

## INITIAL DEVELOPMENT AND APPLICATION OF A NEW METHOD TO ANALYZE ORGANIZATIONAL RISK

Heather A. Stoner, Amy L. Alexander, Bruce Skarin, & William J. Salter  
Aptima, Inc.  
Woburn, MA

Large, complex organizations implement a range of policies and procedures to protect against risk. Sometimes, however, catastrophic accidents occur, such as with the Challenger and Columbia shuttles. "Procedural drift" has been suggested as a source of such accidents (Rasmussen, Pejtersen, & Goldstein, 1994), a process whereby procedures gradually shift, based on operating experience, until an extreme situation causes the shifted procedures to fail. Brief case studies were conducted of the shuttle accidents to explore and refine this hypothesis and to develop methods for investigating it systematically, including a "knowledge map" method, initially described by Lintern (2003). A knowledge map is a theoretically-driven variant of a work domain analysis (Rasmussen et. al., 1994), which incorporates procedural task analysis and organizational controls to analyze the overarching goal of the system. This enables visualization of how procedures occur within the organization, describing collaborations, decisions, and management procedures. Secondary goals are achieved through tasks, and tasks are further broken down into sub-tasks to describe how the system, at its procedural level, acts within the global constraints of the system. Since the knowledge map method does not generate runnable models, a system dynamics model was also developed to model the processes and dynamics of procedural drift, particularly as it affects safety. This combination of methods shows considerable promise in analyzing the shuttle accidents in particular and procedural drift and organizational risk in general.

### INTRODUCTION

Large, complex organizations implement a range of policies and procedures to protect against risk. These both provide constraints and induce what we term *organizational affordances*. Organizational affordances allow for certain practices that lay the groundwork for *procedural drift*, or migration toward the boundaries of (explicitly mandated) safe operations over time (Rasmussen, Pejtersen, & Goldstein, 1994). In many of today's organizations local practice drifts away from demands of global constraints. Such drift is usually in response to a combination of operational experiences and on-going pressures (e.g., budgets or schedules), and is potentially a major threat to safety in today's complex, socio-technical systems. This drift can be found at the core of both NASA shuttle accidents, and was given the name of *normalization of deviance* by Vaughn (1996). This term captures the insight that practices and events once considered deviations (from procedures, expectations, or requirements) often become treated as normal over time.

Past and contemporary approaches to safety seek to eliminate the drift generated by local pressures through the use of tight control in the form of rules and procedures (Lintern, 2003). But many complex systems are too large and complex to be accurately defined *a priori*. In this paper we conducted brief case studies of the Challenger and Columbia shuttle accidents, to explore and refine the procedural drift hypothesis and to develop methods for investigating it systematically. Focusing on the Columbia accident, we explored the applicability of two methods: knowledge map and system dynamics (SD) modeling.

### METHOD

Knowledge map, initially described by Lintern (2003) is a theoretically-driven variant of WDA (Rasmussen et. al., 1994). It enables the description of the environment in which the NASA shuttle program was operating without consideration of any given operator or position. Its utility derives from its ability to view the environment in which the organization operates to help determine when normal and abnormal situations occur. A key difference between knowledge map and WDA is that the former allows descriptions of tasks within its structure. While WDA is useful for constraints of a system, it cannot describe the organization or procedures that shift over time. A knowledge map is theoretically driven by a WDA, but includes procedural task analysis and organizational controls. A task analysis (Kirwan & Ainsworth, 1992) is a method that breaks down the mental and physical steps that the operator goes through. Fitting the procedural task analysis into this method enables visualization of how procedures are carried out.

