The Process of Tailoring Models for *a priori* Safety and Risk Management for use within Industry

Steele, K.[1] and Pariès, J.[2]

[1] Foundation for an Industrial Safety Culture, 5 rue Paulin Talabot, 31106 Toulouse, France
Dédale, 15 place de la Nation, 75011 Paris, France
Conservatoire National des Arts et Métiers, 41 rue Gay Lussac, 75005 Paris, France
ksteele@dedale.net
[2] Dédale, 15 place de la Nation, 75011 Paris, France
jparies@dedale.net

Abstract. Models of risk or safety as they are published in a general form are used by industry as a framework or template for building a specific, relevant version of the model to fit a particular organisation or activity. Thus choosing a model is only the first step, somehow the model must be made useful and usable for the application, customised for the specific situation. The ultimate usefulness of a model therefore depends not only on its conceptual robustness and appropriate fit for the application but also on how well it is adapted for use. The process of tailoring or adapting models *a priori* for organisational use therefore presents a significant challenge for safety researchers endeavouring to have an impact on industrial practice. The pragmatic issues of applying traditional risk and accident models are not always thoroughly addressed and, in some cases, are ignored completely in the literature. This problem will be even more difficult, and more important to overcome, for the implementation of complex, systems models which capture safety and successful performance as well as (or instead of) failures and unwanted events.

1 ADAPTING THE TOOL TO FIT THE JOB

The explicit use of models is an integral part of accident investigation and safety management work. As simplified representations of the phenomenon of focus, models are a more manageable means of visualizing and understanding the world. They are an indispensable tool for dealing with complex, large-scale, safety-critical systems like those of interest to Resilience Engineering research.

1.1 Different Models

Accident Models. Perhaps the most commonly recognised use of models is in accident investigation. Reason's Swiss cheese model, for example, provides a metaphor for the progression of an accident scenario, causes, contributing factors, as well as safety barriers. Such accident models can, and do, provide the underlying structure for investigative activities and recommendations.

Inherent in accident models are assumptions about the boundaries of the system, the relationships between system elements, the typology of possible causes or contributing factors, and even the underlying notion of the nature of accidents (i.e. in the case of the Swiss cheese model: accidents result from linear propagations of failures).

A generic accident model may be redrawn and re-labelled to illustrate the features of the particular accident and the operational environment as deemed appropriate for illustration or explanation after the events. Thus the model is tailored or customised in order to be locally relevant and useful.

The Human Factors Analysis and Classification System (HFACS) is an example of the generic model (once again using Swiss cheese as the example) being customised for use in industry, however this classification scheme is specified to the level of detail appropriate for airline flight operations, and not beyond to the more specific company level [Shappell & Weigmann, 2000].

While this customisation process is by no means straightforward, (there are many subjective judgments required) when the unfortunate time comes to apply the model the problem is, at least, bounded to the specific accident scenario.

Risk Models. This is not the case when customising models for risk analysis, since this tailoring process is carried out *a priori*, that is to say, based on an analysis of the system in the absence of any specific event or accident. It is thus an unbounded problem: the limits to the system which can be modeled are arbitrary. While it may not be as evident and the term “model” is not used as frequently, risk matrices and fault trees, etc. are also models, in many senses. They are simplified representations of the system which specify the relationships between the components, and reveal underlying assumptions about the notion of safety.

Safety Models. One of the tenets of the Resilience Engineering movement is that to make progress on safety there is a need for models of performance and system behaviour that goes beyond accidents, risk, faults, failures, errors, etc. to capture the positive or successful aspects of performance and the operational controls and constraints. According to Diane Vaughan [1999], it is the same adaptive human behaviour responsible for both success and failure, thus it is not possible to separate the concepts. For the purposes of this discussion, the term “safety model” will be used to describe models of the system which include normal operations and successful performance, rather than focussing exclusively on risks, hazards, failures, or accidents.

1.2 The Context of this Paper

The focus of the authors’ research is the methodology for customising SaMBA (Safety Model Based Analysis) developed by Dédale [Bieder & Pariès, 2003; Pariès & Bieder, 2003] and used in the aviation industry (previously referred to as AMSMA and ERASM). SaMBA is a safety management tool which must be adapted, *a priori*, to support risk and safety management as well as to provide a framework for understanding and leveraging event reports. The model itself constitutes a comprehensive safety architecture of the target system or activity so the construction process is necessarily resource intensive, however, once created, it will be continually used and updated through the incident reporting process.

