Proceedings of the 2005 Winter Simulation Conference
M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, eds.

USING SYSTEM DYNAMICS FOR SAFETY AND RISK MANAGEMENT IN COMPLEX ENGINEERING SYSTEMS

Nicolas Dulac
Nancy Leveson
David Zipkin
Stephen Friedenthal
Joel Cutcher-Gershenfeld
John Carroll
Betty Barrett

Massachusetts Institute of Technology
Cambridge, MA 02139, U.S.A.

ABSTRACT

This paper presents a new approach to modeling and analyzing organizational culture, particularly safety culture. We have been experimentally applying it to the NASA manned space program as part of our goal to create a powerful new approach to risk management in complex systems. We describe the approach and give sample results of its applications to understand the factors involved in the Columbia accident and to perform a risk analysis of the new Independent Technical Authority (ITA) structure for NASA, which was introduced to improve safety-related decision-making.

1 THE PROBLEM

Traditionally accidents are treated as resulting from an initiating (root cause) event in a chain of directly related failure events. This traditional approach, however, has limited applicability to complex systems, where interactions among components, none of which may have failed, often lead to accidents. The chain-of-events model also does not include the systemic factors in accidents such as safety culture and flawed decision-making. Technical risk management requires more than simply looking at the technical parts of systems. A new, more inclusive approach is needed that encompasses the technical aspects, as well as the managerial, organizational, social, and political aspects of the system and its environment.

To accomplish this goal, we use a new foundational model of accident causation plus formal modeling and analysis of both the physical and organizational aspects of systems. The modeling involves various types of executable and analyzable models, including system dynamics models.

Our approach rests on the hypothesis that safety culture can be modeled, formally analyzed, and engineered. Models of the organizational safety control structure and dynamic decision-making and review processes can potentially be used for: (1) designing and validating improvements to the risk management and safety culture; (2) evaluating and analyzing risk; (3) detecting when risk is increasing to unacceptable levels (a virtual "canary in the coal mine"); (4) evaluating the potential impact of changes and policy decisions on risk; (5) performing "root cause" (perhaps better labeled as systemic factors or causal dynamics) analysis; and (6) determining the information each decision-maker needs to manage risk effectively and the communication requirements for coordinated decision-making across large projects.

2 INTRODUCTION TO STAMP

STAMP (System Theoretic Accident Model and Processes) views accidents as the result of flawed processes involving interactions among people, societal and organizational structures, engineering activities, and physical system components (Leveson 2004). 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. In the Space Shuttle Challenger loss, for example, the O-rings did not adequately control propellant gas release by sealing a tiny gap in the field joint. In the Mars Polar Lander loss, the software did not adequately control the descent speed of the spacecraft—it misinterpreted noise from a Hall effect sensor as an indication the spacecraft had reached the surface of the planet.

Accidents such as these, involving engineering design errors, may in turn stem from inadequate control over the

development process, i.e., risk is not adequately managed in the design, implementation, and manufacturing processes. Control is also imposed by the management functions in an organization—the Challenger accident involved inadequate controls in the launch-decision process, for example—and by the social and political system within which the organization exists. The role of all of these factors must be considered in hazard and risk analysis.

Note that the use of the term "control" does not imply a strict military-style command and control structure. Behavior is controlled or influenced not only by direct management intervention, but also indirectly by policies, procedures, shared values, and other aspects of the organizational culture. All behavior is influenced and at least partially "controlled" by the social and organizational context in which the behavior occurs. The connotation of control systems in engineering centers on the concept of feedback and adjustment (such as in a thermostat), which is an important part of the way we use the term here. Engineering this context can be an effective way of creating and changing a safety culture.

Systems are viewed in STAMP as interrelated components that are kept in a state of dynamic equilibrium by feedback loops of information and control. A system is not treated as a static design, but as a dynamic process that is continually adapting to achieve its ends and to react to changes in itself and its environment. The original design must not only enforce appropriate constraints on behavior to ensure safe operation, but must continue to operate safely as changes and adaptations occur over time. Accidents, then, are considered to result from dysfunctional interactions among the system components (including both the physical system components and the organizational and human components) that violate the system safety constraints. The process leading up to an accident can be described in terms of an adaptive feedback function that fails to maintain safety as performance changes over time to meet a complex set of goals and values. The accident or loss itself results not simply from component failure (which is treated as a symptom of the problems) but from inadequate control of safety-related constraints on the development, design, construction, and operation of the socio-technical system.

While events reflect the effects of dysfunctional interactions and inadequate enforcement of safety constraints, the inadequate control itself is only indirectly reflected by the events—the events are the result of the inadequate control. The system control structure, therefore, must be examined to determine how unsafe events might occur and if the controls are adequate to maintain the required constraints on safe behavior.