Knowledge map supports navigation of a complex socio-technical system through (summarizing Lintern, 2003):

- Functional clustering of information related to a single issue;
- Representation of meaning at both higher and lower levels of detail;
- Explicit representation of hidden interdependencies; and
- Linking items between levels to reveal the means of

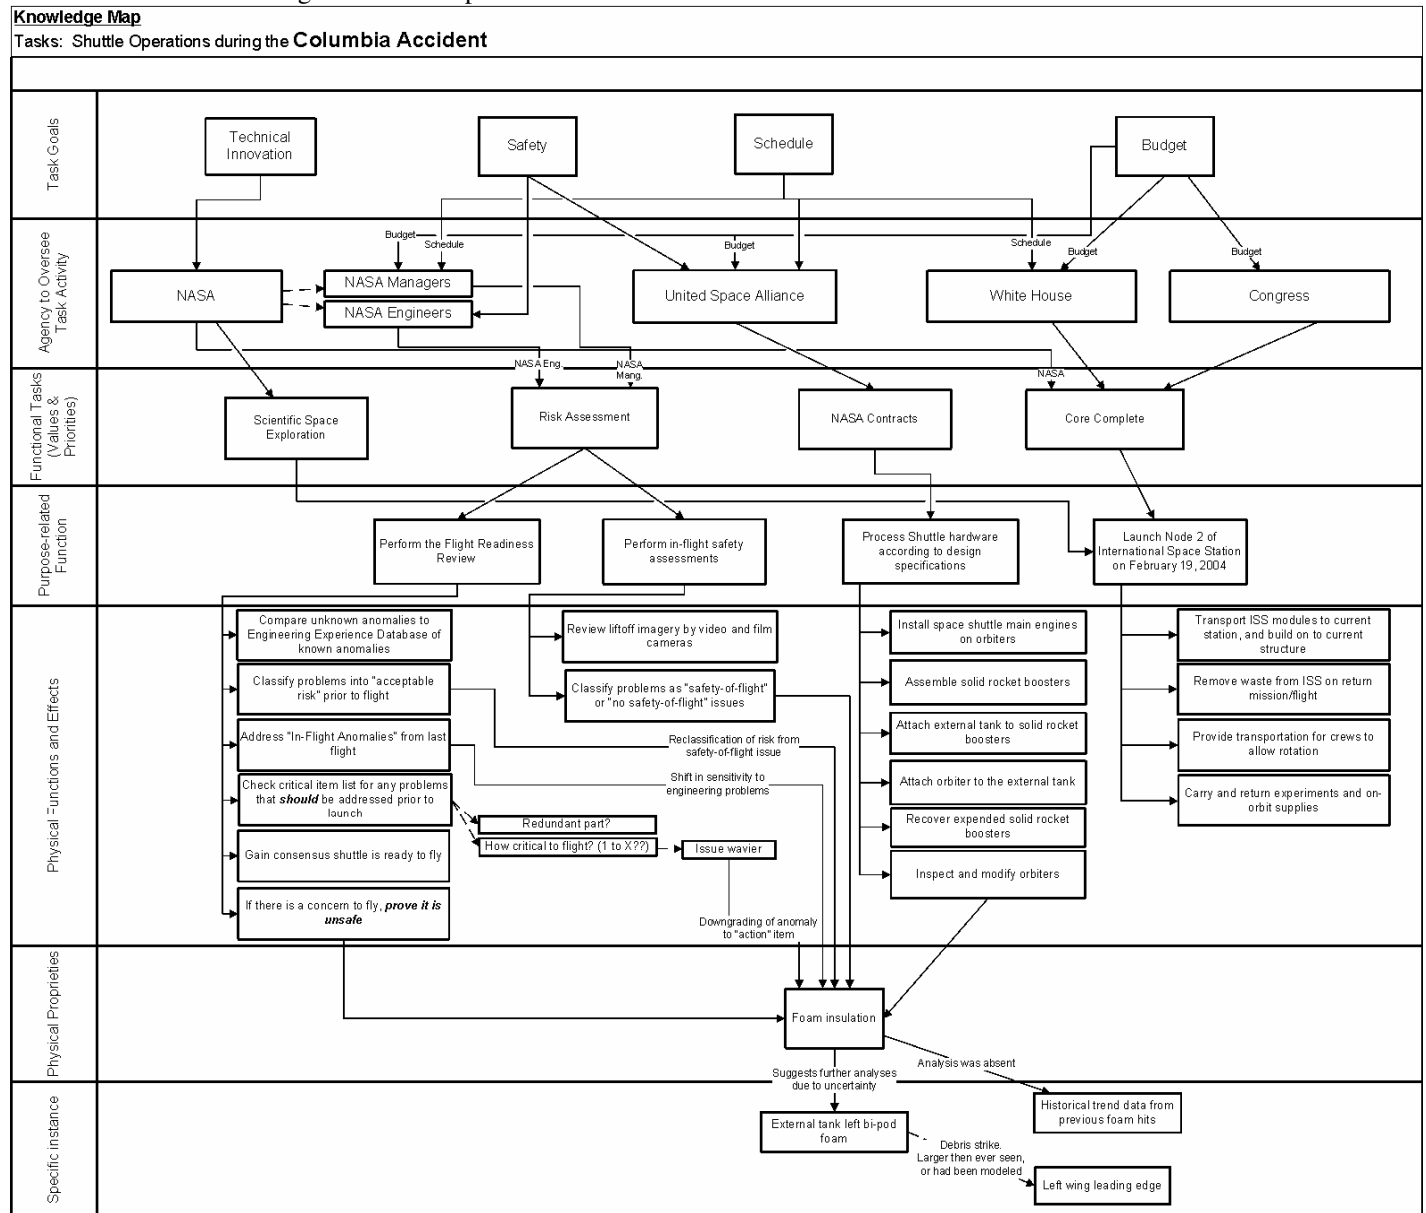
achieving goals.

