In R. Jensen, (Ed), Proceedings of the Twelfth International Symposium on Aviation Psychology, April 14-17, 2003, Dayton, Ohio, 719-724 [CD ROM]

TYRANNY IN RULES, AUTONOMY IN MAPS: CLOSING THE SAFETY MANAGEMENT LOOP

Gavan Lintern
Aptima Inc
12 Gill St, Suite 1400
Woburn MA 01801, USA
lintern@aptima.com

The development of complex, safety-critical socio-technical systems proceeds through a rational design stage in which an extensive set of procedures is established. However, all major system accidents have, as one contributing factor, a failure of operational personnel to adhere to certain critical procedures. The typical response is to develop more detailed constraints in an attempt to prevent reoccurrence of that sort of accident. This approach exemplifies the worst of rule-based safety management. It is retroactive and fails to recognize the strength of lessons learned by operational personnel in practice. Safety management must embrace a proactive strategy that takes account of the strength of on-the-job adaptation. Nevertheless, rational design cannot be dismissed entirely. It does produce a globally coherent rule set that can be degraded by local adaptations. In this paper I discuss an innovative knowledge map that can represent both local and global constraints, one that could facilitate rational design and also permit operational personnel to see the effects on global constraints of their local adaptations. A component of this system is an audit management process that that will maintain global coherency as it enhances the robustness of local procedures by feeding lessons learned in practice back into a global system update.

Safety Management: Local or Global?

…the problem of rules created by those who do not have to live the life

John Irving, discussing a dominant theme of his novel, The Cider House Rules

The drift of local practice away from demands of global constraints is a major threat to safety in today's complex, socio-technical systems. Current approaches to safety seek to eliminate this drift through use of tight control in the form of rules and procedures. However, adaptation of local practice is inevitable and represents an opportunity to enhance safety if fed back into system redesign in a manner that takes account of the global constraints.

Redesign through adaptation to local practice might be accomplished via an innovative and revolutionary knowledge map that depicts both global and local constraints in a way that can be assimilated by designers and by operational personnel. The structure of that knowledge map could be based on an Abstraction-Decomposition map, a form of representation that conforms to the deep structure of human problem solving. This knowledge map could potentially represent a paradigm shift in safety management, one that is proactive and also responsive to operational expertise.

The changing nature of socio-technical systems

Human collaborative systems are inevitably open to the generation of new properties and the issue I address in this paper is how to deal with that ***openness***. Current rule-based approaches to safety management generally fail to take account of two pervasive properties of complex socio-technical systems, firstly that the human participants are constantly changing the system, and secondly that this process of change, induced by operational experience, has enormous (generally untapped) potential for enhancing safety.

Most approaches to safety management attempt to lock the system down so that it does not generate new properties. This is done by the imposition of detailed rule sets. That strategy can work well in the case of orderly, non-critical systems (even if they are open) and it can appear to work for a considerable time even in open, safety-critical systems. However, open systems are infinitely generative. Thus, we cannot construct a rule set that will incorporate all possibilities. Worse, we may try to be comprehensive, but that attempt can produce such a large rule set that its very size confounds those who must work with it.

The openness of a complex socio-technical system is a source of latent pathogens (Reason, 1997) that can amplify the effects of seemingly normal events to the point that they reverberate through the system in ways never imagined by designers or operators. The fundamental assumption of the argument I present here is that we have neglected this openness and that we continue to pay a price for that neglect.

The potency of operational experience

Once deployed, rule sets become established as ***the formal way*** of doing things. There is generally no recursive mechanism to feed ***lessons learned in practice*** back into the redesign or retuning of the system. Procedures developed from a rational analysis of requirements rather than from within practice itself are often clumsy, fragile and

incomplete.

A contrast to rational analysis can be found within aviation where aircrews develop procedures as they work out how to accomplish specific tasks. Procedures developed in this manner constitute abbreviated descriptions of expert performances. They provide a detailed and well-crafted plan of action that is robust and efficient (Lintern & Naikar, 2001). Aviation has led the way in the development of robust procedures from distillation of actual practice. Nevertheless, local adaptation via procedures developed in practice is contrary to the philosophy of rational design and often generates informal mechanisms that directly oppose the expressed goals and values of safety management (McDonald, Corrigan & Ward, 2002).