One significant difference between SaMBA and risk models that simplifies the tailoring process is that SaMBA does not focus on creating an inventory of *all* potential hazards, but rather on the (finite) means by which the organisation keeps control of the process.

It was in the context of this research, during a search for other customisation methods for applying models in practice, that the underemphasis of this important tailoring process was recognised as a topic relevant to discussions of resilience.

2 UNDERSTANDING 'HOW' AS WELL AS 'WHAT'

There is some middle ground between methods or guidelines which are overly prescriptive and those which are underspecified. This paper is not advocating proceduralisation at the expense of all else or trying to stifle creative local solutions. It is clear, however, that busy practitioners do not have time to waste listening to grand academic theories if they are completely impractical and lack consideration of commercial constraints. To use an example from the aviation industry, airline safety managers usually have only one question after hearing academics give presentations: "How do we accomplish this?" There is a demand for a bridge between academia and industry, and that means putting thought into how theory can be applied in practice. Clear guidance on using models is needed to avoid confusion and ensure the users reap the potential benefits.

It is also important for there to be some minimum standard and transparency associated with the application of a model. A recognised standard, such as compliance with certain regulation or use of a standard tool creates shared understanding and expectations.

2.1 Model Customisation

Risk Models. Unlike accident models, risk models are intended to be used proactively (at least as far as possible given that the frequency data may be based on past incidents or accidents) and thus need to be adapted *a priori*. This process of modeling the system 'from scratch' seemed particularly relevant to the researchers from the outset. However, the results of an initial literature search did not yield much of value for the purposes.

While there is a large corpus of work on the topic of risk management and analysis, and in particular on how to analyse the data in the models, there is relatively little practical guidance on how to identify the hazards initially to build the risk inventory. More concerning is that amongst the rare documents reviewed which did describe a method, it was not uncommon to find a process beginning with a statement like: "once the hazards have been identified…" and then moving on to explain the subsequent steps in more detail.

This is worrying because of the potentially limitless number of hazards which could be identified and included in the analysis, the unbounded nature of the problem, as mentioned above. No matter how thorough the search, a specific hazard inventory or fault tree will always be incomplete. ‘Guidance’ such as this mentioned above does not encourage a thorough search. This challenge is further compounded by the dynamic nature of the industry since the technology and operating context are continually changing. There will always be surprises.

Some catastrophic events in aviation which are now recognised hazards but were surprises at the time are the United accident at Sioux City, Iowa in 1989, the 9/11 terrorist attacks, and the Alaska Airlines crash in 2000. It wasn’t considered possible to have a total hydraulics failure before the Sioux City accident, thus no training for this emergency procedure, nor were there any existing regulations for the spatial separation of the hydraulic lines. The terrorist attacks in September 2001 certainly came as a surprise to most of the world (regardless of whether the United States government knew of the possibility) and there were no structural or procedural provisions in place for such a scenario. Even following the Alaska Airlines crash, it is unlikely that many airlines have included in their risk models the possibility that the inspection interval approved by the Federal Aviation Authority and the engineers at the equipment manufacturer may be insufficient [Dekker, 2005].

There are other steps in the risk assessment procedure which are known weaknesses and which are difficult to specify, namely the subjective process of judging severity and frequency. These, in combination with this problem of completeness of the model certainly weaken the sense of security associated with the ‘objective’ numbers which probabilistic risk assessment (PRA) produces.

It may be that the few thoroughly-documented, well-established methods for hazard identification which do exist (e.g. FMEA, HAZOP) are robust and versatile enough for the needs of most risk analysis methods, and it may be a wise strategy to have common approaches rather than a different solution for each company.

It may also be that the problem of identifying risks using this kind of model is actually not that complicated, since the underlying safety paradigm is more mechanistic than ecological, based on the normative concepts of error, deviations, and failure. If the possibilities are limited to linear combinations of ‘human errors’ and mechanical failures, the inventory process may not require overly detailed guidelines to be acceptably thorough. If this is the case, then the process of customising risk models based on this worldview is too different from the problem of implementing SaMBA to be transferable for the purposes of this research.

2.2 Challenges for Customising Resilience Models

The Problem of Knowing What to Look for. The introduction of the concepts ‘latent failure’ and ‘system accident’ in recent decades made hazard identification and risk analysis more challenging. Although the concept of latent failures is not difficult to understand in theory, identifying such failures in practice before an accident is another matter. This is due in part to the subjective nature of what is ‘safe’, since it is highly contextual and the only conclusive proof of the ‘absence’ of safety is hindsight in the wake of an accident or near-miss.