A knowledge map analyzes the *overarching goal* of the system (i.e., shuttle operations), and breaks it into *secondary goals* according to individual organizations—describing collaborations, decisions, and management procedures. These are achieved through *functional tasks*, and these tasks are further broken down into *sub-tasks* to describe how the system, at its procedural level, acts within the global constraints of the organization. The lower levels describe models of functionality, reliability, and interdependence.

**RESULTS: THE COLUMBIA ACCIDENT**

The Columbia accident was the focus of our work. Information for this effort was primarily derived from the Columbia Accident Investigation Board report issued in

2003. The Columbia accident occurred on February 1, 2003 upon re-entry, although the original physical cause for the accident occurred on January 16, 2003 during ascent. At 81.7 seconds after launch, one large and two smaller pieces of foam separated from the left bipod ramp and struck the wing’s leading edge. Velocity at impact to the wing was approximately 220 m/s. Historical data had shown no other foam instances of this magnitude. Nevertheless, Mission Control decided that there was “no concern for [Reinforced Carbon-Carbon] or [Thermal Protection System] tile damage,” and that there was “absolutely no concern for re-entry.” However, as we know, the orbiter’s skin had been penetrated, exposing the underlying airframe to extreme temperatures, resulting in burn-through and the eventual disintegration of Columbia. Figure 1 illustrates a knowledge map for the Columbia accident.



**Figure 1. The Columbia Accident Knowledge Map.**

The *task goals* (highest level) describe system and global goals. These can be tangible or intangible, and are often difficult to measure. The shuttle program's goals during Columbia included: conducting innovative research, ensuring safety, staying on schedule, and working within a limited budget.

The *overseeing agency* (second level) describes the acting organizations with specific vested interest or control over tasks, goals, or the effects of the system. For Columbia, NASA as an organization was responsible for conducting innovative research. NASA was then decomposed into its managers and engineers. Managers maintained control of the shuttle schedule and budget, while engineers had control over safety. The United Space Alliance (USA), a contractor, was also responsible for safety while maximizing the utility of adherence to budget and schedule constraints.

The *functional tasks* (third level) describe specific goals of the organizations above. These goals need to be maximized or minimized for that organization's success. They often conflict among different organizations. NASA as a whole leads the way in scientific space exploration but also runs various other NASA programs that change over time (e.g., building the International Space Station). NASA managers and engineers assess risks of shuttle operations, contractors perform according to described obligations, and NASA programs like the U.S. Core Complete are carried out by NASA with input from the White House and Congress.

The *purpose-related function* (fourth level) describes higher-level tasks used to achieve the previously-described values and priorities. An example of a function is 'risk assessment,' characterized by the tasks of performing flight readiness and in-flight reviews. The remaining functional tasks are illustrated in Figure 1.