A STAMP modeling and analysis effort involves creating a model of the system safety control structure: the safety requirements and constraints that each component (both technical and organizational) is responsible for maintaining; controls and feedback channels, process models representing the view of the controlled process by those controlling it, and a model of the dynamics and pressures that can lead to degradation of this structure over time. These models and the analysis procedures defined for them can be used (1) to investigate accidents and incidents to determine the role played by the different components of the safety control structure and learn how to prevent related accidents in the future, (2) to proactively perform hazard analysis and design to reduce risk throughout the life of the system, and (3) to support a continuous risk management program where risk is monitored and controlled.

STAMP uses two types of models: static control structure models and dynamic behavior (system dynamics) models. A new aspect with respect to system dynamics is the typing together of system dynamics and static structure modeling. We believe this will lead to easier construction of system dynamics models and to more complete or expanded modeling and system understanding. The following section provides a short description of our first attempts at using system dynamics to model safety culture and decision-making in the NASA space shuttle program.

3 MODELING SAFETY CULTURE AND DECISION-MAKING AT NASA

3.1 Initial High-Level Diagram

Our interest in using system dynamics as an integral part of a STAMP Risk Analysis process started with a very simple causal loop diagram drawn in the aftermath of the Columbia accident. The objective was to describe the fundamental system adaptation modes responsible for the erosion of safety at NASA (See Figure 1).

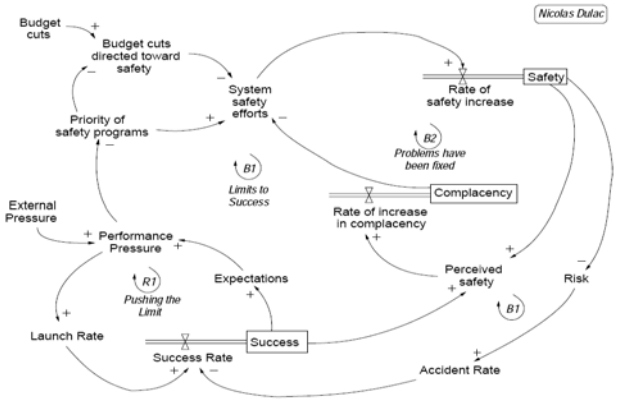

Figure 1: Simplified Model of the Dynamics Behind the Shuttle Columbia Loss

The reinforcing feedback loop labeled R1 or Pushing the Limit, shows how as external pressures increased, per-

formance 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, of course, is an unstable system and cannot be maintained indefinitely—note the larger balancing loop, B1, in which this loop is embedded, labeled Limits to Success. The upper left loop represents part of the safety program. The external influences of budget cuts and increasing performance pressures that reduced the priority of safety procedures led to a decrease in system safety efforts. The combination of this decrease along with loop B2, in which fixing problems increased complacency, which also contributed to reduction of system safety efforts, eventually led to a situation of (unrecognized) high risk. While reduction in safety efforts and lower prioritization of safety concerns may lead to accidents, accidents usually do not occur for a considerable time period (years) so false confidence is created that the reductions are having no impact on safety and therefore pressures increase to reduce safety efforts and priority even further as the external performance pressures mount.

A simple system dynamics model was created out of the causal loop diagram. Model analysis indicated an inherently oscillating behavior where risk is allowed to creep up undetected (see Figure 2) as safety efforts diminish under safety budget cuts and increasing complacency associated with a program perceived to be safe and operational. The major counter-intuitive finding associated with the initial model was that safety should be examined carefully at the when the program seems to be highly successful. The model generated a lot of enthusiasm at NASA colloquiums [more on this], but its simplicity limited the quality of insights that could be extracted from its analysis. Consequently, we decided to create a more complete system dynamics model of the NASA Space Shuttle safety decision-making.

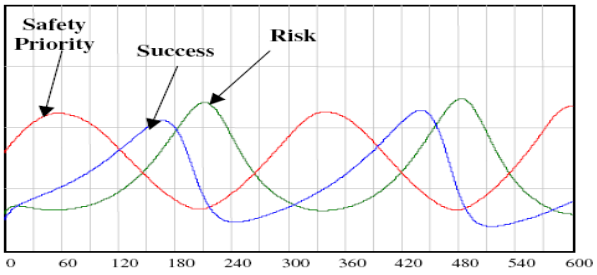

Figure 2: Oscillating Dynamics from Simplified Model

A simple system dynamics model was created out of the causal loop diagram. Model analysis indicated an inherently oscillating behavior where risk is allowed to creep up undetected (see Figure 2) as safety efforts diminish under safety budget cuts and increasing complacency associated with a program perceived to be safe and operational. The major counter-intuitive finding associated with the initial model was that safety should be examined carefully at the point when the program seems to be highly successful. While the model helps to understand the accident, its simplicity limits the quality of insights that can be extracted from its analysis. Therefore, we decided to create a more complete system dynamics model of the NASA Space Shuttle safety decision-making.

