EUROCONTROL

The Systemic Potentials Management: Building a Basis for Resilient Performance

A White Paper

SUPPORTING EUROPEAN AVIATION

FOUNDING MEMBER SESAR JOINT UNDERTAKING

NETWORK MANAGER

Contents

"It is no good getting the right answer to the wrong questions: you've got to get the right question before the right answer can be of any use".

(Attributed to Lord Mountbatten)

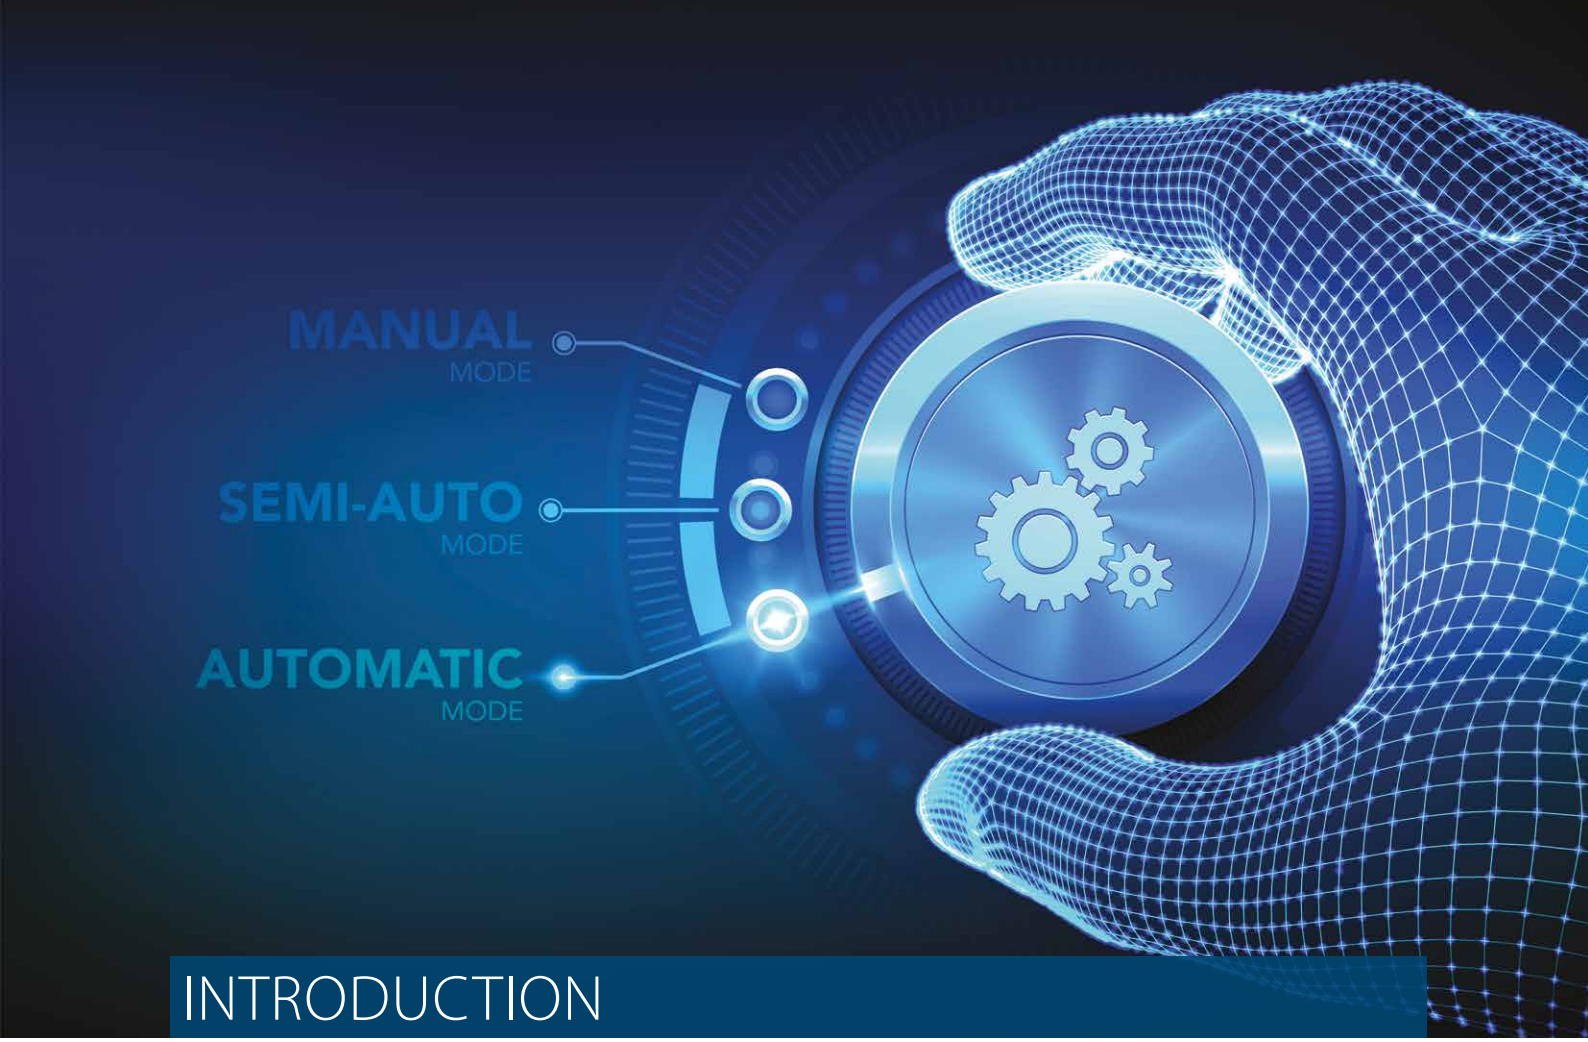

INTRODUCTION

Automation in Air Traffic Management

Air navigation services are characterised by a high and steadily increasing level of automation. Here, as well as in other industries, automation projects generally follow two classical principles. The first is the compensatory principle, proposed by Fitts (1951), according to which functions should be allocated so that humans and machines do what they respectively are best at. Or as FABEC (Functional Airspace Block Europe Central) puts it "Let ATCOs focus on the real, challenging work, to do what they are the best at, and leave the routine work to the machine". The compensatory principle has guided automation efforts for many years, but has also received its share of criticism (Dekker & Woods, 2002). The second principle is the gradual take-over of human work by automation and IT as described by Sheridan (1992). This can be seen in the progression that begins with information processing and presentation, then augments that with various types of alerts and recommendations and ends – for now – in a situation where a growing number of decisions can be made by the machine.

It is a common experience from automation projects in air traffic control and elsewhere that it can take considerable time for a project to move from the initial idea, through system design, to the operational stage. Although precise estimates are scarce, it is certainly in the order of years rather than months. During this time there will inevitably be changes in much that can affect the planned project. External factors such as time, funding, new priorities, political pressures, and public concerns are therefore just as important for managing an ongoing change as technological progress and innovation. On