A White Paper on Resilience Engineering for ATM

Cooperative Network Design

EUROCONTROL

Tempora mutantur, nos et mutamur in illis
("Times change, and we change with them")

INTRODUCTION

In January 2007, a project was launched by EUROCONTROL with the aim of understanding the new area of Resilience Engineering and its relevance to Air Traffic Management (ATM). Resilience Engineering is developing important tools and methods for both system developers and people responsible for the maintenance and management of system safety, in a number of industries.

This White Paper provides an overview of the work carried out and seeks to answer the following questions:

- What is Resilience Engineering?
- Why do we need it in ATM?
- What is the usefulness of performance variability?
- What does Resilience Engineering look like in practice?
- How does Resilience Engineering fit with other safety methods?
- How mature is Resilience Engineering and what is the added value for ATM?

The added value of an Resilience Engineering approach is that it provides a way to address the issues of emergent accidents and the often disproportionate consequences that are an outcome of ever more complex technologies and ever more integrated organisations. Resilience methods are likely to be important tools in the management and assurance of ATM safety in the future.

WHAT IS RESILIENCE ENGINEERING?

Resilience is the intrinsic ability of a system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both expected and unexpected conditions.

Performance variability:
The ways in which individual and collective performances are adjusted to match current demands and resources, in order to ensure that things go right.

Since humans are indispensable in all situations involving change, Resilience Engineering naturally has strong links with Human Factors and Safety Management. It is based on the following premises:

1. Performance conditions are always underspecified. Individuals and organisations must therefore adjust what they do to match current demands and resources. Because resources and time are finite, such adjustments will inevitably be approximate.

2. Some adverse events can be attributed to a breakdown or malfunctioning of components and normal system functions, but others cannot. The latter can best be understood as the result of unexpected combinations of performance variability.

3. Safety management cannot be based exclusively on hindsight, nor rely on error tabulation and the calculation of failure probabilities. Safety management must be proactive as well as reactive.

4. Safety cannot be isolated from the core (business) process, nor vice versa. Safety is the prerequisite for productivity, and productivity is the prerequisite for safety. Safety must therefore be achieved by improvements rather than by constraints.

Adopting this view creates a need for an approach that can represent the variability of normal system performance, and for methods that can use this to provide more comprehensive explanations of accidents as well as identify potential risks.

FROM SAFETY MANAGEMENT TO RESILIENCE ENGINEERING

Risk governance and safety management have traditionally, and with good reason, been concerned with what can go wrong and therefore lead to unwanted outcomes. The set of possible outcomes can schematically be shown as in Figure 1, where the x-axis describes predictability, ranging from very low to very high, and the y-axis describes the value of the outcome, ranging from negative to positive.

The established approaches to risk and safety mainly focus on the things that go wrong, more specifically those areas in figure 1 named disasters, accidents, and incidents – with occasional forays into near misses. The mishaps region describes unwanted outcomes that in practice have been eliminated. But there has traditionally been little or no focus on things that go right, despite the fact that these happen far more often than things that go wrong. If, for instance, the probability of failure is 10^{-4}, then there will be 9,999 normal outcomes for every failure!.

The focus of Resilience Engineering is on the whole set of outcomes, i.e., things that go right as well as things that go wrong – with the possible exceptions of the areas of serendipity and good luck, where we are mostly in the hands of fate.

The aim of Resilience Engineering is not only to prevent things from going wrong, but also to ensure that things go right, i.e., to facilitate normal outcomes. Simply put, the more likely it is that something goes right, the less likely it is that it goes wrong. And there is clearly an added value in trying to facilitate normal outcomes and in discarding the traditional separation between safety and productivity.

