through the allocation of more resources to design or by applying pressure to work faster. If the project is perceived to be over budget, management can place *Requests for Additional Resources* or

deadline extensions (which ultimately translate into more resources) or try to improve completion rate, which also reduces fixed costs (for the project).

Just as at the Agency level, the amount and severity of safety and integration problem reports has an impact on the amount of energy and resources expended by project management to improve system safety. If the *Safety Analysis Completion Rate* falls behind the *Development Completion Rate*, more resources may be allocated to safety in order to catch up. Resources are limited, however, and allocating more resources to safety means that fewer resources will be available for other activities such as design or integration and vice versa. The impact of various resource allocation strategies are discussed in a scenario presented later.

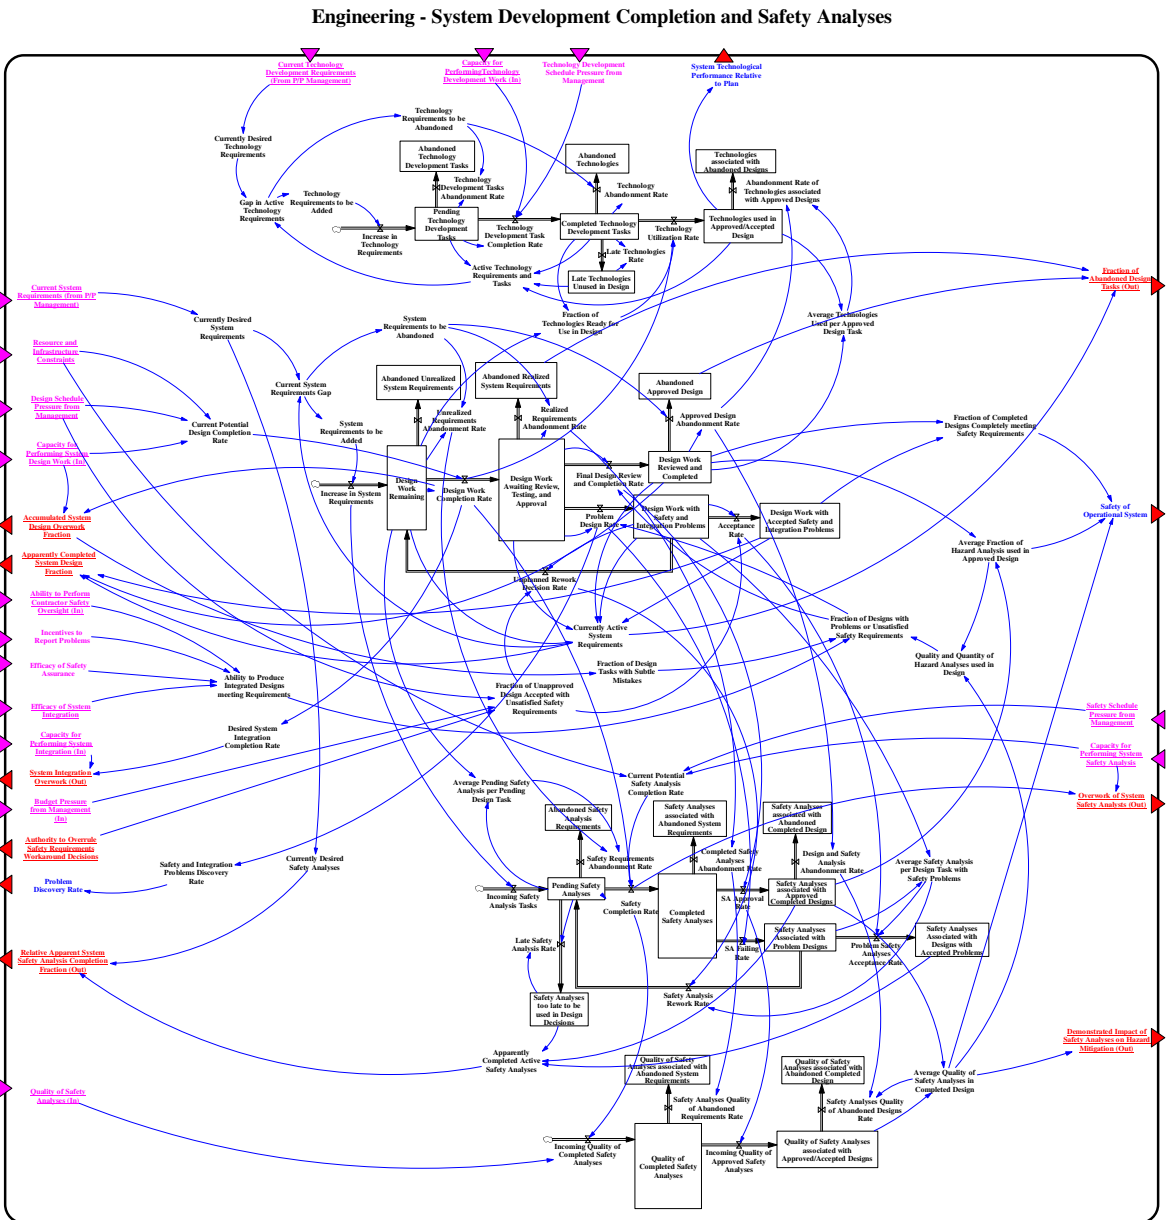


Figure 145: System Development Completion and Safety Analyses Component Structure

The System Development Completion and Safety Analyses component is at the core of the ESMID model. It includes three different flows that have to be synchronized and coordinated to produce a final integrated product. The three flows are: 1) Technology development, 2) System development tasks, and 3) Safety analyses. The timing of these flows is critical. Late technologies cannot be used in design without significant development delays. Similarly, late

safety analyses might delay design or might not be used in design decisions, resulting in an unsafe system. In addition, some development work might have to be redone if problems are found along the way. This rework causes delays in both the design and safety activities and thus further increases schedule pressure. System engineering and integration is responsible for making sure that the three flows of technology, design and safety merge in a synchronized and coordinated manner.