3.2 Detailed Model of Safety Decision-Making in the NASA Manned Space Program

The larger model was created to understand the factors in the Shuttle safety culture and decision-making that contributed to the Columbia loss. The original model was constructed using both Leveson's personal long-term association with NASA as well as interviews with current and former employees, books on NASA's safety culture, such as *Inside NASA* (McCurdy 1994), books on the Challenger and Columbia accidents, NASA mishap reports including: CAIB (Gehman 2003), Mars Polar Lander (Young 2000), Mars Climate Orbiter (Stephenson 1999), WIRE (Branscome 1999), SOHO (NASA/ESA 1998), Huygens (Link 2000), other NASA reports on the manned space program such as SIAT (McDonald 2000) and others, as well as many of the better researched magazine and newspaper articles. A detailed documentation of the original model cannot be provided in this paper, but is available upon request from the author. The initial results from our modeling efforts provided some interesting insights and reinforced our belief that system dynamics modeling should be an integral part of a STAMP analysis. Among the scenarios investigated, a contractor analysis was performed to understand the effect of different levels of contracting on system risk. We found that increased contracting did not significantly change the level of risk until a "tipping point" is reached where NASA was not able to perform the integration and safety oversight that is their responsibility. After that point, risk escalates substantially (see Figure 3).

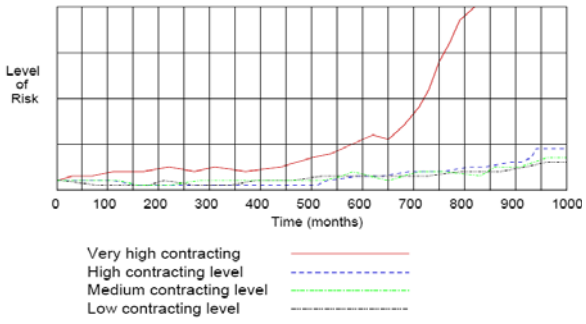

Figure 3: Contractor Scenario Analysis Results

Another scenario investigated the impact on the model behavior of increasing the independence of safety decision

makers through an organizational change like the Independent Technical Authority (ITA). This analysis approximated the effect of the ITA by modifying parameters in the system such as: better reporting, better safety reviews, and more power and authority to safety decision-makers. The results show that significantly higher risk mitigation potential could be achieved by a successful implementation of the ITA program (Figure 4).

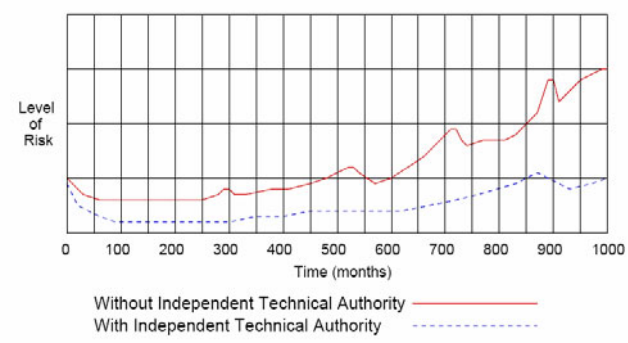

Figure 4: Simplified ITA Scenario Analysis Results

Based on this first attempt at performing detailed modeling of safety culture and decision-making in the manned space program, we were asked by NASA to assist in a planned assessment of the new ITA by using our modeling approach to identify metrics and measures of effectiveness for the assessment. To accomplish this goal, we modified the original model to include a structure that better captures the effects of the ITA program. The objective was to perform a structured analysis of the risks associated with the implementation of the new NASA ITA program. The model was created based on information we obtained from the ITA Implementation Plan and our personal experiences at NASA. The remaining of this paper discusses the entire risk analysis process, as well as the system dynamics model, analysis and results using the ITA program implementation risk analysis as an example.

4 THE STAMP-BASED RISK ANALYSIS PROCESS

We followed a traditional system engineering and system safety engineering approach (see Figure 5), but adapted to the task at hand (organizational risk analysis).

The first step in a STAMP-based risk analysis is to identify the high-level hazard(s) independent technical authority was designed to control and then the general requirements and constraints necessary to eliminate that hazard(s). For ITA:

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

System Safety Requirements and Constraints:

1. Safety considerations must be first and foremost in technical decision-making.
2. Safety-related technical decision-making must be done by eminently qualified experts, with broad participation of the full workforce.
3. Technical decision-making must be credible (executed using credible personnel, with safety analyses available and used throughout the system lifecycle).
4. The Agency must provide avenues for the full expression of technical conscience (for safety-related technical concerns) and provide a process for full and adequate resolution of technical conflicts as well as conflicts between programmatic and technical concerns.

Each of these high-level requirements was then refined into more detailed requirements.