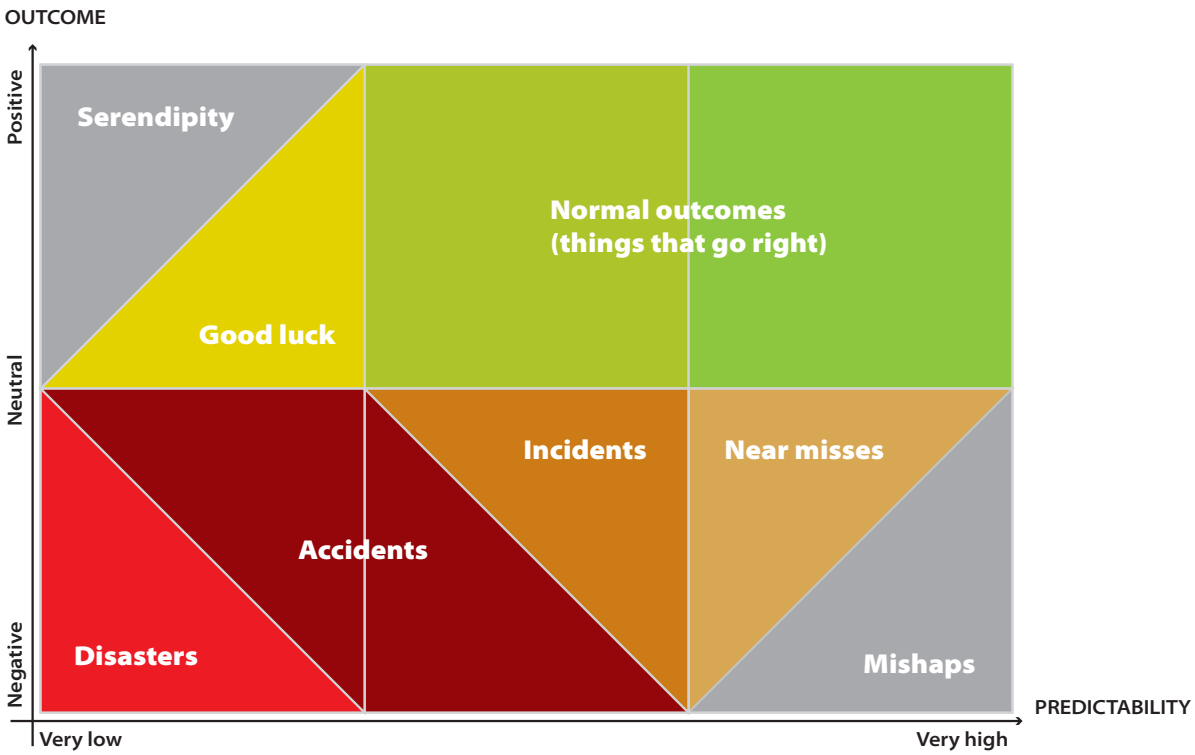

Figure 1: The set of possible outcomes

WHY DO WE NEED RESILIENCE ENGINEERING?

A simple answer to this question is given by the types of accidents that can occur in complex yet 'well-defended' systems, of which Air Traffic Management (ATM) is a good example. While accidents such as Überlingen (2002) with hindsight can be explained by failure modes emerging from the weak interaction of multiple factors, few would have considered this situation credible ahead of time. Traditional thinking makes it difficult to describe and understand how multiple factors can come together at a single point of time and in a position of airspace to produce something as disastrous as a mid-air collision.

Today, most ANSPs have safety nets, Safety Management Systems (SMS), safety assessment and assurance processes, and many are also improving their safety culture. These efforts add layers of safety to already safe systems. While this will reduce the likelihood of accidents even further, it also means that those that do slip through these 'nets' will be complex and multi-faceted. Such accidents will be due more to coincidences among the variability of functions and human performance in different parts of the system, than to manifest failures and incorrect human actions.

Accident analysis and risk assessment methods have usually been developed in response to problems following major technological developments or to cope with 'new' types of accidents. Figure 2 shows the distribution of some well-known methods used to address technical, human factors, and organisational issues, respectively. It is noteworthy that human factors methods came onto the scene after the accident at Three Miles Island in 1979, and that organisational methods were developed following the Chernobyl and Challenger accidents in 1986.