the way to deployment, compromises must repeatedly be made, plans and technologies must be adapted, and original ideas may become obsolete, be omitted or changed. It is very tempting – and all too easy – to trivialise anything that is not seen as directly contributing to the planned progress and pre-defined results. But failing to pay attention and prepare for unplanned disruptions is an inherent problem of automation projects and of change management in general. A higher level of automation nevertheless has consequences not just for the safety of ATM and the role of the front-line operational staff, but also for the functioning of the system as a whole.

Relying on a narrow perspective can lead to a waste of time, money and energy, and also adversely affect an organisation's potentials for resilient performance. The consequences of this have become obvious as the world-wide Covid-19 crisis has shown how important it is that ANSPs, as well as other aviation stakeholders, are able to dynamically revise their business models and strategies as circumstances change, to continuously anticipate and adjust to changes that threaten their existence, and to do so before the need becomes desperately obvious.

Weak Signals Project

Towards the end of 2011, EUROCONTROL started a project called "Weak Signals in ANSP's safety performance" in collaboration with DFS – German Air Navigation Service Provider – and several universities and research institutes. The purpose of the project was to establish what kind of information or weak signals operators[1] usually rely on. The term *weak signals* was used as a contrast to the *strong signals,* which represents the information that *a priori* is defined as necessary for effective safety management. Strong signals are what people know they must pay attention to and therefore also what they notice – as discussed by the *What-You-Look-For-Is-What-You-Find* principle (Lundberg, Rollenhagen & Hollnagel, 2009). Weak signals are sources of information that are not covered by traditional Safety Management Systems, comprising what they are not prepared for, what they do not expect, and what they otherwise fail to notice.