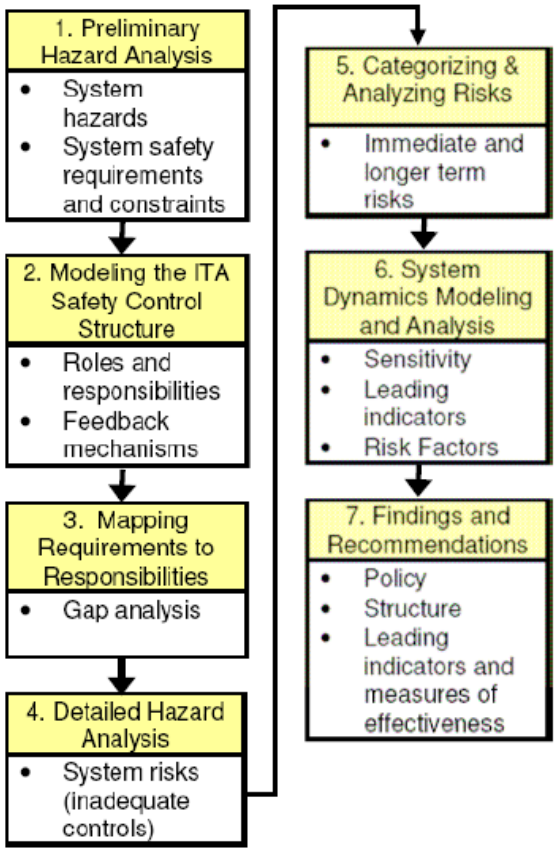

Figure 5: STAMP-Based Risk Analysis Process

The next step was to create a structural model of the safety control structure in the NASA manned space program, augmented with the independent technical authority

as designed. This model includes the roles and responsibilities of each organizational component with respect to safety. We then traced each of the above system safety requirements and constraints to those components responsible for their implementation and enforcement. In this process, we identified some omissions in the organizational design and places where overlapping control responsibilities could lead to conflicts or require careful coordination and communication.

We next performed a hazard analysis on the safety control structure, using a new hazard analysis technique based on STAMP. A STAMP hazard analysis works on both the technical (physical) and the organizational (social) aspects of systems. There are four general types of risks in the ITA concept:

1. Unsafe decisions are made 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 they do not have adequate impact on system design, construction, and operation.

The hazard analysis applied each of these types of risks to the NASA organizational components and functions involved in safety-related decision-making and identified the risks (inadequate control) associated with each. The resulting list of risks was quite long (250), but most appeared to be important and not easily dismissed. To reduce the list to one that could be feasibly assessed, we categorized each risk as either an immediate and substantial concern, a longer-term concern, or capable of being handled through standard processes and not needing a special assessment.

We then used our system dynamics models to identify which risks were the most important to measure and assess, i.e., which provide the best measure of the current level of system risk and are the most likely to detect increasing risk early enough to prevent significant losses. This analysis led to a list of the best leading indicators of increasing and unacceptable risk.

The analysis also pointed to structural changes and planned evolution of the safety-related decision-making structure over time that could strengthen the efforts to avoid migration to unacceptable levels of organizational risk and avoid flawed management and engineering decision-making leading to an accident. The following section provides a description of the contribution of system dynamics modeling to the entire ITA risk analysis.

5 DYNAMIC RISK ANALYSIS OF THE INDEPENDENT TECHNICAL AUTHORITY

5.1 Model Description

One of the significant challenges associated with modeling a socio-technical system as complex as the Shuttle program is creating a model that captures the critical intricacies of the real-life system, but is not so complex that it cannot be readily understood. To be accepted and therefore useful to decision makers, a model must have the confidence of the users and that confidence will be limited if the users cannot understand what has been modeled. We addressed this problem by breaking the overall system dynamics model into nine logical subsystem models, each of an intellectually manageable size and complexity (see Figure 6).

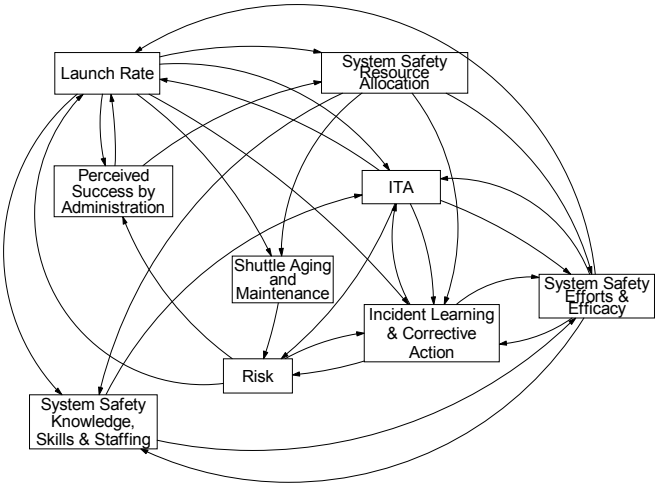

Figure 6: The Nine Subsystem Models and their Interactions

The subsystem models were built and tested independently. Extensive partial model testing was used in order to increase our confidence that the model behavior would be accurate. It was also verified that the behavior of each subsystem module passed the intent rationality test (Morecroft 1985, Sterman 2000). The behavior of each subsystem model was shown to be in accordance with the open-loop behavior rationally expected when critical feedback loops are removed. For example, in the absence of external pressures to modify the resources allocated to safety efforts (e.g., schedule and budget pressures), the System Safety Resource Allocation model should output a constant level of safety resources. Once validation and confidence in the behavior of each subsystem model was established, subsystem models were connected to one another so that important information could flow between them and emergent properties arising from their interactions could be included in the analysis. The model was built in a modular

fashion, which made it easy to test and modify individual subsystem models independently, and then re-integrate them.