The challenge then becomes how to take stock of all company policies, activities, structure, etc. and determine with precision where the dangers are lurking. Which of the normal practices in an inherently risky business [Dekker, 2005] will be an ingredient in the next accident? This “weak signal” problem is plaguing industry and researchers have not yet provided sufficient answers [Axelsson, 2006]. Thus the challenge of *how* to actually identify hazards in practice is only becoming more difficult, as the notion of what can be considered a hazard is expanding to include ‘normal’ practice.

Complexity and Clashing Paradigms. As the nature of accidents has evolved and become more complex [Amalberti, 2001] the models have in turn become more complex and detailed [Steele, in press]. As mentioned, the emphasis is necessarily shifting from modeling accidents, risk, failure, and errors to modeling performance, success, and normal operations proactively [Svedung & Rasmussen, 2002; Hollnagel, 2004; Dekker, 2005; Pariès, 1999, 2006]. The models emphasise relationships such as interactions, constraints, or resonance between system elements [Hollnagel, 2004; Leveson, Daouk, Dulac, & Marais 2003; Leveson et al. 2006].

Although AcciMap [Svedung & Rasmussen, 2002], STAMP: Systems-Theoretic Accident Model and Processes [Leveson, et al. 2003], and FRAM: Functional Resonance Accident Model [Hollnagel; 2004] all use the term accident in their names, these models capture normal operations and may be useful for understanding success as well.

Such models are less intuitive to understand (than traditional PRA for example) and more complicated and involved to tailor for industry’s use. Although many in the industry recognise the limitations of current approaches and are eager for something new, the difference in fundamental philosophies makes the new models incompatible with the past worldview, further increasing the challenge. The onus is therefore on the developers of the models or other researchers to provide practitioners with reasonable, realisable methods for understanding and using new tools. Resilience Engineering must strive to develop well thought-out processes for applying new theories and models in practice.

2.3 Examples of Success

Sidney Dekker’s Field Guides [2001, 2006] are an example of a past success story. These books target the layperson, convincingly summarising some of the less traditional perspectives on error and accident causality making them seem like common sense. The Cognitive Reliability and Error Analysis Method (CREAM) by Erik Hollnagel [1998] is another example of an innovative method which has been designed to be ‘usable’. Both Dekker and Hollnagel provide clear guidance on how to put these ideas into practice. They are not the only researchers doing this of course, it is just not always an easy thing to find the balance between overly rigid guidelines and underspecification.

The written format is not the only option, training courses and safety consulting are popular modes for the transfer of such skills and knowledge, but the information must originate from somewhere and have solid theoretical support. Unity, rather than fragmentation, within the Resilience community would give future endeavours better momentum.

3 CONCLUSIONS

Choosing a model to use for safety management activities is only the first step, somehow the model must be made useful and usable for the application. The ultimate usefulness of a model therefore depends not only on its conceptual robustness and appropriate fit for the application but also on how well it is adapted for use. The process of tailoring or adapting models *a priori* for organisational use, therefore, presents a significant barrier for safety researchers to influence industrial practice. The pragmatic challenges of applying traditional risk and accident models are not always thoroughly addressed and, in some cases, completely untreated within the literature. This problem will be even more difficult, and more important to overcome, for the implementation of new complex, systemic models which focus on safety, successful performance, interactions, and constraints rather than just failures and accidents.

A Means or an End? Although it is not central to this discussion, the question arises when considering the act of adapting models of whether this collaborative activity (normally involving senior personnel) could have significant benefits for the organisation beyond the resulting model itself, as mentioned by Steve Epstein elsewhere in this publication and Kanki, Marx, & Hale [2004,1275]. The tailoring process up to this point has been viewed as a means to an end rather than something of value in itself, but during the investigation of the method it is relevant to consider the spin-off benefits. This process may comprise an integral safety-management exercise, proving beneficial for creating consensus and aligning visions of the safety model. Regularly repeating some parts of the process to keep the model up-to-date is required to some extent for the maintenance of SaMBA. This could be a valuable "broadening check" [Woods, 2003, personal communication, 2004] both at the individual and organisation levels, facilitating sensitivity to organisational drift.

Lessons beyond Modeling. Parallels can be drawn here with the challenges of putting theory into practice in the case of regulated aviation safety interventions such as operational incident reporting systems or safety management systems. The reasons for the failure of many incident reporting systems to accomplish their ambitious goals are manifold (although it is not a simple matter to objectively judge whether such interventions do actually 'improve safety'). One key reason for this failure is the lack of helpful guidance on how to effectively transform a large quantity of event data into organisational change which can be expected to prevent accidents or 'improve safety'. Proof-of-concept and sound guidance for implementation is imperative for such interventions; otherwise wasting resources on ineffective activities may actually have a detrimental effects safety.