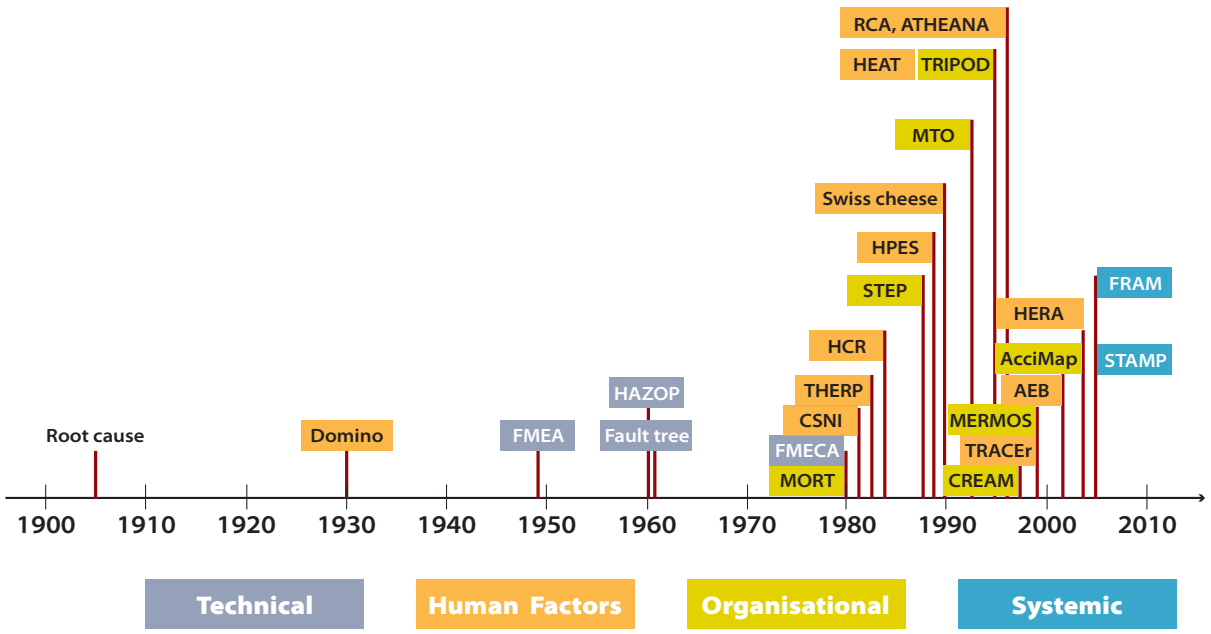

Figure 2: Accident Analysis and Risk Assessment Methods

Methods must clearly be powerful enough to address the problems found in real-life applications. Established ways of thinking about accidents, such as the Swiss Cheese analogy of holes in safety barriers that 'line up', will therefore at some time be unable to prevent, predict, and explain new types of accidents. The variability of normal performance may for instance combine to produce effects that go beyond what the Swiss Cheese analogy was intended to describe, even if the holes are allowed to grow or shrink dynamically and the blocks of cheese allowed to move around. In order to address these more complex phenomena, Resilience Engineering uses the principle of **resonance** to represent how the variability of normal performance can combine dynamically in ways that may lead to disproportionate (non-linear) effects. Safety assessment methods have historically developed from technical methods, via human factors methods, to organisational methods. While many current methods combine the technical, human, or organisational levels, resonance must be treated at the system level.

Resonance:
A principle that explains how disproportionate large consequences can arise from seemingly small variations in performance and conditions.

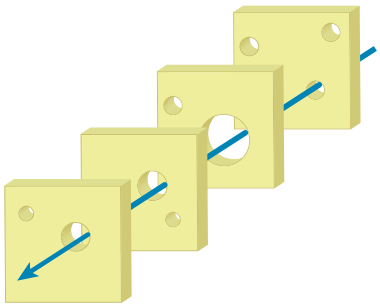

Figure 3: 'Swiss Cheese' Model

THE USEFULNESS OF PERFORMANCE VARIABILITY

The first premise of Resilience Engineering makes clear that performance variability is inevitable but also useful. Procedures and instructions are always incomplete, except for extremely simple situations. Following procedures and instructions to the letter will therefore both be inefficient and unsafe. To compensate for this incompleteness, individuals and organisations habitually adjust their performance to match the current demands, resources, and constraints. The ability to do so is at the heart of successful performance. But since information, resources, and time are finite, the adjustments will inevitably be approximate. Performance variability is thus unavoidable, but should be recognised as a source of success as well as of failure.

This has, however, not always been so. The concern for how human factors could affect safety began with the accident at the Three Mile Island nuclear power plant in 1979. At that time, the focus was naturally on control room operations (the 'sharp end'), and methods were therefore developed to support that. The situation in the 1980s with regard to the focus of safety is depicted in Figure 4.

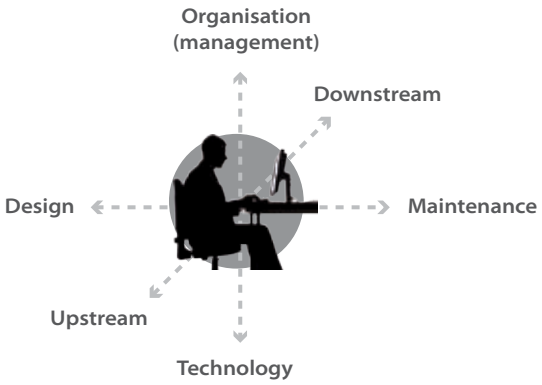

Figure 4: Safety Focus anno 1984

Socio-technical systems have since the 1980s become steadily more complex due to rampant technological and societal developments. The idea of a socio-technical system is that the conditions for successful organisational performance – and conversely also for unsuccessful performance – depends on the interaction between social and technical factors. The scope of safety assessment must therefore be extended in several directions. A 'vertical' extension is needed to cover the entire system, from technology to organisation. One 'horizontal' extension is needed to increase the scope to include both design and maintenance. A second 'horizontal' extension is needed to include both upstream and downstream processes. Altogether, the situation today therefore looks more as shown in Figure 5.