The first flow is that of technology development. Once technology requirements have been defined, technology development work has to be completed on time for a specific technology to be available in design. Some technologies might be abandoned along the way because of changing requirements or external factors. Other technologies might not be ready on time for design or deployment. In some cases, a technology might not live up to its promises, thus requiring more investments, resources, or time. Some of these technologies might have to be abandoned or replaced by an already available technology. Good technology requirements planning should include off-ramps to minimize the impact of technology abandonment.

The second flow is that of development tasks completion. Initial system requirements translate to remaining development tasks to be accomplished. Development tasks get completed according to the current development capacity, which is a function of the resources and workforce available as well as the workforce productivity. These factors, in turn, affect workforce experience, training, overwork, and schedule pressure. Once development tasks are completed, they are submitted for review, testing and approval. Subsequently, some tasks will be completed, while others will be found to have safety or integration problems requiring changes and rework—hence the Rework Cycle. The percentage of tasks to be reworked will depend on many factors including: *Availability of Information from Hazard Analyses*, *Efficacy of System Integration*, *Ability to Perform Contractor Safety Oversight*, efficacy of testing, *Efficacy of Safety Assurance*, *Design Schedule Pressure from Management*, mistakes made by overworked or burned out development personnel and *Incentives to Report Problems*. At any point in the development cycle, changes in requirements may necessitate the abandonment of some requirements or completed designs, as well as the introduction of new requirements to be transformed into remaining development tasks. A scenario to be discussed later addresses the impact of changes along the system development life cycle.

The completion of safety analysis tasks mirrors the completion of development tasks through a non-conserved co-flow structure [Sterman, 2000]. However, the timing of development and safety completion may not be synchronized in every case. If the information from safety analyses is not available at the time when design decisions have to be made, two outcomes are possible: either (1) a decision will be made without the proper safety information or (2) development will be delayed while waiting for safety analyses to be completed. Neither outcome is optimal. If safety analysts work hand-in-hand with engineering and system engineering and integration is performed correctly, the safety and development flows should be tightly connected and the safety analyses should be performed at the same time as development tasks. While this synchronization of safety analysis and design task completion is an ideal situation, it may not always reflect the way things are done in the ESMD or in typical system development activities. Because it is very difficult to have the safety information available exactly when it is needed, a good approach is to try to anticipate safety analysis needs in order to have a head start over development tasks. Anticipating needs may not always be possible because of the highly iterative nature of development activities, but it should be attempted when possible. Otherwise, the workforce may choose to accelerate the completion of safety analyses by cutting corners and reducing the fidelity and *Quality of Safety Analyses Performed*. In addition to keeping track of safety analyses performed, a coflow structure is used to monitor the *Quality of Safety Analyses Performed* and the *Average Quality of Safety Analyses in the Completed Design*.

One of the major variables calculated in the System Development Completion and Safety Analyses Component is the ultimate *Safety of the Operational System*. This variable is a dimensionless, multi-attribute utility function that is meant to characterize how safe the operational will be. Its inputs are the *Average Fraction of Hazard Analyses used in the Approved Design*, the *Fraction of Completed Designs Meeting Safety Requirements*, and the *Average Quality of Hazard Analyses used in the Completed Design*. The variable is meant to serve as a relative measure between simulation runs of how well tasks crucial to the safety of the system are performed in the design process.