ACKNOWLEDGEMENTS

This work was supported by the Foundation for an Industrial Safety Culture (FonCSI).

REFERENCES

Amalberti, R. (2001). The paradoxes of almost totally safe transportation systems. *Safety Science*, *37*, 109-126.

Axelsson, L. (2006). Structure for management of weak and diffuse signals. In Hollnagel, E., Woods, D., & Leveson, N. (eds.). *Resilience Engineering: Concepts and Precepts*. Ashgate Publishing Ltd: Aldershot, UK.

Bieder, C. & Pariès, J. (2003). AMSMA – Aviation maintenance safety management assistant: a top-down approach to derive general safety lessons from reported events. *Proceedings of the JRC/ESReDA Seminar on Safety Investigation of Accidents.* Petten, The Netherlands, 12-13 May.

Dekker, S.W.A. (2006). *The Field Guide to Understanding Human Error*. Ashgate Publishing Ltd: Aldershot, UK.

Dekker, S.W.A. (2005). *Ten questions about human error: A new view of human factors and system safety*. Lawrence Erlbaum Publishers: New Jersey.

Dekker, S.W.A. (2002). *The Field Guide to Human Error Investigations*. Ashgate Publishing Ltd: Aldershot, UK.

Hollnagel, E. (2004). *Barriers and Accident Prevention*. Ashgate Publishing Ltd: Aldershot, UK.

Hollnagel, E. (1998). *Cognitive Reliability and Error Analysis Method (CREAM)*. Elsevier Science Ltd: Oxford, UK.

Kanki, B., Marx, D., & Hale, M. (2004). Socio-Technical probabilistic risk analysis: Its capabilities and limitations. In Sptizer, C., Schmocker, U., & Dang, V. (eds.) *Probabilistic Safety Assessment and Management*, 2004(3), 1271-1275. Springer.

Leveson, N., Daouk, M., Dulac, N., & Marais, K. (2003). *A Systems-Theoretic Approach to Safety Engineering.* Retrieved 7 November, 2003 from the internet: http://sunnyday.mit.edu/accidents/external2.pdf

Leveson, N., Dulac, N., Zipkin, D., Cutcher-Gershenfeld, J., Carroll, J., & Barrett, B. (2006). Engineering resilience into safety critical systems. In Hollnagel, E., Woods, D., & Leveson, N. (eds.). *Resilience Engineering: Concepts and Precepts*. Ashgate Publishing Ltd: Aldershot, UK.

Pariès, J. (1999). Shift in aviation safety paradigm is key to future success in reducing air accidents. *Presentation to the ICAO Regional Symposium Régional on Human Factors and Aviation Safety,* Santiago du Chili. ICAO Journal vol 54, N°5. Montréal.

Pariès, J. (2006). Complexity, emergence, resilience. In Hollnagel, E., Woods, D., & Leveson, N. (eds.). *Resilience Engineering: Concepts and Precepts*. Ashgate Publishing Ltd: Aldershot, UK.

Pariès, J., & Bieder, C. (2003). A Complementary Approach to Support Risk Management and Decision-Making. *Communication to the 6th Australian Aviation Psychology Symposium.* Sydney, Australia.

Shappell, S., & Weigmann, D. (2000). The Human Factors Analysis and Classification System – HFACS. DOT/FAA/AM-00/7. *www.nifc.gov/safety_study/accident_invest/humanfactors_class&anly.pdf*, accessed on 02-09-06.

Steele, K. (in press). *Past progress and future challenges for modeling safety and accidents.* Proceedings of the 50th conference of the European Association for Aviation Psychology. Potsdam, Germany.

Svedung, I., & Rasmussen, J. (2002). Graphic representation of accident scenarios: Mapping system structure and the causation of accidents. *Safety Science,* 40, 397-417.

Vaughan, D. (1999). The dark side of organizations: Mistake, misconduct, and disaster. *Annual Review Sociology* 25, 271-305.

Woods, D. (2003). *Creating foresight: How resilience engineering can transform NASA's approach to risky decision making.* Testimony on The Future of NASA for Committee on Commerce, Science and Transportation. Retrieved 7 November, 2003 from the World Wide Web: http://csel.eng.ohio-state.edu/woods.