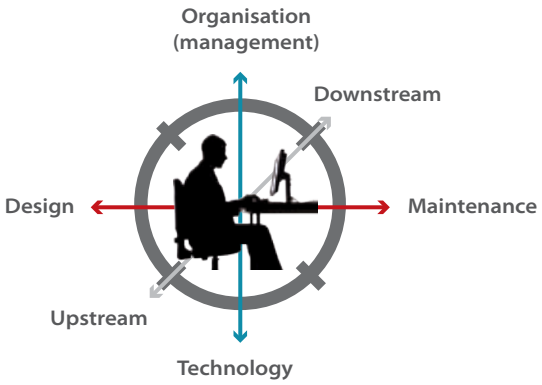

Figure 5: Safety Focus anno 2009

Upsteam-downstream means integration into the company's business. In aviation, gate-to-gate would be an example, i.e. you can no longer consider previously separate functions as separate. There are dependencies to what went before (upstream) and what comes after (downstream). In general production it is 'just in time' (JIT) in order to remove the need for inventories (raw materials, spare parts etc).

As a result of these developments, safety methods must today address systems that are larger and more complex than the systems of yesteryear. Because there are many more details to consider; because some modes of operation may be incompletely known; because of tight couplings among functions; and because systems may change faster than they can be described, the net result is that many systems, ATM included, are underspecified or intractable. For these systems it is clearly not possible to prescribe tasks and actions in every detail. This means that performance must be variable or flexible rather than rigid. In fact, the less completely the system is described, the more performance variability is needed.

	Tractable system	Intractable system
Number of details	Description are simple with few details	Description are elaborate with many details
Comprehensibility	Principles of functioning are known	Principles of functioning are partly unknown
Stability	System does not change while being described	System changes before description is completed
Relation to other systems	Independence	Interdependence
Metaphor	Clockwork	Teamwork

Table 1: Differences between tractable and intractable systems

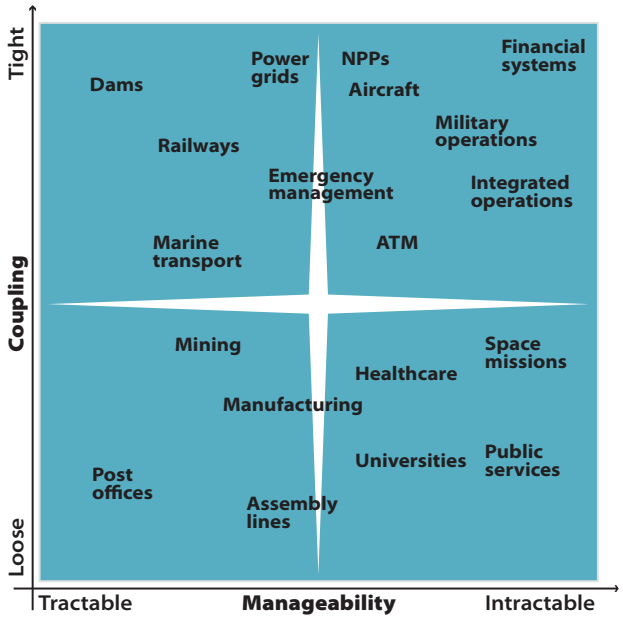

Figure 6: System Characteristics

It is useful to make a distinction between tractable and intractable systems. Tractable systems can be completely described or specified, while intractable systems cannot. The differences between the two types of systems are summarised in Table 1.

Most established safety methods have been developed on the assumption that systems are tractable. As this assumption is no longer universally valid, there is a need to develop methods to deal with intractable systems. Resilience Engineering is one answer to this problem.

Figure 6 shows the result of characterising different socio-technical systems on the dimensions of 'Coupling' and 'Manageability'. Existing methods are not well suited for systems in the upper right quadrant (intractable and tightly coupled), which makes this a primary focus for Resilience Engineering.