Engineering - Technical Personnel Resources and Experience

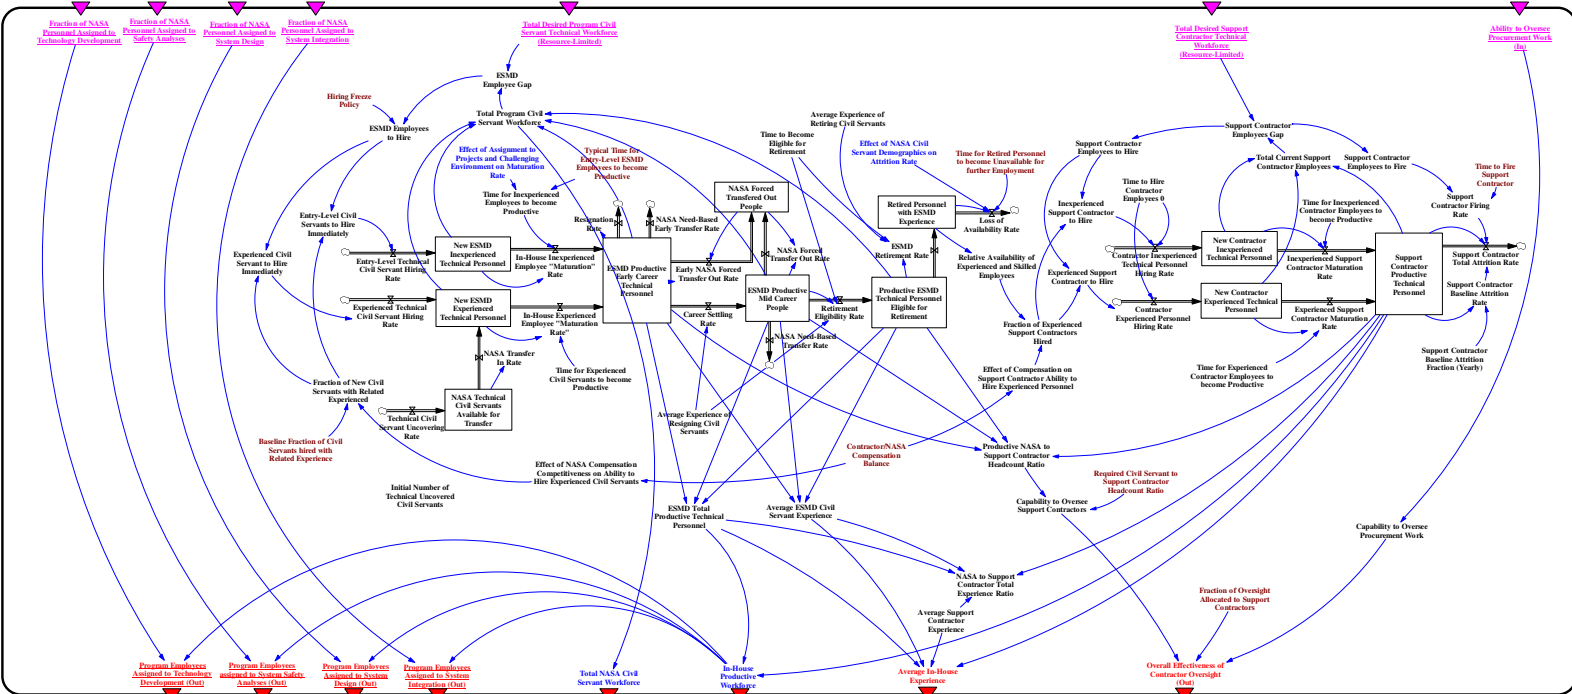


Figure 146: Engineering – Technical Personnel Resources and Experience Structure

The purpose of the Engineering (Technical Personnel Resources and Experience) component is to keep track of the human resources working on ESMD projects. This component was initialized and calibrated using employment data available on the NASA web site. The objective is to monitor the availability and characteristics of the technical workforce responsible for the development of the exploration system. The component considers the number of people hired for entry-level positions and for experienced positions, as well as transfers between ESMD and other NASA directorates such as the Space Operations Mission Directorates (SOMD). It also keeps track of the experience of NASA technical employees as well as attrition rates, retirements, and employees choosing to use their early retirement option in order to work as consultants or as contractor employees. The scope of this component includes the entire technical workforce working in-house at NASA, that is, NASA technical civil servants, and support contractors working at NASA centers. Employees working for the procurement contractors are not explicitly modeled in this component (the budget allocations for procurement are explicitly modeled in the NASA Administration and ESMD Component of the model.). All these factors have a critical impact on NASA's ability to develop a safe and successful exploration system.

This component receives inputs from higher-level components, such as the desired size of the civil servant technical workforce and the budget available to hire support contractors. The allocation of technical human resources into the five different bins discussed previously is also provided as an input to this component. The difference between the desired workforce size and the current workforce size drives the hiring rate. Civil servants can either be hired at the entry-level or at the experienced level. Additionally, if civil servants from another NASA directorate are available, they will be transferred to ESMD and have priority over new hires. Once civil servants are hired or transferred to ESMD, it takes time for them to become fully productive. According to interview data, it takes approximately 3 months for an experienced hire to become productive, while it takes up to two years for entry-level hires to become productive. The employment data shows that once civil servants become productive, they will leave the NASA workforce in one of two ways: either (1) they will stay at NASA until they retire or (2) they will stay for a few years and then make an early career decision to work in the private sector or to work for a different government agency. The data shows that very few mid-career NASA employees leave the civil servant workforce. However, it happens

regularly that NASA employees are transferred to and from projects and directorates. The model accounts for this possibility and allows analysts to investigate the impact of transfers on system development. The component also takes into account the fact that hiring civil servants is more difficult than hiring support contractors. The only requirement to hire a support contractor is to have the budget available. More important, laying off support contractors is much easier. Civil servants are frequently transferred, but firing a civil servant is rare and reductions in workforce are very difficult to do on the government side, which is a disadvantage in large-scale system development. This reality, combined with the fact that recent administrations have had a desire to reduce the size of the government workforces, creates a strong bias toward hiring more contractors.

The output of this component includes the number of in-house technical employees working in the five areas mentioned previously: 1) Technology, 2) Integration, 3) Safety, 4) Development, and 5) Other ESMD projects. The output also includes the *Average ESMD Civil Servant Experience* and *Average Support Contractor Experience* which both have impacts on productivity, and capability to oversee contractors (support and procurement).