Procedural Drift

Success in dealing with the issues of openness and the fragility of rational procedures will constitute a much-desired paradigm shift in safety management. Much as a martial arts expert uses the energy of an opponent to advantage, ***lessons learned in practice*** could be fed back into redesign of the system, thereby improving safety by enhancing robustness of procedures while, at the same time, accommodating to the openness of the system.

The Problem

The essential problem I confront here is that design of any new system is generally driven by rational considerations of designers who either are not practitioners or who are not currently involved in practice (Lintern, 1995). The rational system, once deployed, will be reshaped in practice by local pressures. In a distributed system, local practice will drift to become disconnected from global constraints. This is possibly the major threat to safety in today's complex socio-technical systems (Rasmussen, Pejtersen & Goodstein, 1994; Reason, 1997).

All past and contemporary approaches to safety seek to eliminate the drift generated by local pressures through use of tight control in the form of rules and procedures. The approach I offer here seeks to exploit that drift, to permit it to function as a local means of developing robust and efficient procedures, but to guide that drift by maintaining explicit connection to global constraints. Thus, the strengths of operational practice would be coordinated with the strengths of rational design to enhance system design, operational practice, and system redesign.

A Case Study

A stimulus for this approach, one that illustrates the need and the challenges, is an analysis by Snook (2000) of the destruction of two US Army Black Hawk helicopters over Northern Iraq by two USAF F-15s on 14th April 1994 during **Operation Provide Comfort**. All on board the Black Hawk helicopters, which included a number of UN peacekeepers, perished in this accident. The accident occurred despite AWACS coverage and despite a host of carefully designed systems that should have prevented it.

The F-15s involved in this accident were assigned the task of sanitizing the operational area, i.e. of ensuring there were no enemy aircraft and that it was safe for other allied flights. Although the F-15 flight was to be the first into the area that day, the two Black Hawks were already there. The F-15 pilots asked at three different times whether there had been adjustments to the Air Tasking Order (which did not identify the Black Hawk operation) and were advised there was not. One of these requests went to the AWACS team who knew of the Black Hawk operation. The AWACS team followed the engagement without raising the possibility that these two helicopters, read by the F-15 pilots as hostile, were in fact US aircraft. All this unfolded against a backdrop of no enemy incursions into this space in a considerable time.

Procedural Drift in Complex Systems

Analyses of this accident (Snook, 2000; Leveson, Allen & Storey, 2002) reveal the challenges facing the design and operation of complex, socio-technical systems. Although the original design of procedures (as embedded in the Operations Plans for Operation Provide Comfort) appeared to be sound, local pressures of operational practice induced a drift to locally efficient but globally inconsistent procedures. Snook (2000) argues that this process is inevitable and posits an engine that cycles through four states:

1. Planners assume a ***tightly coupled*** system in which interdependent processes affect each other directly and immediately. Given that assumption, planners over-design the system as a conservative approach to reducing the possibility of accidents from interactions of tightly coupled processes. Finally, planners assume that operational personnel will follow procedures as specified.
2. Operational personnel initially assume that all rules are justified and that failure to follow the rules will have severe consequences (beyond those of disciplinary action). However, the system is predominantly ***loosely coupled*** and the rules not well tuned to operational practice. Operational personnel come to believe that strict adherence to the rules is unnecessary and they subsequently implement local adaptations, which then become the locally accepted ways of doing things. Snook refers to this process as ***Practical Drift***. Following Johnston (2003), ***Procedural Drift*** is preferred in

this paper as a term better suited to aviation.

3. While the system is **predominantly** loosely coupled, it is not entirely so. Occasional circumstances bring normal processes into an unusual (but not extraordinary) confluence of tightly coupled systems. Because the global rationality of the **system-as-designed** has been degraded, the **local adaptations** permit the now tightly coupled processes to interact in unfortunate ways, often resulting in an incident or accident.
4. The management response to any ensuing accident is to re-establish global rationality by writing and then more strictly enforcing an enhanced rule set. This effort reestablishes global control but increases the force that generates Procedural Drift.