THE REASONS FOR PERFORMANCE VARIABILITY

To predict how resonance may lead to accidents, we must be able to describe and model the characteristic variability of the system. A main source of performance variability is the underspecification of work, as described in the previous section. In addition to the underspecification, human performance can vary for several other reasons:

Physiological and/or fundamental psychological factors (e.g., affecting perception and vigilance).

Higher level psychological factors such as ingenuity, creativity, and adaptability.

Organisational factors, as in meeting performance demands, stretching resources, substituting goals, etc.

Social factors, as in meeting expectations of oneself or of colleagues, complying with informal work standards, etc.

Contextual factors, for instance if the workplace is too hot, too noisy, too humid, etc.

Other factors such as the unpredictability of the domain, e.g., weather conditions, number of flights, pilot variability, technical problems, etc.

The challenge for resilience engineering is to represent the variability of a system, such as ATM, in a way that makes it possible to identify what may affect performance either adversely or positively. This must overcome two obstacles:

1. Because ATM is a **complex** and intractable system, its functions, their interactions, and potential variances, can only be specified approximately.

2. In a time of traffic growth, planned changes to ATM's fundamental infrastructure and technical systems must be reconciled with strong but varying financial pressures and demands.

   The ATM system environment feels the effect of the resulting **instability** and increasing variability as ANSPs try to ride out the changes whilst remaining safe and profitable. This requires them to be flexible, to rely on human ingenuity and skill, and to make use of performance variability rather than to constrain it.

HOW IS IT DONE?

Resilience Engineering provides the conceptual basis for a new perspective on safety. The leading prototype currently is the Functional Resonance Assessment Method (FRAM). This method is based on four principles:

- **The equivalence of success and failures.** Failures do not stand for a breakdown or malfunctioning of normal system functions, but rather represent the adaptations necessary to cope with the underspecification found in complex real-world systems.

- **The principle of approximate adjustments.** To get anything done people must adjust their performance to the current conditions. Because resources and time are finite, such adjustments will inevitably be approximate.

- **The principle of emergence.** Both failures and normal performance are emergent phenomena: neither can be attributed to or explained simply by referring to the (mal)functions of specific components or parts.

- **The principle of functional resonance.** FRAM replaces the traditional cause-effect relation with resonance. This explains how the variability of a number of functions every now and then may resonate, i.e., reinforce each other, leading to excessive variability in one or more downstream functions. The consequences may spread through the system by means of tight couplings rather than easily identifiable cause-effect links.

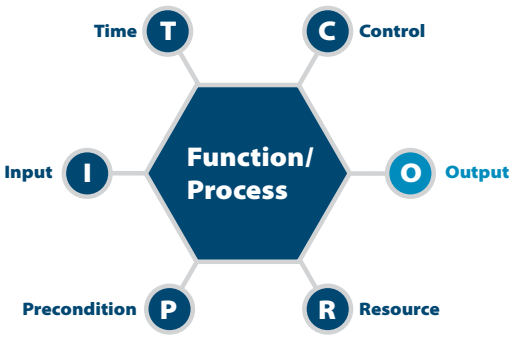

Figure 7: The six aspects of a FRAM function

FRAM is a developing method. The following is an illustration of what a FRAM analysis might look like for an overflight scenario:

A FRAM analysis consists of five steps:

1. **Define the purpose of the analysis.** FRAM can be used for safety assessment (looking at future events) as well as for accident investigation (looking at past events).

2. **Identify and describe the relevant system functions.** A function, in FRAM terms, constitutes an activity or task which has important or necessary consequences for the state or properties of another activity. Each function is characterised by six aspects: Input, Output, Preconditions, Resources, Time and Control, as depicted in Figure 7.

3. **Assess and evaluate the potential variability of each function.** FRAM uses a distinction between foreground and background factors, which may all affect performance variability. Foreground factors are directly associated with the functions being modelled and may vary significantly during a scenario, while background factors refer to common conditions that may vary more slowly. Both sets of factors should be calibrated as far as possible using information extracted from accident databases.

4. **Identify where functional resonance may emerge.** This step finds the possible ways in which the variability from one function can spread through the system. In case of functional resonance, the combinations of this variability may lead to situations where the system loses its capability to safely manage variability.

5. **The fifth and last step is the development of effective countermeasures.** These usually aim at dampening performance variability in order to maintain the system in a safe state, but can also be used to sustain or amplify functional resonance that leads to desired or improved outcomes.