The *physical functions and effects* (fifth level) describe specific tasks in detail, and can be viewed as a hierarchical task analysis. One task evaluated for Columbia was the flight readiness review. It was decomposed into composite tasks of comparing unknown anomalies to the engineering experience database; classifying problems into 'acceptable risk' prior to flight; addressing 'in-flight anomalies' from the last flight; reviewing the critical items list; and proving conditions were unsafe to fly. In-flight safety assessments included reviewing liftoff imagery by video and film cameras, then classifying anomalies as a "safety-of-flight risk" or not.

The *physical properties* (sixth level) describe the grouping of instances' descriptions represent both physical and intentional constraints within the system and its tasks. At this level, "foam insulation" is called out, and further

instantiated at the final levels. Here we focused on the direct known cause of the Columbia accident; however, more could be included.

The *specific instance* (seventh level) describes the physical proprieties that compose the upper levels of abstraction by describing specific instances or objects needed to achieve those tasks. These are specific to the local instance, rather than global.

### System Dynamics Modeling

The SD model simulates the behaviors of procedural drift, illustrating how small deviations from procedures over time, coupled with continued experience of safe events, can incrementally (and unknowingly) lead to dangerous system states. Such a model can allow for the anticipation of risks and potential safety threats inherent to changes in organizational procedures. Within NASA, our model showed non-linear relationships between the major factors of interest; in particular, accumulating safe experience with foam strikes, which drives normalization of deviance (Vaughn, 1996), combined with schedule pressures, moved the system into a high-risk condition. In that high-risk condition, environmental and stochastic factors produced a catastrophic accident. The numeric approach to simulating SD models provides approximations of non-linear equations, allowing for the examination of such complex behavior. To derive an effective solution, plausible explicit assumptions of linked factors are formed, thus providing reasonable means for examining initial behaviors. The utility of SD is that the precise values in any assumption do not have a significant effect on the basic relationships defined within the model.

The purpose of linking a SD model to the knowledge map is to represent the influence over time of the factors identified in the map on the overall safety level of the system (Figure 2). Moreover, the model shows how factors can interact to create greater risk than when viewed in isolation. With the addition of each factor identified and its relations to others (Figure 3), the "risk profile" over time changes significantly. If only design problems are considered, the risk changes very little. The addition of detection methods helps reduce risk, but when the influence of normalization of deviance (Vaughn, 1996) is included, the risk climbs slowly over time. Finally, including the influence of cost and schedule pressure helps to demonstrate how the system can be pushed into an unsafe region despite efforts to control risk.