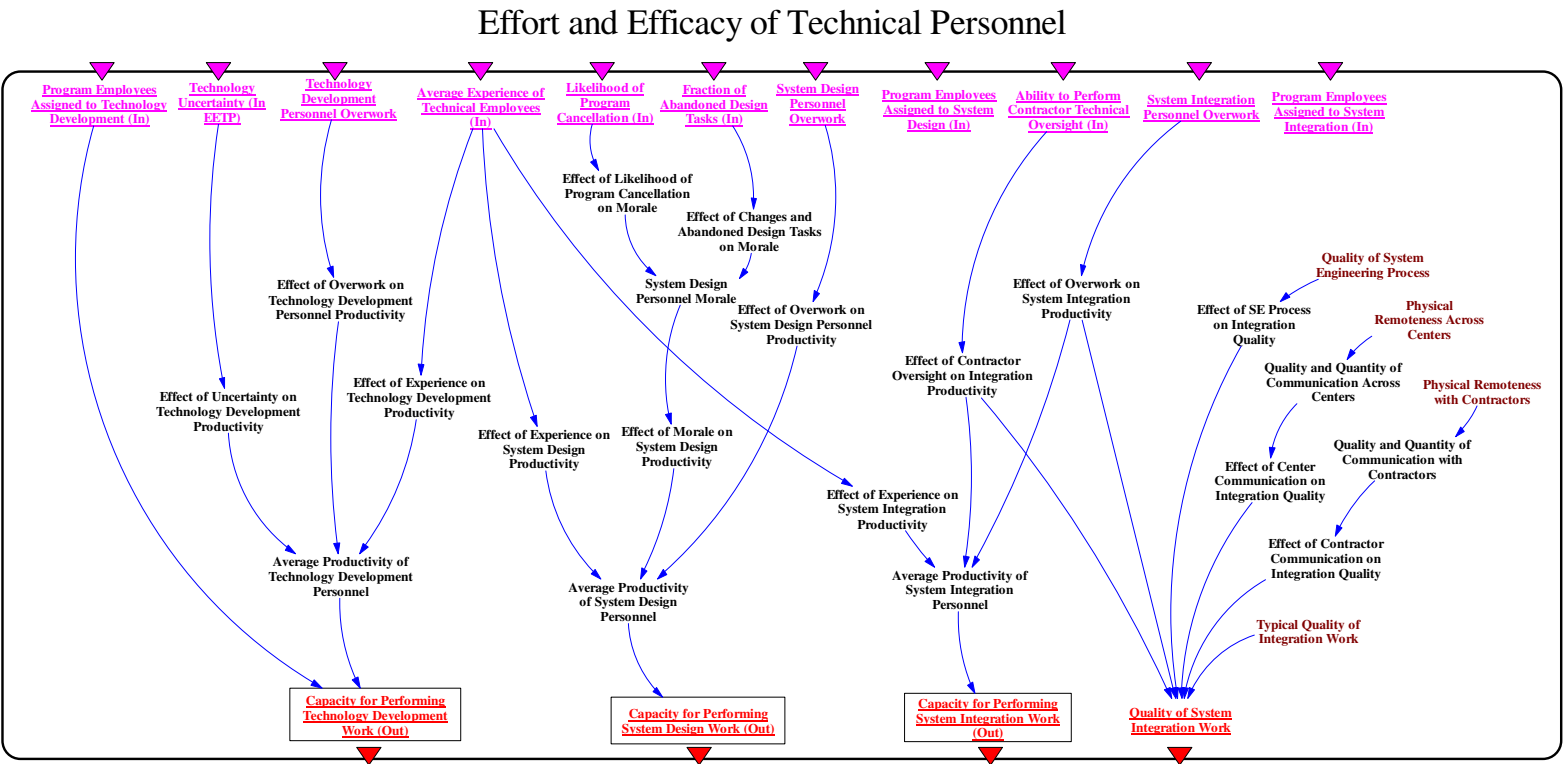


Figure 147: Engineering – Effort and Efficacy of Technical Personnel Component Structure

The purpose of the Engineering (Effort and Efficacy of Technical Personnel) component is simply to collect information from various sources in the model and output the total capacity of in-house workforce to perform development work in areas of technology development, system integration, and system development. In order to accomplish this, the component needs such inputs as the number of employees assigned to different areas, the overwork of employees in those areas and other inputs that affect the motivation and productivity of employees such as the *Likelihood of Program Cancellation*, *Requirements and Design Changes*, and project abandonment. This component also computes a value for the average *Quality of System Integration Work*, which is a function of many factors such as the *Ability to Perform Contractor Oversight*, the *Quality of the System Engineering and Integration Process* and the *Quality and Quantity of Communication* across NASA Centers and contractor offices.