This engine might be seen as a behavioral pump with four cylinders (Figure 1) in which the motive force is drawn from the ecology of the system, where rational logic is overcome by what we might call an ***eco-logic***. From this perspective, Procedural Drift is pervasive in complex socio-technical systems that are predominantly loosely coupled. Rasmussen, et al. (1994) view this as an inevitable migration towards the boundaries of safe operation where serious consequences can result if occasional but normal circumstances bring processes into an unusual confluence of tightly coupled systems.

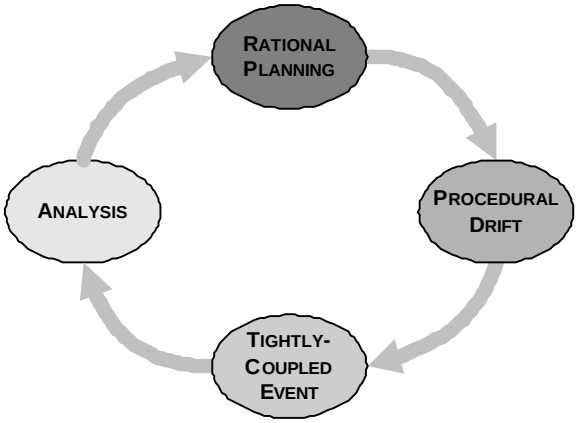

Figure 1. Local practice becomes disconnected from global constraints through a cycle of Rational Planning à Procedural Drift à Tightly Coupled Event à Analysis à Rational Planning à ……

Procedural Drift in Operation Provide Comfort

"Well Homer, you're the only one who's read those rules and so you're the only one feeling guilty."

Arthur Rose, Crew Boss (played by Delroy Lindo) in the film version of The Cider House Rules

Procedural Drift was widespread in Operation Provide Comfort, seemingly influencing local operations in all corners of the system. For example, Army and Air Force operations were poorly integrated and the natural consequence was a drift to localized operational procedures that were not mutually compatible. Army pilots were unaware of the correct procedure in relation to Identify Friend or Foe (IFF) and failed to fully understand the implications of the sanitizing role that the USAF had in ensuring the operational area was clear of enemy aircraft. The failure to coordinate IFF codes was identified as one of the many significant events in the destruction of the Black Hawks.

Systemic Issues

... operators would not always follow the written procedures … because the desired goal would not be achieved … (they were) criticized for "lack of procedural compliance". The operators decided …. they would do exactly what the procedure said … became stuck in an infinite loop … criticized… yet again …for 'malicious procedural compliance.'

Kim Vicente (1999), Cognitive Work Analysis, p xv

The processes used to develop procedures for Operation Provide Comfort are typical of design approaches to complex, large-scale socio-technical systems. Johnston (2003) describes a number of aviation contexts that further illustrate the pervasive problems:

- Systems are over-designed with an unnecessarily complex overlay of rules and procedures.
- The extensive documentation that publishes rules and procedures seems comprehensive but is not.
- The polite fiction is maintained that operational personnel are fully conversant with this documentation whereas casual analysis suggests that no one could possibly be fully conversant with such an extensive (and fluid) set of documentation.
- It is assumed that rational planning can produce robust and effective procedures. However, procedures developed by rational planning are often clumsy and fragile.
- Although it is assumed that complex socio-technical systems such as Operation Provide Comfort are static, many dynamic forces are at work to force continual change.
- The inevitability of local adaptation is not acknowledged and so there is no global oversight to ensure that local adaptations remain consistent with global constraints.
- Local adaptations emerge from lessons learned in practice, which is widely recognized as a powerful force for tuning effective behavior, but no mechanism is established for feeding the lessons of operational experience back into a global system update.

It is ironic that a design philosophy oriented towards ensuring safety produces so many system features that actually compromise safety.

Today's Typical Response

Safety management appears to be locked in a wrong-headed approach of retrospective analysis followed by development of more intricate control. The typical adjustment following an incident such as the destruction of the Black Hawk helicopters in Northern Iraq is to develop more rules to eliminate the possibility of a repeat incident of that type. This approach ignores Perrow (1984) who has provided evidence that larger, more complex rule sets can actually *increase* the risk of serious incident. Figure 1 supports Perrow's claims by depicting a process in which rational planning only feeds the motive force of procedural drift. However, even after his extensive and insightful analysis, Snook (2000) is at a loss about how to rectify the situation.