The following description provides a high-level listing of some key variables and concepts contained in each subsystem model. A detailed description of the content of each model is impossible given paper length constraints, however, interested readers are invited to request it from the author.

Risk: Incident and accident occurrence, effective vehicle age, quantity and quality of inspections, proactive hazard analysis and mitigation efforts, response of the program to anomalies (symptom fix vs. systemic factor fix response).

System Safety Resource Allocation: Level of resources allocated to system safety, priority of safety program, priority of launch performance, NASA safety history, performance expectations, schedule pressure, budget pressure.

System Safety Knowledge, Skills, and Staffing: NASA and contractors' system safety knowledge and skills, ability to oversee contractor safety activities, number of NASA system safety employees, number of contractor system safety employees, aggregate experience of NASA employees, aggregate experience of contractor employees, age of NASA employees, portion of work contracted out, stability of funding, hiring rate, attrition rate, experience at hire, learning rate.

Shuttle Aging and Maintenance: Age of the shuttle vehicles (in launches), amount of maintenance, refurbishments, and safety upgrades, resources available for maintenance, maintenance requirements, original design lifetime, uncertainty in remaining system life.

Launch Rate: Perception of success by management, performance expectations from management, schedule pressure, launch commitment, launch backlog, launch delays.

System Safety Efforts and Efficacy: Availability and adequacy of system safety resources, availability and effectiveness of safety processes and standards, system safety staff characteristics (number, knowledge, experience, skills, motivation, and commitment), ability of NASA to oversee and integrate contractor safety efforts, quantity and quality of lessons learned.

Incident Learning and Corrective Action: Number of safety-related incidents, fraction of safety problems reported depending on the effectiveness of the reporting process, employee sensitization to safety problems, fear of reporting problems and concerns, risk perceived by engineers and technical workers, fraction of safety problems investigated, thoroughness of investigation process, fraction of problems resulting in no action, fraction of corrective actions that only address the symptoms of the problem, fraction of corrective actions that address the systemic factors that led to the problem, waiver issuance rate, fraction of corrective actions rejected at safety review, quality of lessons learned.

Perceived Success by Management: Accumulation of successful launches, NASA recent safety history, occurrence of serious events and accidents.

Independent Technical Authority: Effectiveness and Credibility of ITA, quality and thoroughness of safety analyses, workload of ITA designees, attractiveness of being a Technical Warrant Holder (TWH), TWH resources and training, ability to attract knowledgeable trusted agents, trusted agent training adequacy, ITA influence and prestige, ability to attract highly skilled and well-respected technical leaders, ITA power and authority.

5.2 Model Analysis

Once the models were thoroughly tested, three types of analyses were performed: (1) sensitivity analyses to investigate the impact of various ITA program parameters on the system dynamics and on risk, (2) system behavior mode analyses, and (3) metrics identification and evaluation.

5.2.1 ITA Model Sensitivity Analysis

In order to investigate the effect of ITA parameters on the system-level dynamics, a 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 7 and 8 show the results of the 200 individual traces, for the variables *ITA Effectiveness and Credibility* and *System Technical Risk*.

The initial sensitivity analysis shows that at least two qualitatively different system behavior modes can occur. The first behavior mode (behavior mode #1 in Figure 7) is representative of a successful ITA program implementation where risk is adequately mitigated for a relatively long period of time (behavior mode #1 in Figure 8). More than 75% of the runs fall in that category. The second behavior mode (behavior mode #2 in Figure 7) 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 (behavior mode #2 in Figure 8).

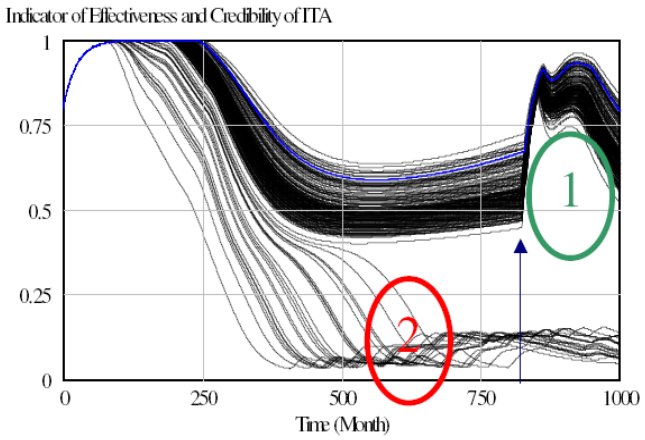

Figure 7: Sensitivity Results for Effectiveness and Credibility of ITA

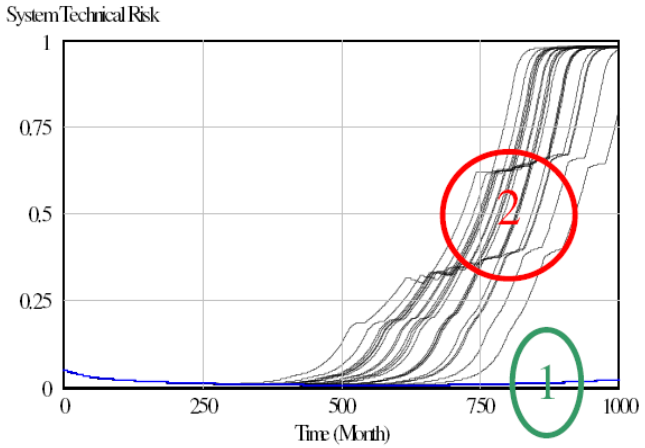

Figure 8: Sensitivity Results for System Technical Risk

5.2.2 System 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 parameters involved. Using this information, we were able to identify some potential metrics and indicators of increasing risk as well as potential risk mitigation strategies.