An example of what this looks like, applied to an over-flight scenario, is in the following. The analysis of system functions (Step 2) produced the following list:

- Provide ATC clearance to pilot
- Monitoring
- Planning
- Strip marking
- Coordination
- Update flight data processing system
- Provide meteorological data to controller
- Provide flight and radar data to controller
- Controller-pilot communication
- Sector-to-sector communication

By characteristing each function using the six aspects described in Step 2 above (cf., Figure 7), the following instantiation of the model was produced.

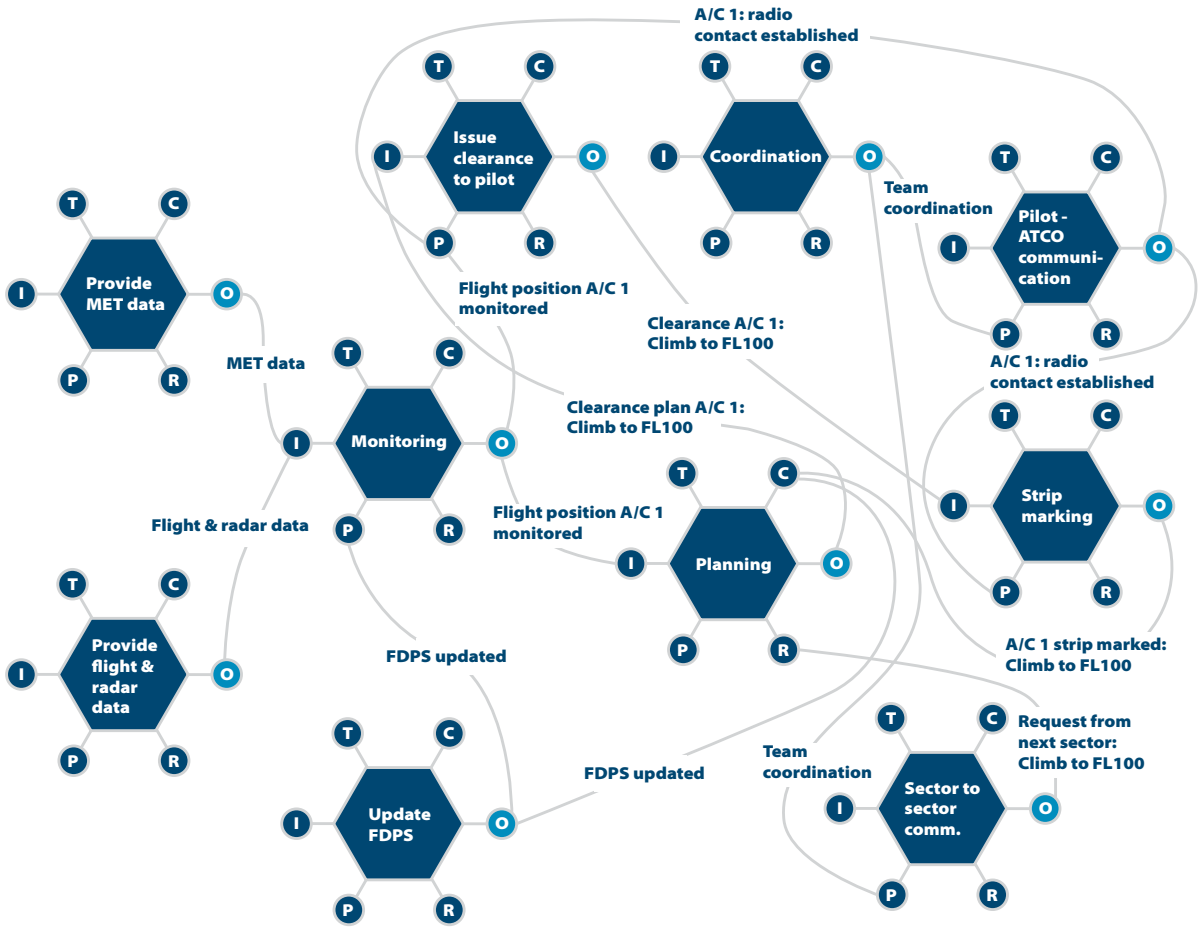

Figure 8: Instantiation of the FRAM model for the overflight scenario

HOW DOES RESILIENCE ENGINEERING FIT WITH OTHER SAFETY METHODS?

Resilience Engineering is a developing approach, one that provides an additional perspective on safety assessment and management and offers potential resilience assessment techniques, such as FRAM, to complement existing tools. Adopting a Resilience Engineering view does not require that existing practices are discarded wholesale. But it does mean that they are looked at in a different way, which in turn may change how they are applied, as well as the way in which their results are interpreted.