Scientific approaches to safety research emphasize this retroactive, control-based philosophy. Leveson, et al. (2002), who reviewed the loss of the Black Hawk helicopters over Northern Iraq, have developed a model in which accidents are viewed as resulting from a lack of constraints imposed on the system design and operations, and are attributed to failures at several levels of the socio-technical control structure. The stated goal of the Leveson et al modeling effort is to identify (retroactively) the factors that might be addressed to prevent future accidents.

One issue identified by Leveson, et al. (2002) was the failure of coordination among AWACS controllers, which they attributed at least partially to training deficiencies. However, appropriate training had been designed into the system (Snook, 2000). AWACS crews are deployed as teams from the US and, prior to deployment, are required to participate (as a team) in two simulator-training sessions. This crew participated in only one session, which three of the crew's senior members did not attend (because they *did not have to*). Training before deployment had become devalued because the generally benign state of AWACS operations had permitted a drift towards more relaxed procedures (cf Rasmussen, 1997).

The failure in this case was not one of inadequate control but of Procedural Drift. No system of constraints can prevent effects of the sort exemplified by the degradation of training for AWACS crews in Operation Provide Comfort. Retroactive, constraint based safety management has been with us for a long time. Further scientific development of this sort of strategy constitutes a misguided effort to do the wrong thing more effectively.

Safety Management: An Open Systems Approach

A proactive, open systems approach to safety management requires the representation of design rationale in an appropriately structured knowledge map. This knowledge map will be comprehensive and integrated. It will reveal global and local constraints and also the interplay between them. It will be used by designers to represent design rationale and by operational personnel for their comprehensive assimilation of system constraints. Because the map will reveal design rationale, operational personnel will be able to assess in real time the effects of any local adaptations they implement and the relationship of these local adaptations to the global constraints.

This map will also support redesign based on the explicit assumption that lessons learned within operational practice (procedural drift, local adaptation) could enhance rather than degrade safety. Operational personnel will be encouraged to note their local adaptations within the framework of the knowledge map and an audit team will have the assigned responsibility of identifying local adaptations and confirming that they remain consistent with global constraints (Figure 2).

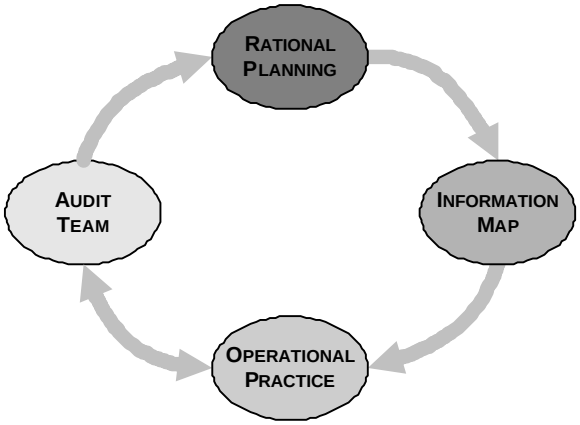

Figure 2. The proactive, open systems approach to safety management. Note the two-way process between operational practice and the audit team.

The Abstraction-Decomposition Map

The effectiveness of the proposed system will depend crucially on the way the representation of information is structured. The proposed structure will mirror an **Abstraction-Decomposition map**, which is a knowledge representation tool developed by Jens Rasmussen (Rasmussen, et al., 1994; Vicente, 1999). This map will be a unified system that incorporates information from all design documents.

The map in Figure 3 illustrates the essential elements of an Abstraction-Decomposition map:

- Different types of process and function terms at each of five levels of abstraction.
- Means-End relations between levels of abstraction (two-headed arrows between levels), i.e. the reason a function is provided propagates down from

intentional constraints while causes of functional changes propagate up from physical constraints.

- Decompositions shown by dashed, single-headed arrows within levels.