Engineering - Effort and Efficacy of System Safety Analysts

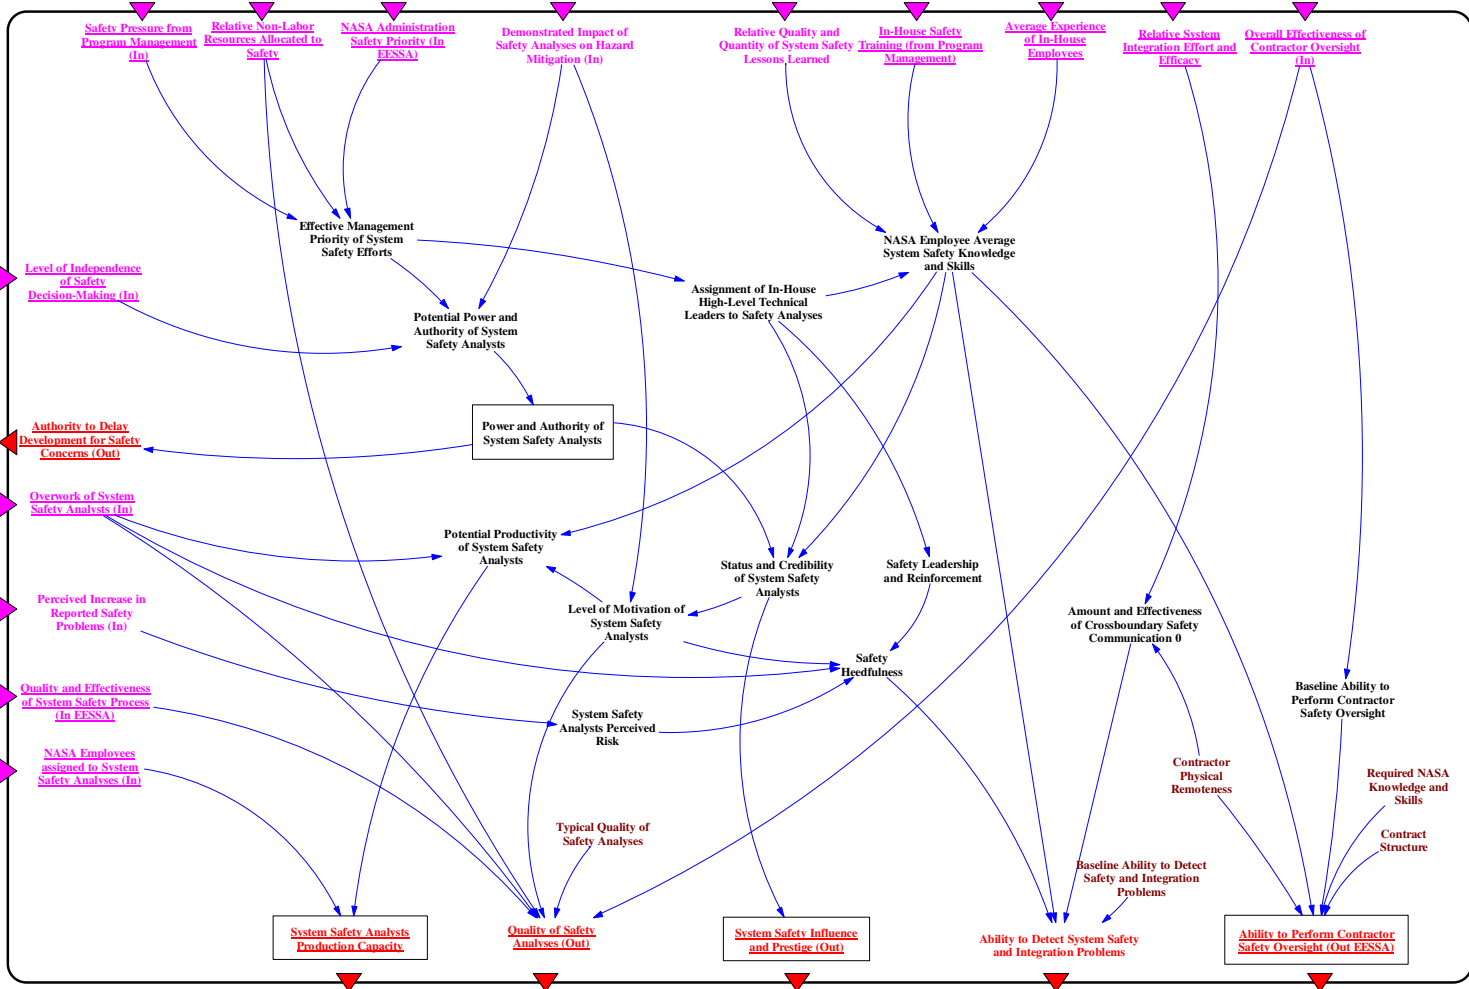


Figure 148: Safety and Mission Assurance - (EESSA) Component Structure

The focus of the Safety and Mission Assurance component is on the effort and efficacy of in-house employees working on safety analyses. The purpose of the component is to determine the capacity of safety analysts to work hand-in-hand with other engineers and technical people in order to produce high-quality, useful safety information to be used in making design decisions. Many soft factors such as the *Power and Authority of System Safety Analysts* and the *Status and Credibility of System Safety Analysts* will have a large influence on the impact of system safety analysts on the safety of the final system. Consequently, all of these factors have to be included in this component, even though they are difficult to quantify and their influence might not be completely understood. The outputs of this component include the *Capacity for Performing Safety Analyses*, current *Quality of Safety Analyses*, the *Influence and Prestige of the Safety Organization* and their *Authority to Delay System Development for Safety Concerns*, and their *Ability to Detect System Safety and Integration Problems*.

REFERENCES

- Ackoff, R. L. (1971). "Towards a system of systems concept." Management Science **17**(11): 661–671.
- Albee, A., et al. (2000). Report on the Loss of the Mars Polar Lander and Deep Space 2 Missions. JPL Special Review Board, Jet Propulsion Lab (JPL).
- Apostolakis, G. E. (2004). "How Useful is Quantitative Risk Assessment?" Risk Analysis **24**(3): 515-520.
- Ashby, W. R. (1956). An Introduction to Cybernetics. London, UK, Chapman and Hall.
- Ayyub, B. M. (2003). Risk Analysis in Engineering and Economics, Chapman & Hall.
- Bedford, T., R. Cooke (2001). Different Types of Uncertainty. Probabilistic Risk Analysis: Foundations and Methods. Cambridge, UK, Cambridge University Press: 17-38.
- Bedford, T., R. Cooke (2001). Probabilistic Risk Analysis: Foundations and Methods Cambridge, UK, Cambridge University Press.
- Belanger, P. (1995). Control Engineering: A Modern Approach. Oxford, UK, Oxford University Press.
- Boisjoly, R. (1987). "Ethical decisions - Morton Thiokol and the space shuttle Challenger disaster." Journal of the American Society of Mechanical Engineering: 1-13.
- Branscome, D. R. (1999). WIRE Mishap Investigation Board Report. Washington, DC, NASA.
- Burt, R. (1995). Structural Holes: The Social Structure of Competition. Cambridge, MA, Harvard University Press.
- Carroll, J. S., J.W. Rudolph , S. Hatakenaka (2002). "Learning from experience in high-hazard organizations." Research in Organizational Behavior **24**: 87-137.
- Checkland, P. (1981). Systems Thinking, Systems Practice. New York, NY, John Wiley & Sons.
- Cicerone, R. J. (2001). Climate Change Science: An Analysis of Some Key Questions. Washington, D.C., National Academy of Sciences Committee on the Science of Climate Change.

Cook, R., Jens Rasmussen (2005). "Going Solid: A Model of System Dynamics and Consequences for Patient Safety." Quality and Safety in Healthcare **14**: 130-134.

Cooke, D. L. (2003). "A System Dynamics Analysis of the Westray Mine Disaster." System Dynamics Review **19**(2): 139-166.

Creswell, J. W. (1994). Research Design: Qualitative and Quantitative Approaches, Sage Publications.

Davis, T. (2006). Select Bipartisan Committee to Investigate the Preparation for and Response to Hurricane Katrina. Washington, D.C., U.S. House of Representatives.

Dekker, S. (2005). Resilience Engineering: Chronicling the Emergence of Confused Consensus. Resilience Engineering: Concepts and Precepts. D. D. W. a. N. L. Erik Hollnagel, Ashgate Publishing.

DoD (2000). MIL-STD-882D - Standard Practice for System Safety. D. o. Defense.

Dulac, N., Nancy G. Leveson (2004). An Approach to Design for Safety in Complex Systems. INCOSE04. Toulouse, France.

Dulac, N., Nancy G. Leveson (2004). Incorporating Safety in Early System Architecture Trade Studies. International System Safety Conference San Diego, CA.

Edwards, W. (1977). "How to Use Multiattribute Utility Measurement for Social Decisionmaking." IEEE Transactions on Systems, Man, and Cybernetics **7**: 326-340.

Fischhoff, B. (1975). "Hindsight != Foresight: the effect of outcome knowledge on judgment under uncertainty." Journal of Experimental Psychology: Human Perception and Performance **1**: 288-299.

Ford, D. (1995). The Dynamics of Project Management: An Investigation of the Impacts of Project Process and Coordination on Performance. Dynamic Engineering Systems. Cambridge, MA, MIT. **Ph.D.**

Ford, D., Serman J. D. (1998). "Dynamic Modeling of Product Development Processes." System Dynamics Review **14**(1): 31-68.

Ford, D. N., John D. Serman (1998). "Expert knowledge elicitation to improve formal and mental models." System Dynamics Review **14**(4): 309-340.

Forrester, J. W. (1961). Industrial Dynamics, Pegasus Communications.

Forrester, J. W. (1969). Urban Dynamics. Cambridge, MA, Productivity Press.