The strong signals represent the information that is acknowledged as necessary – either based on a theory or based on experience. Strong signals are especially the distinctive and disruptive events – usually in the form of reported accidents and incidents – that should be reduced or eliminated. Such strong signals are well-defined, attract attention and are therefore difficult to miss. Weak signals, on the other hand, are ambiguous and characterized by a low visibility because they have limited predictability, fail to attract attention, and therefore often go unnoticed or get rationalised away. The weak signals comprise the many small *events* that lie below the threshold of reporting or severity, but also the usually unacknowledged *performance patterns* – the habits, the routines, and the common trade-offs – that most of the time lead to the expected outcomes, but which every now and then give rise to unexpected and unwanted results. The weak signals are therefore very similar to the "dynamic non-events" that Weick (1987, 2011) argued were the foundation of reliable performance.

The Significance of Weak Signals

Signal detection theory is concerned with the ability to differentiate between information-bearing patterns (called signals) and random patterns that distract from the information (called noise). This is usually treated as a question of the detection threshold, but in a psychological rather than a physical sense. Here countless studies have shown that an operator is never a passive receiver of information, but rather an active decision-maker who makes challenging perceptual judgments under conditions of uncertainty. Weak signals are therefore more about the meaningfulness of the signal rather than about the strength.

Outside signal detection theory the idea of weak signals arose in the mid-1970s in connection with the growing interest for 'strategic management' and 'strategic surprises'. Here Ansoff (1975) put forward a conceptual framework and a practical procedure that a firm operating in a turbulent environment could plan for strategic surprises by responding to weak signals. Out of the extensive literature on weak signals the EUROCONTROL and DFS project team decided to choose the following working definition from Schoemaker & Day (2009, p. 86):

> *"A seemingly random or disconnected piece of information that at first appears to be background noise but can be recognized as part of a significant pattern by viewing it through a different frame or connecting it with other pieces of information."*

A weak signal is in other words something that is missed because it is not recognised as meaningful in the current context, and it is a psychological rather than a physical phenomenon. The ability to recognise patterns is crucial for the proper functioning of joint cognitive systems in any dynamic setting. Patterns must, however, not be idiosyncratic but should represent a social consensus within the given frame of reference. Patterns can point to solutions, as in recognition-prime decision making (Klein, 1993), yet still allow problems to be solved individually by fitting them to the current conditions and context (Woods et al., 2021). System design and system management must therefore pay attention to how people make use of patterns or weak signals in their work.

One consequence of this definition of weak signals is that the same piece of information would be meaningful rather than random, were it properly placed in context. Another is that experienced operators have learnt to notice weak signals, which is why it is important to pay attention to what actually happens when work goes well. Weak signals are, however, not just the pieces of information that do not fit the apparent patterns and configurations of the available information, but also the observable regularities or patterns in how activities are carried out. These regularities in turn arise from habits, routines, roles, or rituals that may or may not be recognised by the performers themselves. Weak signals are furthermore not only important for front-line operators, but play a role in what people do everywhere. Cases where the outcome is an unexpected useful discovery are admired as examples of serendipity. But in most other cases the use of weak signals is simply called recognition, as in the ways people rely on their experience both to make sense of the situations they are in and to decide what to do (Klein, 1993).

1 Air traffic controllers in Tower, Approach, Area Control and Upper Area Control.

Tangible and intangible processes

Front-line operators (i.e., controllers) manage the air traffic, hence the safety of work at the celebrated sharp end. Here the process being managed is tangible in the sense that it is concrete and easy to understand.[2] The aircraft have a physical reality and their movements are subject to – and constrained by – known principles (laws of aerodynamics). Appropriate strong signals can therefore be defined, although they often need to be supplemented by weak signals. Yet exactly the same reasoning applies to people in other organisational positions, at all levels, even at the very top. The people who work at the equally famous blunt end can be seen as managing the organisation as such, both to ensure that the primary activity (e.g., air traffic management) is acceptable and to ensure that the secondary activity (how the organisation itself performs) is acceptable. In the latter case the process being managed is intangible in the sense that it is vague or abstract, that it has few concrete correlates or concomitants, and that it is therefore difficult or impossible to define or understand. Here there is little basis for defining strong signals, so weak signals will have to suffice.