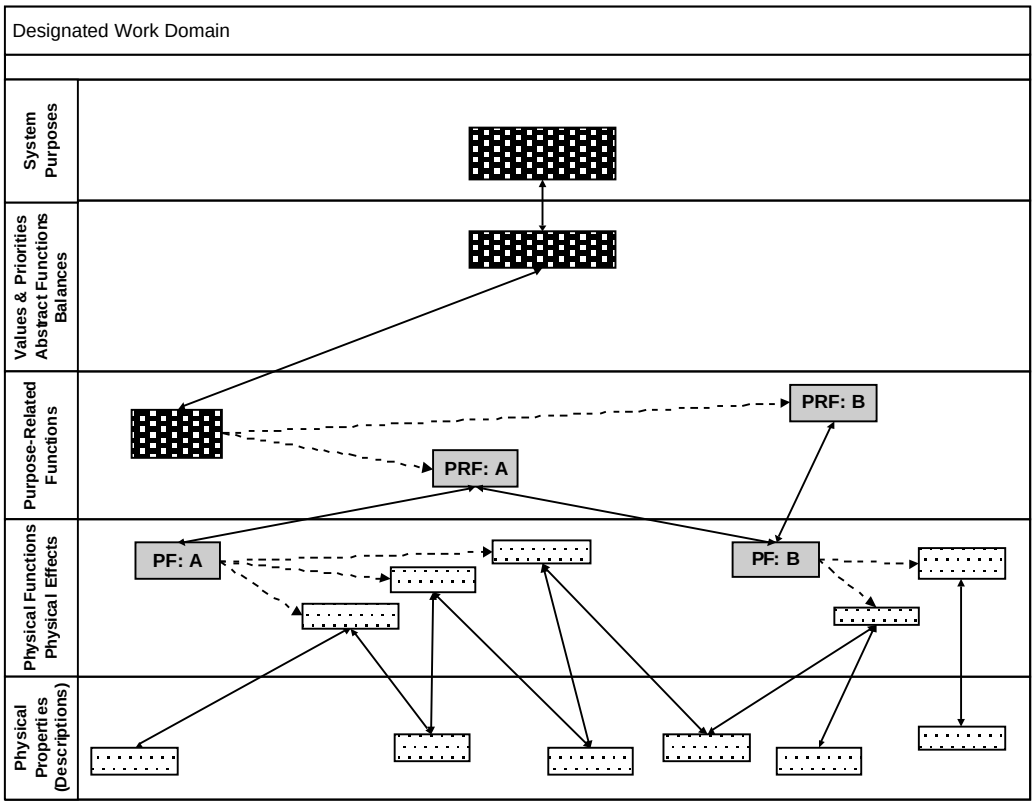

Figure 3, An Abstraction-Decomposition Map with functions at different levels of abstraction and decomposition

- Interdependencies shown by the crossings of the means-end relations (note the crossing means-end link from Physical Function B to Purpose-Related Function A).

Why Abstraction-Decomposition?

The Abstraction-Decomposition map is thought to be an important means of describing a complex, information-rich workspace because expert trouble-shooters and expert problem-solvers are known to navigate through an Abstraction-Decomposition space (as represented in figure 3) as they go about their work (Rasmussen, et al., 1994). Typically, they start with purposes or values at the system level and then work towards decompositions at physical object and physical function levels. Also, typically, the trajectory through the space is irregular and opportunistic. The general claim made by Rasmussen et al. (1994) and supported in our own work (Lintern, Miller & Baker, 2002) is that this pattern of cognitive activity is characteristic (at least implicitly) of natural reasoning behavior.

An Abstraction-Decomposition map reflects the manner in which people navigate through the deep structure of a problem (Rasmussen, 1995). Information should be presented in a way that maps into that deep structure. In the past, because of our reliance on typewriters and face-to face communication, permanence and locality were essential organizing principles and functionality was necessarily subordinate. Electronic documents and distributed communications now permit us to do things differently but we typically still adhere to those legacy principles through the constraints of our own legacy thinking (instructions and advisories for Operation Provide Comfort were distributed across at least 12 documents). However, we can and we should relinquish that legacy thinking and bring functional abstraction and functional granularity to the fore as the prime organizing principles in support of operational personnel as they explore the constraints of their work domain.

How this will work

The Abstraction-Decomposition map will support navigation through the deep structure of a complex socio-technical system because:

- Its functional clustering of information will help workers integrate the information they need for a problem in the way they need it.
- Representation at higher and lower levels of detail

is important. It will help operational personnel develop an overview of the system and then lead them to those parts relevant to their specific issues.