- Forrester, J. W. (1972). World Dynamics. Cambridge, MA, Productivity Press.
- Forrester, J. W. (1985). "'The' Model Versus a Modeling 'Process'." System Dynamics Review **1**(1): 133-134.
- Forrester, J. W. (1989). The System Dynamics National Model: Macrobehavior from Microstructure. Computer-Based Management of Complex Systems: International System Dynamics Conference. P. M. M. E. O. K. Zahn. Berlin, Germany, Springer-Verlag.
- Forrester, J. W. (1992). "Policies, decisions, and information sources for modeling." European Journal of Operational Research **59**: 42-63.
- Freeman, L. C. (2004). The Development of Social Network Analysis: A Study in the Sociology of Science. Vancouver, BC, CAN, Empirical Press.
- Freudenburg, W. R. (1988). "Perceived Risk, Real Risk: Social Science and the Art of Probabilistic Risk Assessment." Science **242**(4875): 44-49.
- Gehman, H. (2003). Columbia Accident Investigation Report, NASA.
- Hastings, D., H. McManus (2004). A Framework for Understanding Uncertainty and its Mitigation and Exploitation in Complex Systems. Engineering Systems Symposium, Cambridge, MA.
- Hollnagel, E. (2002). Understanding Accidents - from Root Causes to Performance Variability. 2002 IEEE 7th Conference on Human Factors and Power Plants.
- Hollnagel, E., David D. Woods and Nancy Leveson (2005). Resilience Engineering: Chronicling the Emergence of Confused Consensus, Ashgate Publishing.
- Howard, R., Matheson, J. (1984). Readings in the Principles and Practice of Decision Analysis. Menlo Park, CA, Strategic Decision Systems.
- Johnson, W. G. (1973). MORT - The Management Oversight and Risk Tree, U. S. Atomic Energy Commission.
- Johnson, W. G. (1980). MORT Safety Assurance System. New York, NY, Marcel Dekker.
- Jones, E. E., Harris V.A. (1967). "The attribution of attitudes." Journal of Experimental Social Psychology **3**: 1-24.
- Keeny, R. L. a. H. R. (1993). Decision Making with Multiple Objectives: Preferences and Value Tradeoffs. Cambridge, UK, Cambridge University Press.
- La Porte, T. R. (1996). "High Reliability Organizations: Unlikely, Demanding, and At Risk." Journal of Contingencies and Crisis Management **63**(4).

La Porte, T. R., Paula Consolini (1991). "Working in Practice But Not in Theory: Theoretical Challenges of High-Reliability Organizations." Journal of Public Administration Research and Theory **1**: 19–47.

Leplat, J. (1987). Occupational accident research and systems approach. New Technology and Human Error. K. D. Jens Rasmussen, and Jacques Leplat. New York, NY, John Wiley & Sons: 181–191.

Leveson, N. G. (1995). Safeware: System Safety and Computers. Reading, MA, Addison-Wesley.

Leveson, N. G. (2000). "Intent Specifications: An Approach to Building Human-Centered Specifications." IEEE Transactions on Software Engineering.

Leveson, N. G. (2001). Evaluating Accident Models using Recent Aerospace Accidents. Cambridge, MA, MIT Department of Aeronautics and Astronautics.

Leveson, N. G. (2003). A New Approach to Hazard Analysis for Complex Systems. International Conference of the System Safety Society. Denver, CO.

Leveson, N. G. (2004). "A New Accident Model for Engineering Safety Systems." Safety Science **42**(4): 237–270.

Leveson, N. G. (2004). "The Role of Software in Spacecraft Accidents." Journal of Spacecraft and Rockets **41**(4): 564- 575.

Leveson, N. G. (2006). A New Approach to System Safety Engineering. Cambridge, MA, Unpublished Manuscript.

Leveson, N. G., Joel Cutcher-Gershenfeld (2004). System Safety Engineering Can Learn from the Columbia Accident. International Conference of the System Safety Society. Providence, RI.

Leveson, N. G., Joel Cutcher-Gershenfeld, Betty Barrett, Alexander Brown, John Carroll, Nicolas Dulac, Lydia Fraile, Karen Marais (2004). Effectively Addressing NASA's Organizational and Safety Culture: Insights from Systems Safety and Engineering Systems. MIT ESD External Symposium, Cambridge, MA.

Leveson, N. G., Mirna Daouk, Nicolas Dulac, Nicolas, Karen Marais (2004). A Systems-Theoretic Approach to Safety Engineering: A Case Study MIT ESD External Symposium, Cambridge, MA.

Leveson, N. G., Nicolas Dulac (2005). Risk Analysis of NASA Independent Technical Authority. Cambridge, MA, MIT.

- Leveson, N. G., Nicolas Dulac (2005). Safety and Risk Driven Design in Complex Systems of Systems. NASA/AIAA Space Exploration Conference. Orlando, FL.
- Lions, J. L. (1996). Ariane 5: Flight 501 Failure. Paris, France, Ariane 5 Failure Inquiry Board.
- Lovelock, J. (1979). Gaia: A New Look at Life on Earth. Oxford, UK, Oxford University Press.
- Lyneis, J. M. (1980). Corporate Planning and Policy Design. Cambridge, MA, Productivity Press.
- Lyneis, J. M. (2000). "System Dynamics for Market Forecasting and Structural Analysis." System Dynamics Review **16**(1).
- Lyneis, J. M., Kenneth G. Cooper, Sharon A. Els (2001). "Strategic Management of Complex Projects: A Case Study Using System Dynamics." System Dynamics Review **17**(3).
- Lyneis, J. M., Kimberly Reichelt (1999). "The Dynamics of Project Performance: Benchmarking the Drivers of Cost and Schedule Overrun." European Management Journal **17**(2).
- Magee, C. L., O. de Weck (2004). Complex System Classification. INCOSE04. Toulouse, France.
- Marais, K. (2005). A new approach to risk analysis with a focus on organizational risk factors. Department of Aeronautics and Astronautics. Cambridge, MA, Massachusetts Institute of Technology.
- Marais, K., Nancy G. Leveson (2003). Archetypes for Organizational Safety. Workshop on Investigating and Reporting of Incidents and Accidents. Williamsburg, VA.
- Marais, K., Nicolas Dulac, Nancy G. Leveson (2004). Beyond Normal Accidents and High Reliability Organizations: Lessons from the Space Shuttle. MIT ESD External Symposium, Cambridge, MA.
- March, J. G., Lee S. Sproull, Michal Tamuz (1991). "Learning from Samples of One or Fewer." Organization Science **2**(1): 1-13.
- Mass, N. (1991). "Diagnosing Surprise Model Behavior: A Tool For Evolving Behavioral And Policy Insights." System Dynamics Review **7**(1): 68-86.
- McCurdy, H. (1994). Inside NASA: High Technology and Organizational Change in the U.S. Space Program, Johns Hopkins University Press.