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.

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 month 850 (denoted by the arrow on the x-axis of Figure 9), 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 9 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. This behavior mode is characterized by an extended period of nearly steady-state equilibrium where risk remains at very low levels.

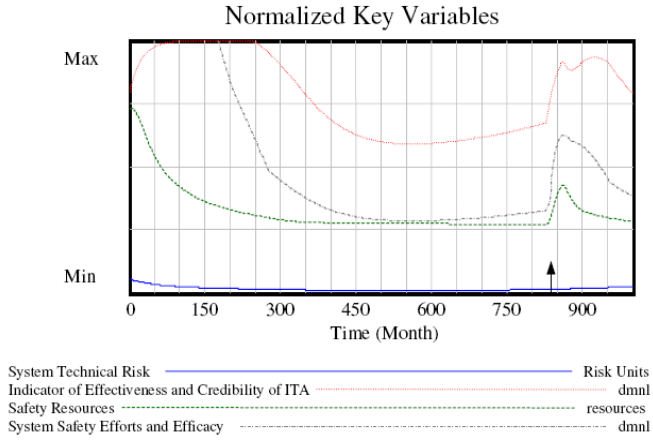

Figure 9: Key Variables for Behavior Mode #1

Behavior Mode #2: Unsuccessful ITA Implementation: In the second behavior mode (behavior mode #2 in Figure 7), 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 dynamics act in the negative direction 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 10).

The key normalized variables for a sample simulation run representative of the second behavior mode are shown in Figure 10. 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 increases System Safety Efforts and Efficacy, but by 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 start 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 accidents 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.

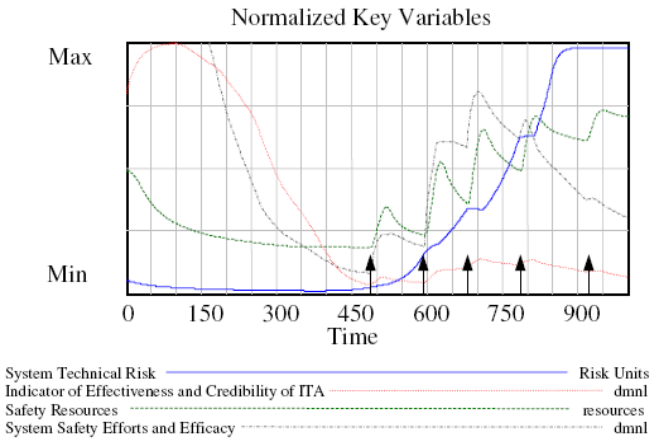

Figure 10: Key Variables for Behavior Mode #2

5.2.3 Metrics Identification and Evaluation

Our models indicate that many good indicators of increasing risk are available. However, many of these indicators become useful only after a significant risk increase has occurred, i.e., they are lagging rather than leading indicators. The requirements waiver accumulation pattern, for example, is a good indicator, but only becomes significant when risk starts to rapidly increase (Figure 11), thus casting doubt on its usefulness as an effective early warning.

Alternatively, the number of incidents/problems under ITA investigation appears to be a more responsive measure of the system heading toward a state of higher risk (see Figure 12). A large number of incidents under investigation results in a high workload for trusted agents, who are already busy working on project-related tasks. Initially, the dynamics are balancing, as ITA personnel are able to increase their incident 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.