The bottom line is that everyone, regardless of what they do and what their role in an organisation is, have to rely on a combination of strong and weak signals, simply because it is impossible completely to specify what the necessary and sufficient strong signals are. The weak signals are pieces of information and patterns of performance that are known from experience to be essential in order to perform with the needed flexibility. Since the purpose of management is to provide the basis for the work of front-line operators, it is clearly as important to understand the Work-as-Done of managers as it is to understand the Work-as-Done of operators. Just as the weak signals that controllers use *could* and have been studied, the weak signals that managers use *should* be studied. The practical question is what the important weak signals for management are and how they can be made easier to recognise.

Weak signals, JND, and change blindness

In the 1860s the psychological study of perception, also known as psychophysics, introduced a concept called the just-noticeable difference (JND). The JND is the amount something must change in order for a difference to be noticeable or detectable at least half the time it happens. If the change is smaller than the JND, it is likely to go unnoticed, hence become a weak signal.

As a simple example, consider the two alarm panels shown in Figure 1 below. In Panel A there is just a single alarm, and it is therefore easy to notice when it is activated. But if the same alarm tile lights up in Panel B, it may be less easy to notice because there already are 22 other active alarms. What is a strong signal in one condition becomes a weak signal in another.

The JND provides an analogy to weak signals in change management. It cannot be used literally because system performance is about distinct outcomes and qualitative changes, whereas psychophysics is about quantitative – and mostly continuous – changes of inputs. In change management, the issue is not the sensitivity of some sensory organ, but rather the ability to notice whether a change of some kind has taken place. Going beyond psychophysics, *change blindness* describes what happens when a distinct change is introduced in a visual field or stimulus but the observer fails to notice it, often because it coincides with some other change (Simons & Rensink, 2005). Since this is a dominant feature of visual perception, it partly explains the weak signals problem for controllers. But beyond the deficiencies in noticing changes in the information that is presented and stable in a situation, there is also an issue with noticing changes that happen over time.

Spatial and temporal patterns

When weak signals are defined as random pieces of information, the allusion is that some of these can be