McCurdy, H. (2001). Faster, Better, Cheaper: Low-Cost Innovation in the U.S. Space Program. Baltimore, MD, Johns Hopkins University Press.

McKernan, J. (1996). Curriculum Action Research: A handbook of methods for the reflective practitioner. London, Kogan Page.

Meadows, D., Dennis Meadows, Jorgen Randers (1992). Beyond the Limits: Confronting Global Collapse, Envisioning A Sustainable Future. Post Mills, VT, Chelsea Green.

Meadows, D., Dennis Meadows, Jorgen Randers, William Behrens (1972). Limits to Growth. New York, NY, New American Library.

Melchers, R. (1999). Uncertainties in reliability assessment. Structural Reliability Analysis and Prediction. Chichester, UK, John Wiley & Sons: 34-45.

Morecroft, J. D. W. (1983). "System Dynamics: Portraying Bounded Rationality." International Journal of Management Science **11**(2): 131-142.

Morecroft, J. D. W. (1984). "Strategy Support Models." Strategic Management Journal **5**(3): 215-229.

Morecroft, J. D. W. (1985). "Rationality in the Analysis of Behavioral Simulation Models." Management Science **31**(7): 900-916.

Morecroft, J. D. W. (1988). "System Dynamics and Microworlds for Policymakers." European Journal of Operational Research **35**(3): 301-320.

NASA (1998). NASA/ESA Investigation Board on the SOHO Mission Interruption. Washington, DC, NASA.

NASA (2005). Technical Authority Implementation Guidance. Washington, DC, NASA.

Nielsen, D. S. (1971). The Cause/Consequence Diagram Method as a Basis for Quantitative Accident Analysis. Denmark, Danish Atomic Energy Commission, RISO Research Establishment.

NTSB (1983). NTSB Accident Report, NTSB.

Oberkampf, W., S. DeLand, B. Rutherford, K. Diegert, K. Alvin (1999). "A new methodology for the estimation of total uncertainty in computational simulation." AIAA Paper 99-1612.

Paich, M. (1985). "Generic Structures." System Dynamics Review **1**: 126-132.

Paté-Cornell, E. M. (1990). "Organizational Aspects of Engineering System Safety." Science **250**: 1210-1217.

- Paté-Cornell, E. M., Dean Michael Murphy (1996). "Human and management factors in probabilistic risk analysis: the SAM approach and observations from recent applications." Reliability Engineering and System Safety **53**(2): 115-126.
- Perrow, C. (1982). The President's Commission and the Normal Accident. The Accident at Three Mile Island: The Human Dimension. C. P. W. David L. Sills, and Vivien B. Shelarski Westview Press.
- Perrow, C. (1983). "The Organizational Context of Human Factors Engineering." Administrative Science Quarterly **28**: 521-541.
- Perrow, C. (1999). Normal Accidents: Living with High-Risk Technologies. Princeton, NJ, Princeton University Press.
- Perrow, C. (2004). Review of NAT-HRO Paper. N. G. Leveson. Cambridge, MA.
- Pidgeon, N., M. O'Leary (1998). Organizational Safety Culture: Implications for Aviation Practice. Aviation Psychology in Practice. N. M. N. Johnston, and R. Fuller. Burlington, VT, Ashgate.
- Ramo, S. (1973). The systems approach. Systems Concepts: Lectures on Contemporary Approaches to Systems. R. F. J. Miles. New York, NY, John Wiley & Sons: 13-32.
- Rasmussen, J. (1997). "Risk Management in a Dynamic Society: A Modelling Problem." Safety Science **27**(2/3): 183-213.
- Rasmussen, J., Annelise Mark Pejtersen, L. P. Goodstein (1994). Cognitive Systems Engineering, Wiley-Interscience.
- Rasmussen, J., Inge Svedung (2000). "Proactive Risk Management in a Dynamic Society." Swedish Rescue Services Agency.
- Rasmussen, J., Inge Svedung (2002). "Graphic Representation of Accident Scenarios: Mapping System Structure and the Causation of Accidents." Safety Science **40**: 397-417.
- Reason, J. (1995). "A System Approach to Organizational Error." Ergonomics **38**(8): 1708-1721.
- Repenning, N. P. (2001). "Understanding Fire Fighting in New Product Development." Journal of Product Innovation Management **18**(5): 285-300.
- Repenning, N. P., John D. Serman (2001). "Nobody Ever Gets Credit for Fixing Problems that Never Happened: Creating and Sustaining Process Improvement." California Management Review **43**(4): 64-88.