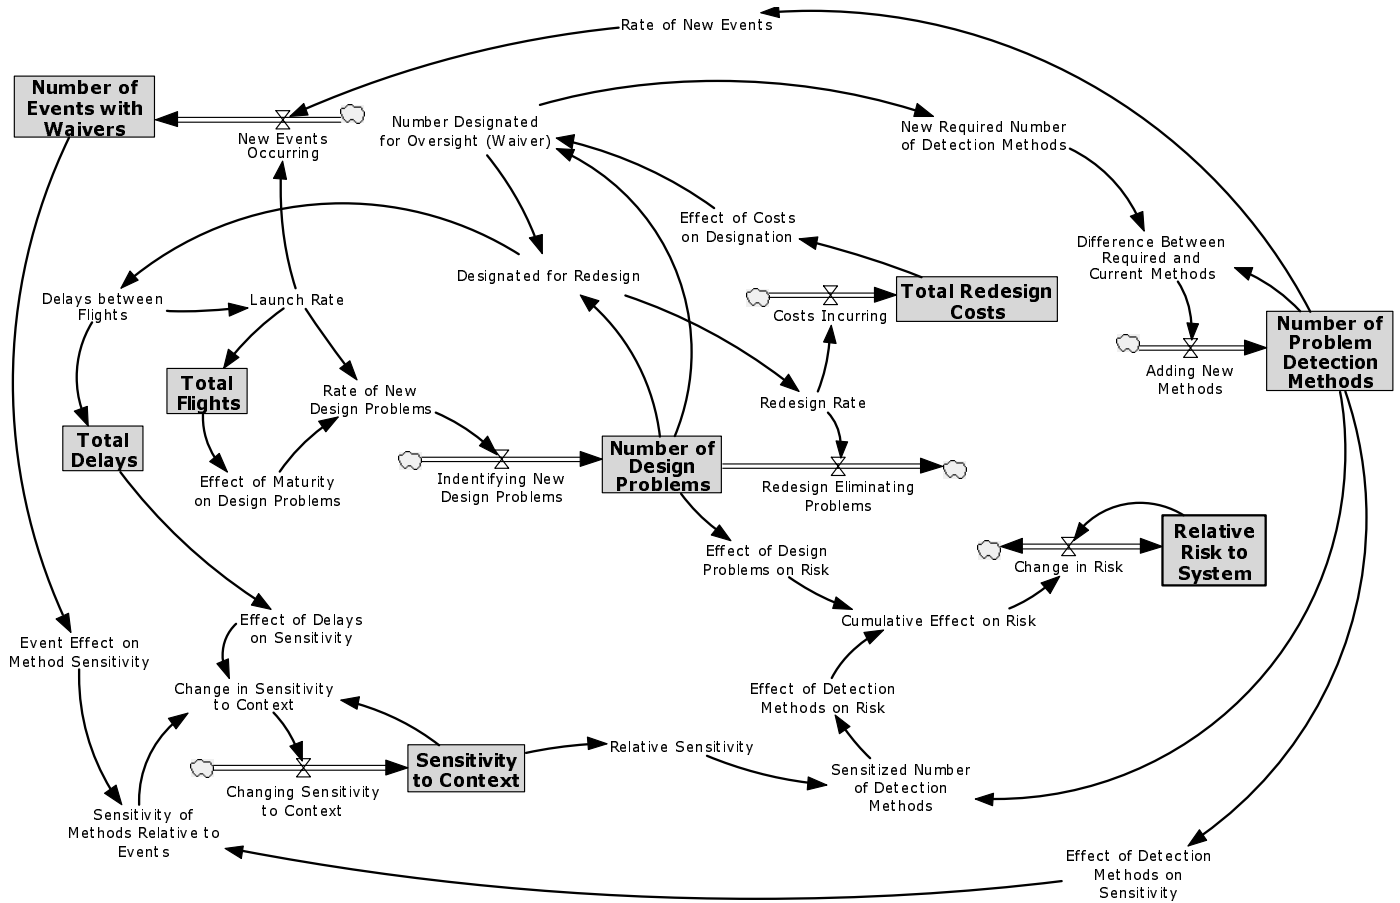


Figure 2. System Dynamics model.

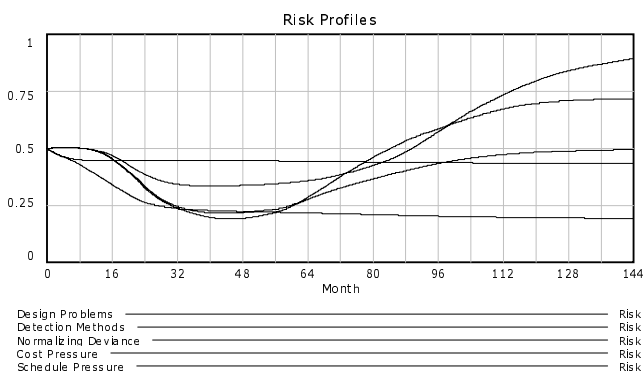


Figure 3. System Dynamics model output for each additional influence factor.

**DISCUSSION**

Prior to both shuttle accidents, the combined factors of normalization of deviance and schedule pressure pushed the boundaries of safety and left NASA operating in realms where unexpected circumstances could lead to catastrophe. In the Challenger case, repeated safe missions that showed o-ring blow-by on return led engineers and managers to discount its importance, lowering its risk classification over

time. When Challenger was launched under extreme environmental conditions (low launch temperature), the potential risk of those conditions on o-ring integrity, and thus on mission safety, was not given sufficient weight. In the Columbia case, years of successful launches with foam shedding led to insufficient concern about the potential danger of the large amount of shedding observed soon after launch.