Figure 1: Noticeable and unnoticeable differences

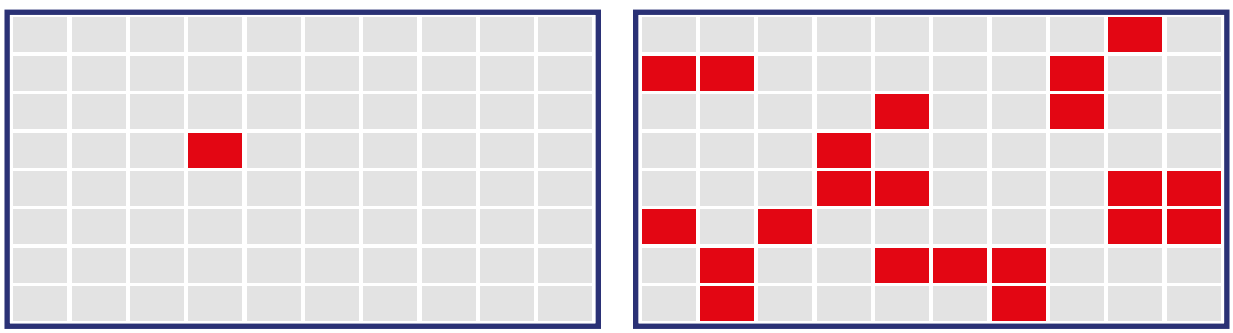

A B

2 The literal meaning of tangible is something that is perceptible by touch.

Figure 2: Illustration of six Gestalt principles

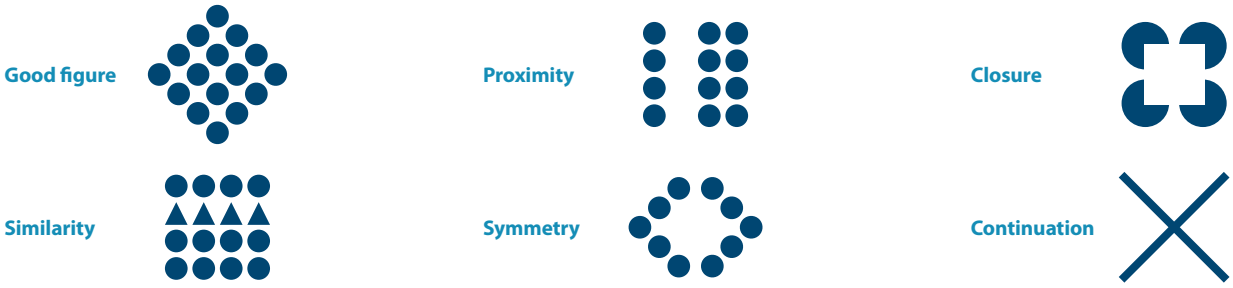

missed because of the inability to recognise an overall pattern or because of the failure to notice a change. The weak signal is in this case something that does not immediately correspond to a *Gestalt* – or even worse, something that is mistakenly perceived as being present because it corresponds to an expected *Gestalt* or configuration even though it is not. *Gestalt* usually refers to a spatial organisation or configuration of individual items (of information) coming from various sources or placed at different locations. Also, Gestalt psychology has formulated a number of principles or laws – proximity, similarity, figure-ground, continuity, closure, and connection – that describe how the human brain perceives visual elements. Six well-known Gestalt principles are illustrated in Figure 2.

But a pattern or Gestalt may also be temporal, such as the order in which different things happen. As long as things happen reasonably fast, relative to the human ability to keep things in mind, temporal and spatial patterns can be considered as analogous. But if things happen very slowly, temporal patterns gradually disappear. This is easy to illustrate by music. If a melody is played very slowly, it is nearly impossible to recognise even when the listener knows what it is. An extreme example of that is the composition by John Cage entitled As Slow As Possible. The music is shown in Figure 3 – here spatial patterns are clearly recognisable, even for people who cannot read music. But when played very slowly, the patterns disappear because we are unable to hear adjacent the sounds "together" – almost as a kind of temporal change blindness. (The piece is currently performed at an organ at the St. Burchardi Church in Halberstadt, Germany. The performance commenced on September 5, 2001, with a pause lasting until February 5, 2003. The first chord was then played until July 5, 2005. The planned duration of the performance is no less than 639 years!)

Another example of temporal patterns is provided by time-lapse photography. This is a technique where the frequency at which film frames are captured is much lower than the frequency that is used to view the sequence. When played at normal speed, time appears to be moving faster and very slow developments, such as the blossoming of a flower – or even of paint drying? – thereby become visible.

More generally, if events fall within the span of working memory, they can be perceived together, in the sense that when the "last" event happens, the "first" can still be remembered.[3] The working memory may be thought of as a mental frame, where the leading edge is the "now", and where it only is possible to perceive what is within the frame. This can be illustrated as shown in Figure 4. The primary activities take place on a time-scale measured in minutes or hours and refer to what front-line operators, for instance ATCOs, must pay attention to and manage. This will typically be a tangible process where something happens so frequently that it is easy to recognise possible patterns.

Figure 3: As Slow As Possible

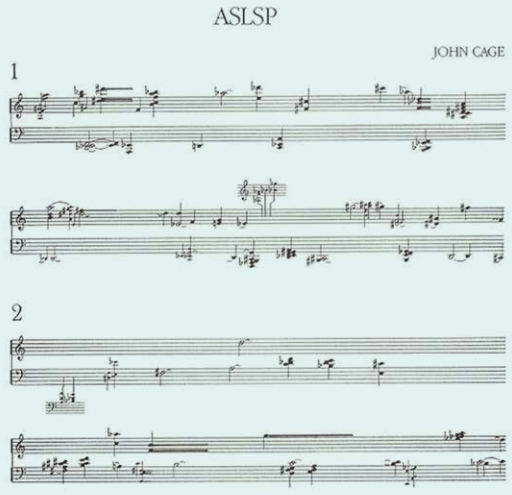

Extract from ASLSP by John Cage
© 1985 by Henmar Press Inc., New York
Reproduced by permission of the Publishers

3 The precise nature of working memory and its temporal capacity is no simple matter, cf., Cowan (2008). For the present discussion we simply acknowledge that there is a limit to how much a person can consider "together" and propose that the working memory span is in the order of 2-3 hours. Things that happen with longer intervals do not "exist" together in working memory and patterns can therefore be difficult or impossible to see.

Figure 4: Working memory span compatibility with primary activities.