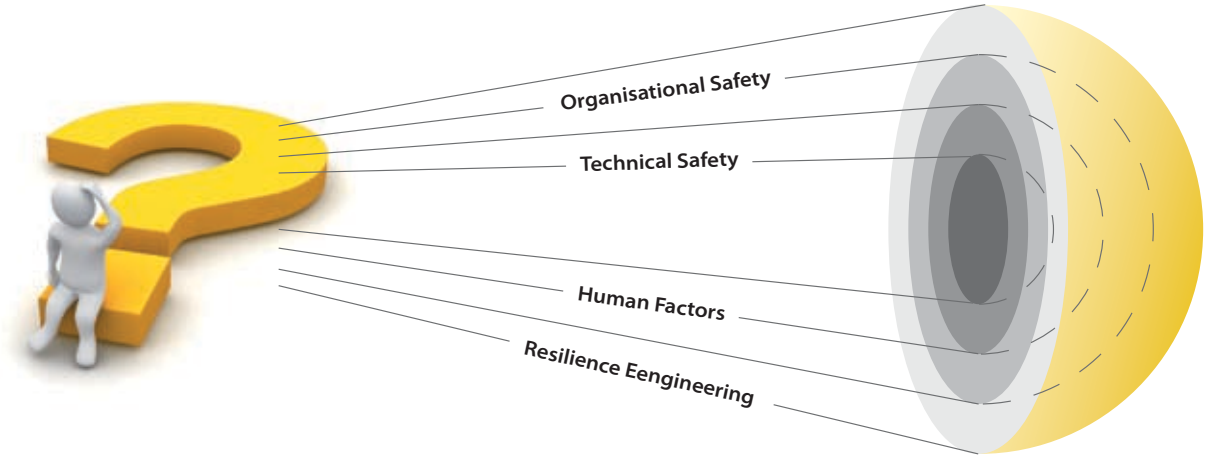

Figure 9: Widening perspectives on safety

HOW MATURE IS RESILIENCE ENGINEERING?

At present, the use of Resilience Engineering in ATM is at the feasibility stage of development. The research team working on the project will complete a first main case study for ATM in 2009, which will show how this approach works in detail and what it can deliver. The example will be a study of a Minimum Altitude Safety Warning (MSAW) system. Following that, further test cases should be conducted to develop and mature the approach to make it an additional tool for safety cases. While this inevitably will require some time, the goal is to have the methodology fit for use prior to SESAR's implementation phase in 2013.

In other domains, such as general aviation, offshore production, and healthcare, Resilience Engineering has been used longer and has provided a substantial body of knowledge and experience. This has been documented in a number of books and marked by several international symposia and workshops, as well as commercial applications, e.g. maintenance of heavy machinery, risk modeling of helicopter safety, and patient safety.

WHAT IS THE ADDED VALUE OF RESILIENCE ENGINEERING?

This key question can obviously best be answered after Resilience Engineering has been seen 'in action' in ATM. Since it addresses a difficult problem, more than one case study may be needed to demonstrate its full potential. However, there is an obvious value in having a technique which can explain and predict the complex accident scenarios that defeat existing barriers and defences, and that also tries to facilitate normal outcomes.

In summary, the questions posed at the beginning of this brochure can be answered as follows:

- Resilience Engineering goes beyond classical safety issues by recognising that safety and production are intrinsically linked. The aim of Resilience Engineering is to develop models and methods to improve an organisation's ability to succeed under varying conditions, thereby enhancing both safety and daily operations.

- Resilience Engineering is needed because many current systems are underspecified, hence exceed the premises of existing safety analysis methods.

- Performance variability enables individuals and organisations to adjust what they do to match current resources and demands. Performance variability is the reason why things go right, and sometimes also the reason why things go wrong.

- Resilience Engineering complements existing safety methods. It offers a different perspective, but is not intended to be a wholesale replacement.

- Resilience Engineering has been actively developed since 2004, and has been successfully applied to accident investigation and risk assurance in several industrial domains, eg. maintenance of heavy machinery, risk modeling of helicopter safety, and patient safety.

This has been documented in a number of books and marked by several international symposia and workshops, as well as commercial applications, e.g. risk modeling of helicopter safety, patient safety, and maintenance of heavy machinery. In the latter cases one technique was to implement a new alarm process to identify signs that things were getting worse, as a basis for timely remedial interventions.