We represented a simplified version of these dynamics for Columbia in a knowledge map and an SD model. However, both accidents showed the *same dynamics of normalization of deviance and procedural drift* over time. We take this convergence as supporting evidence for the potential utility of the methods described.

Of course, our case studies were relatively narrow in scope and focused on modeling the accident dynamics identified in *post hoc* accident analyses (Columbia Accident Investigation Board, 2003; Vaughn, 1996). Nevertheless, we believe that they show promise, certainly as tools for systematic *post hoc* analysis and, potentially, even as predictive models for *a priori* risk identification.

## CONCLUSIONS

The knowledge map method yielded five suggestive insights into the possible nature and processes behind *normalization of deviance*:

1. **Incompleteness.** No one method can completely describe a complex system. Since the knowledge map method does not describe change over time, an SD model was also developed to model the processes and dynamics of procedural drift. Whether the combination of these methods, more exhaustively applied, can help to prevent accidents in complex organizations of course remains an open question, but we believe further work is appropriate given our results.
2. **Procedural drift occurs over differing temporal scales.** Procedural drift, such as foam strike risk reclassifications, accumulates over long temporal scales (years, in both NASA cases), yet the impacts of any given instance occur on a compressed time scale (less than 24 hours for Challenger, over a week or so for Columbia). Modeling must be sensitive to these differences in temporal granularity. We believe that full integration across these scales may prove quite difficult, but that modeling at the larger scale of procedural drift can serve as an early warning system.
3. **Context matters.** Interaction with the extraneous environment requires rich understanding of the system, of that environment, and of how it supports human performance. This was most clear for the Challenger accident, where (external) changes in ambient temperature proved decisive. But for Columbia, too, differences in context of the foam strike were critical: it was larger and hit a particularly vulnerable place on the shuttle.
4. **Boundary states are critical.** Understanding system boundary states is critical to ensure that the system does not deviate from normal into the unanticipated. The low temperatures and large foam strike in a vulnerable spot mentioned above exemplify this.
5. **Formal signs of drift.** Procedural drift is informal, yet waivers provide a formal means of identifying it. Not *all* instances of drift provide such indicators, but it is vital to attempt to identify such formal means when possible.

The methods outlined above merit further development and application, in three primary directions:

1. **More *post hoc* analyses:** Our methods could be applied to accidents in other domains – industry, commercial aviation, medicine – which would help to determine if the methods have general utility.
2. **Closer fusion of knowledge map and SD:** A tighter coupling of the two distinct methods is required to make this an approach of broader application.
3. **Prediction:** Finally, it may be valuable to analyze an organization as it currently operates to investigate and refine the potential predictive utility of our approach. Its initial *a priori* application should probably be as an early warning heuristic model that suggests where risk may be in an organization and its procedures.

## ACKNOWLEDGMENTS

This research was supported by contract NASA NNA04AA56 for which Dr. Patricia Jones was the Technical Point of Contact. Any opinions, findings, and conclusions or recommendations in this publication are those of the authors and do not necessarily reflect the views of NASA.

## REFERENCES

- Columbia Accident Investigation Board (August, 2003). *Columbia Accident Investigation Board Final Report, Vol. 1*. Washington, DC: National Aeronautics and Space Administration.
- Kirwan, B. & Ainsworth, L. K. (1992). *A guide to task analysis*. Philadelphia, PA: Taylor & Francis.
- Lintern, G. (2003, April 14-17). Tyranny in Rules, Autonomy in Maps: Closing the Safety Management Loop. *Proceedings of the Twelfth International Symposium on Aviation Psychology*. Dayton, Ohio.
- Rasmussen, J., Pejtersen, A. M., & Goodstein, L. P. (1994). *Cognitive systems engineering*. New York: John Wiley & Sons, Inc.
- Vaughn, D. (1996). *The Challenger launch decision: Risky technology, culture, and deviance at NASA*. University Of Chicago Press.