- Repenning, N. P., John D. Sterman (2002). "Capability Traps and Self-Confirming Attribution Errors in the Dynamics of Process Improvement." Administrative Science Quarterly **47**: 265-295.
- Rijpma, J. A. (1997). "Complexity, tight-coupling and reliability: connecting normal accidents theory and high reliability theory." Journal of Contingencies and Crisis Management **5**(1): 15-23.
- Rijpma, J. A. (2003). "From Deadlock to Dead End: The Normal Accidents-High Reliability Debate Revisited." Journal of Contingencies and Crisis Management **11**(1): 37-45.
- Roberts, K. H. (1990). "Managing high reliability organizations." California Management Review **32**(4): 101-114.
- Roberts, K. H. (1990). "Some characteristics of one type of high reliability organization." Organization Science **1**(2): 160-176.
- Roberts, K. H., P. Madsen, V. Desai, D. Van Stralen (2005). "A case of the birth and death of a high reliability healthcare organisation." Quality and Safety in Healthcare **14**: 216-220.
- Rochlin, G. I. (1991). Iran Air Flight 655 and the USS Vincennes: Complex, Large-Scale Military Systems and the Failure of Control. Social Responses to Large Technical Systems: Control or Anticipation. T. R. La Porte, Kluwer Academic Publishers.
- Rochlin, G. I., Todd R. La Porte, Karlene H. Roberts (1987). The Self-Designing High Reliability Organization, Naval War College Review.
- Rogers, W. P. (1986). Report of the Presidential Commission on the Space Shuttle Challenger Accident. Washington, D.C., Government Printing Office.
- Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. Advances in experimental social psychology. L. Berkowitz. New York, NY, Academic Press. **10**: 173-220.
- RTF Website, N. Shuttle Return to Flight Website.
- Rudolph, J. W., Nelson P. Repenning (2002). "Disaster Dynamics: Understanding the Role of Quantity in Organizational Collapse." Administrative Science Quarterly **47**: 1-30.
- Sagan, S. D. (1993). The Limits of Safety: Organizations, Accidents, and Nuclear Weapons. Princeton, NJ, Princeton University Press.
- Sagan, S. D. (2004). "The Problem of Redundancy Problem: Why More Nuclear Security Forces May Produce Less Nuclear Security." Risk Analysis **24**(4): 935-946.

- Salge, M., Peter M. Milling (2006). "Who is to blame, the operator or the designer? Two stages of human failure in the Chernobyl accident." System Dynamics Review **22**: 89-112.
- Sarter, N. D., D. Woods (1996). "'How in the world did I ever get into that mode?': Mode error and awareness in supervisory control." Human Factors **37**: 5-19.
- Schein, E. H. (1969). Process consultation: Its role in organizational development. Reading, MA, Addison-Wesley.
- Schulman, P. R. (1993). "The negotiated order of organizational reliability." Administration and Society **25**(3): 353-372.
- Senge, P. M. (1990). The Fifth Discipline: The Art and Practice of the Learning Organization. New York, NY, Doubleday Currency.
- Shewhart, W. A. (1939). Statistical Method from the Viewpoint of Quality Control. New York, NY, Dover.
- Shrivastava, P. (1992). Bhopal: Anatomy of a crisis. London, UK, P. Chapman.
- Simon, H. A. (1957). Rationality and Decision-Making. Models of Man. New York, NY, John Wiley.
- Simon, H. A. (1976). Administrative Behavior. New York, NY, The Free Press.
- Slovic, P. (1999). "Trust, Emotion, Sex, Politics, and Science: Surveying the Risk Assessment Battlefield." Risk Analysis **19**(4): 689-701.
- Starbuck, W. H., F.J. Milliken (1988). "Challenger: fine-tuning the odds until something breaks." Journal of Management Studies **125**: 319-340.
- Stephenson, A. (1999). Mars Climate Orbiter Mishap Investigation Board Report. Washington, DC, NASA.
- Sterman, J. D. (1985). "A Behavioral Model of the Economic Long Wave." Journal of Economic Behavior and Organization **6**(1): 17-53.
- Sterman, J. D. (1986). "The Economic Long Wave: Theory and Evidence." System Dynamics Review **2**(2): 87-125.
- Sterman, J. D. (1989). "Misperceptions of Feedback in Dynamic Decision Making." Organizational Behavior and Human Decision Processes **43**(3): 301-335.
- Sterman, J. D. (1989). "Modeling managerial behavior: misperceptions of feedback in a dynamic decision making experiment." Management Science **35**(3): 321-339.

- Sterman, J. D. (2000). Business Dynamics: Systems Thinking and Modeling for a Complex World. Boston, MA, Irwin McGraw-Hill.
- Sterman, J. D., Repenning, N., Kofman, F. (1997). "Unanticipated Side Effects of Successful Quality Improvement Programs: Exploring a paradox of organizational improvement." Management Science **43**(4): 503-521.
- Strauss, A., J. Corbin (1994). Grounded theory methodology: An overview. Handbook of Qualitative Research. N. K. D. a. Y. S. Lincoln. Thousand Oaks, CA, Sage: 273–285.
- Suokas, J., Veikko Rouhiainen (1993). Quality Management of Safety and Risk Analysis, Elsevier Science Publishers.
- Turner, B. A. (1978). Man-Made Disasters. London, UK, Wykeham Publications Ltd.
- Vaughan, D. (1996). The Challenger launch decision: risky technology, culture, and deviance at NASA. Chicago, IL, University of Chicago Press.
- Vicente, K. J. (1999). Cognitive work analysis : toward safe, productive and healthy computer-based work Mahwah, NJ, Lawrence Erlbaum Associates.
- Von Bertalanffy, L. (1968). General system theory. New York, NY, George Braziller.
- Wasserman, S., & Faust, K. (1994). Social Networks Analysis: Methods and Applications. Cambridge, UK, Cambridge University Press.
- Weaver, W. (1958). A Quarter Century in the Natural Sciences. The Rockefeller Foundation Annual Report, Rockefeller Foundation.
- Weick, K. E. (1987). "Organizational Culture as a Source of High Reliability." California Management Review(Winter): 112-117.
- Weick, K. E., K. Sutcliffe, D. Obstfeld (1999). "Organizing for High Reliability." Research in Organizational Behavior **21**: 81–123.
- Weick, K. E., Karlene H. Roberts (1993). "Collective Mind in Organizations: Heedful Interrelating on Flight Decks." Administrative Science Quarterly **38**(3): 357–381.
- Weinberg, G. M. (1975). An Introduction to General Systems Theory. Toronto, Canada, John Wiley and Sons.
- Wolstenholme, E. F. (1990). System Enquiry - A System Dynamics Approach., John Wiley.
- Wolstenholme, E. F. (2003). "Toward the Definition and Use of a Core Set of Archetypal Structures in System Dynamics." System Dynamics Review **19**(1): 7-26.

Woods, D. D., Richard I. Cook (2002). "Nine Steps to Move Forward from Error." Cognition Technology and Work 4: 137-144.

Young, T. (2000). Mars Program Independent Investigation Board Report. Washington, DC, NASA.