- Hidden interdependencies are the latent pathogens (Reason, 1997) of a system. Designers almost never make these evident. The abstraction-decomposition map represents them explicitly.
- The abstraction dimension of the abstraction-decomposition map shows causal relationships. Rules are developed for a reason and those reasons are generally discussed openly in design meetings. They are, however, not well documented and rarely find their way into operational manuals (cf Johnston, 2003). The abstraction dimension shows reasons through the means-end links. The reason for a function at one level is shown by its connection to one or more functions at the next highest level.
- Means of achieving requirements are revealed through links to functions at the next lowest level.

The Abstraction-Decomposition map establishes the basis for our information system but is not, in itself, open to adjustment on the basis of new properties. In the approach outlined here, it is opened to new properties via feedback from operational experience and via system audits.

Implementation

The Abstraction-Decomposition map might be implemented as an information table with an electronic surface on which it would be possible to develop computer representations of information structures. The information table would have a graphical interface that would rely heavily on iconic representation of critical properties and would use many of the standard tools of graphics programs (e.g. icon libraries, electronic pens, default shapes, connectors) and many of the standard means of computer interaction that permit intuitive and direct selection (touch activation, drag and drop, selection, pointing and linking).

Conclusion

The Abstraction-Decomposition map is a form of representation that explicitly and concurrently reveals global and local constraints. It can enhance safety by helping designers develop a comprehensive and internally consistent rule set, by helping operational personnel understand the relationship of their activities to global constraints and to local requirements of other interdependent functions, and by helping maintain consistency between global and local constraints as a system is redesigned in response to lessons learned in operations. Unlike any other approach to safety management, this approach not only honors the openness of a socio-technical system but also exploits that openness to further enhance safety.

Acknowledgement

From one perspective, Jens Rasmussen had nothing to do with this paper. From another perspective, he had everything to do with it. Thank you Jens, for your inspiration.

References

Johnston, N. (2003). The Paradox of Rules: Procedural Drift in Commercial Aviation. In R. Jensen, (Ed), *Proceedings of the Twelfth International Symposium on Aviation Psychology*, April 14-17, 2003, *Dayton, Ohio* [CD-ROM].

Leveson, N.G.; Allen, P. & Storey, M-A (2002). The Analysis of a Friendly Fire Accident using a Systems Model of Accidents. *International Conference of the System Safety Society*, Denver.

Lintern, G. (1995). Flight instruction: The challenge from situated cognition. *The International Journal of Aviation Psychology, 5*, 327-350.

Lintern, G., Miller, D. & Baker K. (2002). *Proceedings of the 46th Human Factors and Ergonomics Society Annual Meeting*. (p. 531-535). Santa Monica, CA: Human Factors and Ergonomics Society.

Lintern, G. & Naikar, N. (2001). *Analysis of Crew Coordination in the F 111 Mission* (DSTO-CR-0184). Melbourne, Victoria, Australia: Aeronautical & Maritime Research Laboratories, Defence Science & Technology Organisation.

McDonald, N., Corrigan, S. & Ward, M. (2002). Cultural and Organizational factors in system safety: Good people in bad systems. *Proceedings of the 2002 International Conference on Human-Computer Interaction in Aeronautics (HCI-Aero 2002)* p 205-209. Menlo Park, CA: American Association for Artificial Intelligence Press.

Perrow, C. (1984). *Normal accidents: Living with high-risk technologies*. New York: Basic Books.

Rasmussen, J. (1995). Invited Address, 1st Berliner Workshop on Man-Machine Systems; Technical University Berlin, 11-13 October '95

Rasmussen, J. (1997). Risk management in a dynamic society: a modelling problem. *Safety Science, 27*, 183 - 213.

Rasmussen, J., Petjersen, A. M., & Goodstein, L. P. (1994). *Cognitive systems engineering*. New York: John Wiley.

Reason, J. (1997). *Managing the Risks of Organisational Accidents*. Aldershot, UK: Ashgate Aviation.

Snook, S.A. (2000). *Friendly Fire*. Princeton University Press.

Vicente, K. J. (1999). *Cognitive Work Analysis: Towards safe, productive, and healthy computer-based work*. Mahwah, NJ: Lawrence Erlbaum & Associates.