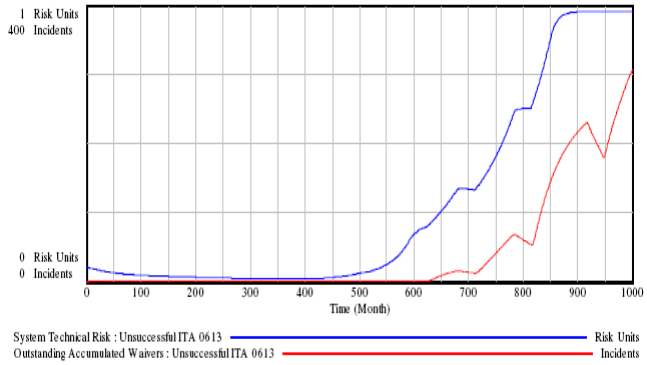

Figure 11: Risk and Requirement Waivers Accumulation

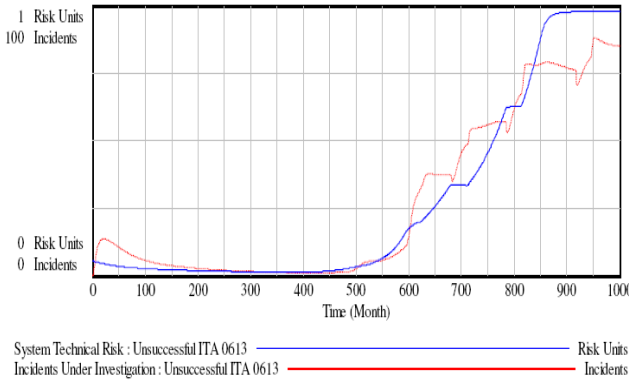

Figure 12: Risk and Incidents/Problems under Investigation

A bottleneck effect is created that makes things worse through a fast acting, negative-polarity reinforcing loop (see Figure 13). This potential bottleneck points to the utility of more distributed technical decision-making.

Using the number of problems being worked is not without its own limitations. For a variety of reasons, the technical warrant holders may simply not be getting information about existing problems. Independent metrics (e.g., using the PRACA database) may have increased accuracy here. It is unlikely that a single metric will provide the information required—a combination of complementary metrics are almost surely going to be required.

Because of its deep structural impact on the system, the health of ITA may be the most effective early indicator of increasing risk. There is a high correlation between the Effectiveness and Credibility of ITA and the location of the tipping point at which risk starts to rapidly increase. However, because the Effectiveness and Credibility of ITA cannot be measured directly, we must seek proxy measures of ITA health. One of the most promising leading indicators of ITA health is the ability to continually recruit the

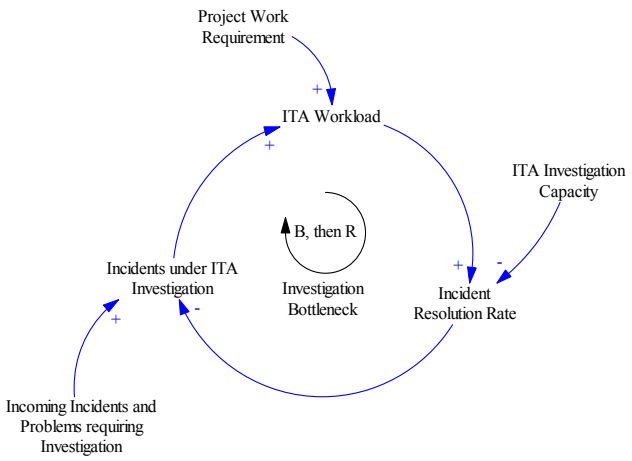

Figure 13: Balancing Loop Becomes Reinforcing as ITA Workload Keeps Increasing

"best of the best". Employees in the organization have an acute perception of which assignments are regarded as prestigious and important. As long as ITA is able to retain its influence, prestige, power and credibility, it should be able to attract the best, highly experienced technical personnel with leadership qualities. By monitoring the quality of ITA personnel (Technical Warrant Holders and Trusted Agents) over time along with turnover and job application data, it should be possible to have a good indication of ITA health and to correct the situation before risk starts to increase rapidly.

In addition to using system dynamics as a tool to identify and evaluate metrics and leading indicators of safety drift, we used the models created in order to perform a first order assessment of the risks identified in the previous steps of the STAMP Based Risk Analysis Process. The list of risks identified included 250 items, approximately 75% of which were related to variables in the system dynamics models. This correlation between risks and model variables allowed us to prioritize risks according to their sensitivity to other model parameters. In order to determine the sensitivity of specific variables, a sensitivity analysis simulation was performed that covered a range of cases including cases where the ITA is highly successful and self-sustained, and cases where the ITA quickly loses its effectiveness. A Low, Medium or High sensitivity rating was assigned depending on the normalized variation percentage of specific model variables during the sensitivity analysis. Figure 7 provides an example of a variable with a high variation to model parameters due to the reinforcing ITA dynamics described above. Figure 14 provides an example of a variable with lower sensitivity to model parameters. This variable provides a measure of the shuttle age relative to its design lifetime. The effective shuttle age is higher at the end of the system lifecycle if the ITA program is successful because risk has been effectively mitigated throughout the system lifecycle, resulting in higher launch rates, especially in the second half of the system life.

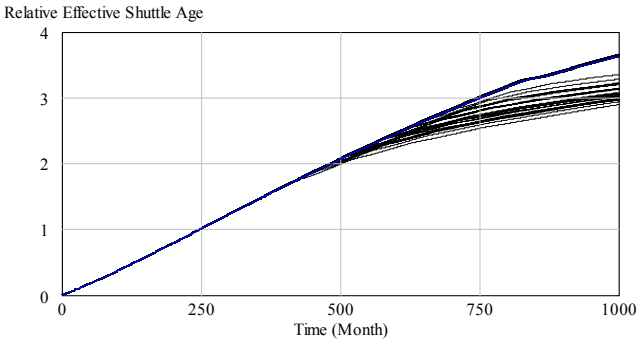

Figure 14: Variable Exhibiting Low Sensitivity to Model Parameters

6 CONCLUSION: USING SYSTEM DYNAMICS IN SAFETY AND RISK ANALYSIS

The objective of the system dynamics part of a STAMP-Based Risk Analysis to provide insight into the dynamic reasons for the adaptation of the safety control structure, or the drift toward an unsafe system state. Once these adaptation mechanisms are better understood, the system dynamics models can be used to help design monitoring systems that will act as a virtual "canary-in-the-mine", alerting decision-makers that the system has reached, or is heading toward an unsafe state.