CONCLUSION

This White Paper has presented the main concepts and practical principles of Resilience Engineering, a developing field which will be important for ATM safety in the future. ATM is among the socio-technical systems that have developed so rapidly that established safety assessment approaches are increasingly challenged to address all the issues and identify all the risks. In complex socio-technical systems, things may go wrong in the absence of manifest failures and malfunctions, and outcomes may often be disproportionately large. To safeguard against such developments, it is necessary to have tools and methods that can deal with the underspecification and tight couplings of complex, highly interactive systems such as ATM. Resilience Engineering offers a conceptual and methodological basis for achieving that goal – it is 'one to watch.' The research and development efforts will continue and further results will be documented and disseminated as a way of demonstrating the added value of the approach. Special emphasis will be put on case studies and guidance on how a smooth integration with conventional safety assurance schemes can be accomplished.

Some additional reading is suggested below.

FURTHER READING:

http://www.resilience-engineering.org/

Hollnagel, E. (2004).
Barriers and accident prevention. Chapter 5 & 6. Aldershot, UK: Ashgate publishers.

Hollnagel, E., Woods, D. D. & Leveson, N. (2006).
Resilience engineering: Concepts and precepts. Aldershot, UK: Ashgate publishers.

Hollnagel, E. (2009).
The ETTO principle: Efficiency-thoroughness trade-off. Why things that go right sometimes go wrong.
Aldershot, UK: Ashgate publishers.

Woltjer, R. & Hollnagel, E. (2008).
Modeling and evaluation of air traffic management automation using the functional resonance accident model (FRAM).
8th International Symposium of the Australian Aviation Psychology Association, Sydney, Australia.

GLOSSARY

Intractable: A system which cannot be described in every detail and where the functioning therefore is not completely understood. Intractable systems are only partly predictable.

Performance variability: The ways in which individual and collective performances are adjusted to match current demands and resources, in order to ensure that things go right.

Resilience: The ability of a system to succeed under varying and adverse conditions.

Resonance: A principle that explains how disproportionate large consequences can arise from seemingly small variations in performance and conditions.

Safety culture: That assembly of characteristics and attitudes in organisations and individuals which establishes that, as an overriding priority, safety issues receive the attention warranted by their significance.

Serendipity: The making of happy and unexpected discoveries by accident or when looking for something else; such as discovery.

For further information contact:

Project Team

Jörg Leonhardt, holds a Master Degree of Human Factors and Aviation Safety from Lund University, Sweden. He is Head of Human Factors in the Safety Management Department of the German Air Navigation Service Provider, DFS.
He is Project Leader the EURCONTROL's FARANDOLE Project, "A resilient approach to evaluate the human contribution to system Safety" and member of the EUROCONTROL HERA User Group.

Joerg.Leonhardt@dfs.de

Erik Hollnagel is Professor and Industrial Safety Chair at MINES ParisTech (France) and Visiting Professor at the Norwegian University of Science and Technology (NTNU) in Trondheim (Norway). He has for many years worked at universities, research centres, and industries in several countries and with problems from several domains, including nuclear power generation, aerospace and aviation, software engineering, healthcare, and land-based traffic. He has published widely and is the author/editor of 17 books, including three books on Resilience Engineering. The latest title from Ashgate is "The ETTO Principle: Why things that go right, sometimes go wrong."

erik.hollnagel@mines-paristech.fr

Luigi Macchi is a PhD student in the Industrial Safety Chair of the Mines-ParisTech University (France). He holds a Psychology degree from the Università degli Studi di Torino (Italy). His PhD adopts the Resilience Engineering perspective and aims to develop a safety assessment methodology accounting for the human contribution to Air Traffic Management safety.

luigi.macchi@mines-paristech.fr

EUROCONTROL Point of Contact

Dr Barry Kirwan leads Safety Research and Development in EUROCONTROL. He has degrees in Psychology, Human Factors and Human Reliability Assessment. He has worked in the nuclear, chemical, petrochemical, marine and air traffic sectors of industry, and lectured at the University of Birmingham in Human Factors. He was formerly Head of Human Reliability at BNFL in the UK nuclear industry, and Head of Human Factors at NATS (UK). For the past nine years he has been working for EUROCONTROL, managing a team of safety researchers and safety culture specialists at the EUROCONTROL Experimental Centre in Bretigny, near Paris. He has published four books and around 200 articles. He is also a visiting Professor of Human Reliability & Safety at Nottingham University in the UK.

barry.kirwan@eurocontrol.int

EUROCONTROL