During the ITA risk analysis process, the system dynamics models contributed many insights that may not have been identified using the static safety control structure alone. For example, the requirements waiver accumulation is often considered a sign of risk increase, but the dynamic modeling provided hints that it may be a lagging indicator with limited effectiveness for early warning. The analysis provided many other candidate indicators (incidents under investigation, quality of ITA designees, …) that could be used as leading indicators of safety drift. More work will be required to evaluate the effectiveness of these indicators on the real system, but the model analysis pointed to candidate indicators that may not have been identified otherwise.

Another interesting insight from the system dynamics analysis is that the performance of ITA is highly dependent on the quality of safety analyses produced by system safety employees at NASA and contractor offices.

In addition, while Technical Warrant Holders (TWHs) are shielded in the design of the ITA from programmatic budget and schedule pressures through independent management chains and budgets, Trusted Agents are not. They have dual responsibility for working both on the project and on TWH assignments, which can lead to obvious conflicts. Good information is key to good decision-making. Having that information produced by employees not under

the ITA umbrella reduces the effective independence of ITA. In addition to conflicts of interest, increases in Trusted Agent workload due either to project and/or TWH assignments or other programmatic pressures can reduce their sensitivity to safety problems. A long list of similar insights was generated from the system dynamics part of the STAMP risk analysis.

The results of our analysis are very encouraging and illustrate the potential for a STAMP-based risk analysis process augmented with system dynamics modeling to significantly improve risk management in complex socio-technical systems.

7 FUTURE WORK

While the process for generating the safety control structure part of the STAMP risk analysis is mature and well documented, the process for creating and using system dynamics models based on it currently requires much effort and domain expertise. Future work will address the creation of system dynamics models based on the STAMP safety control structure and existing system safety archetypes. The model validation (or confidence increase) and insight generation processes will also be addressed in greater detail.

ACKNOWLEDGMENTS

This research was partially supported by a grant from the USRA Center for Program/Project Management Research (CPMR), which is funded by NASA APPL and by the ITA Program within the NASA Chief Engineer's Office.

REFERENCES

Branscome, D.R. (Chair). 1999. WIRE Mishap investigation board report, NASA.

Gehman, Harold (Chair). 2003. Columbia accident investigation report.

Leveson, Nancy. 2004. A new accident model for engineering safer systems. *Safety Science* 42 (4): 237–270.

Link, D.C.R. 2000. Report of the Huygens communications system inquiry board, NASA.

McCurdy, Howard. 1994. *Inside NASA: High technology and organizational change in the U.S. space program.* Johns Hopkins University Press.

McDonald, Harry. 2000. Shuttle independent assessment team (SIAT) report.

Morecroft, J.D.W. 1985. Rationality in the analysis of behavioral simulation models. *Management Science* 31 (7): 900-916.

NASA/ESA investigation board. 1998. SOHO Mission Interruption, NASA.

Stephenson, A. (Chair). 1999. Mars Climate Orbiter mishap investigation board report, NASA.

Sterman, John. 2000. *Business Dynamics: Systems thinking and modeling for a complex world,* McGraw-Hill.

Young, Tom (Chair). 2000. Mars program independent investigation board report, NASA.

AUTHOR BIOGRAPHIES

NICOLAS DULAC is a doctoral candidate in the department of Aeronautics and Astronautics at MIT. His current research interests include system safety, system engineering, visualization of complex systems, hazard analysis in socio-technical systems, organizational safety culture, and dynamic risk analysis.

NANCY LEVESON is Professor of Aeronautics and Astronautics and also Professor of Engineering Systems at MIT. She is a member of the National Academy of Engineering. Her research interests include system engineering, system safety, human-computer interaction and software engineering.

DAVID ZIPKIN is a graduate from MIT's Technology and Policy (TPP) Program.

STEPHEN FRIEDENTHAL is a student in MIT's System Design and Management (SDM) Program.

JOEL CUTCHER-GERSHENFELD is a Senior Research Scientist in MIT's Sloan School of Management and Executive Director of MIT's Engineering Systems Learning Center.

JOHN S. CARROLL is Professor of Behavioral and Policy Sciences at the MIT Sloan School of Management and the MIT Engineering Systems Division. His research focuses on the relationships among individual and group decision making and communication, and organizational learning, change, and culture.

BETTY BARRETT is a research scientist and associate director of the Engineering Systems Learning Center in the Engineering Systems Division at MIT.