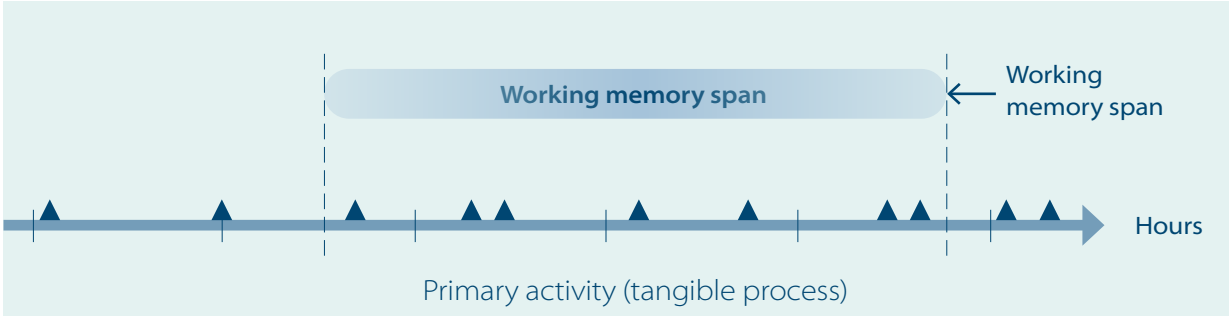

In the case of secondary activities that take place over weeks or months previous events will typically have been forgotten when a new event occurs, as illustrated by Figure 5. Here the frame that represents working memory is too small to contain more than one event at the time. Because of that, it is very difficult to see events "together", unless something can be done to facilitate that. The temporal patterns do exist, but outside the range that we can naturally comprehend.

Temporal patterns are at least as important as spatial patterns when it comes to understanding the weak signals that people rely on in their work. Noticing the order or sequence in which things happen is the basis for inferences about cause-effect relations, hence for finding patterns in what happens. We know that people act with purpose or intention, and that this determines the order in which they do things. In relation to individual performance, this becomes the *modus operandi* – the characteristics of Work-as-Done – that we can recognise in others. In relation to collective performance this becomes the way in which an organisation works both the routines of daily activities and the way it responds to unexpected developments and situations. Temporal patterns are therefore essential sources of information about performance, for individuals as well as organisations.

For the management of tangible processes – such as controlling the movement of aircraft – the speed by which things happen, and therefore also the extent of the temporal patterns, is determined by the nature (and dynamics) of these processes. Since people only can manage the processes for which they can recognise temporal patterns, it follows that their performance in turn will exhibit corresponding temporal patterns that others then can perceive. The situation is, however, completely different for the management of intangible processes that are assumed to be the basis for safety, quality or reliability. For tangible processes, we know what information and which changes are important – such as the movements and positions of aircraft – while for intangible processes we do not. The underlying processes are not only vaguely defined, as illustrated by the many problems in finding meaningful strong signals, they are also slow in the sense that developments and changes take a long time – which means that there are no easily recognisable temporal patterns as discussed above. Yet since experience clearly demonstrates that safety management cannot rely exclusively on the traditional strong signals such as accidents and reportable events, it is necessary that some form of weak signals can be found. Moreover, since the weak signals cannot be derived empirically from observable performance, they must instead be derived analytically from principles and concepts such as resilience. To this, we shall turn next.

Figure 5: Working memory span compatibility with secondary activities.

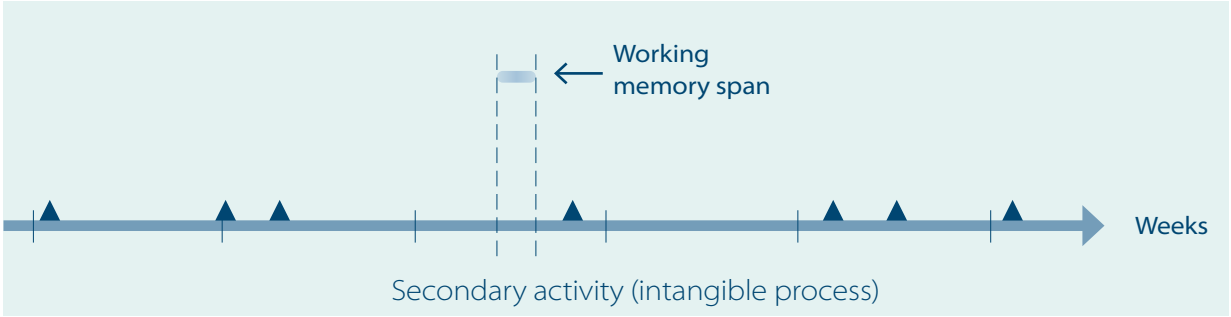

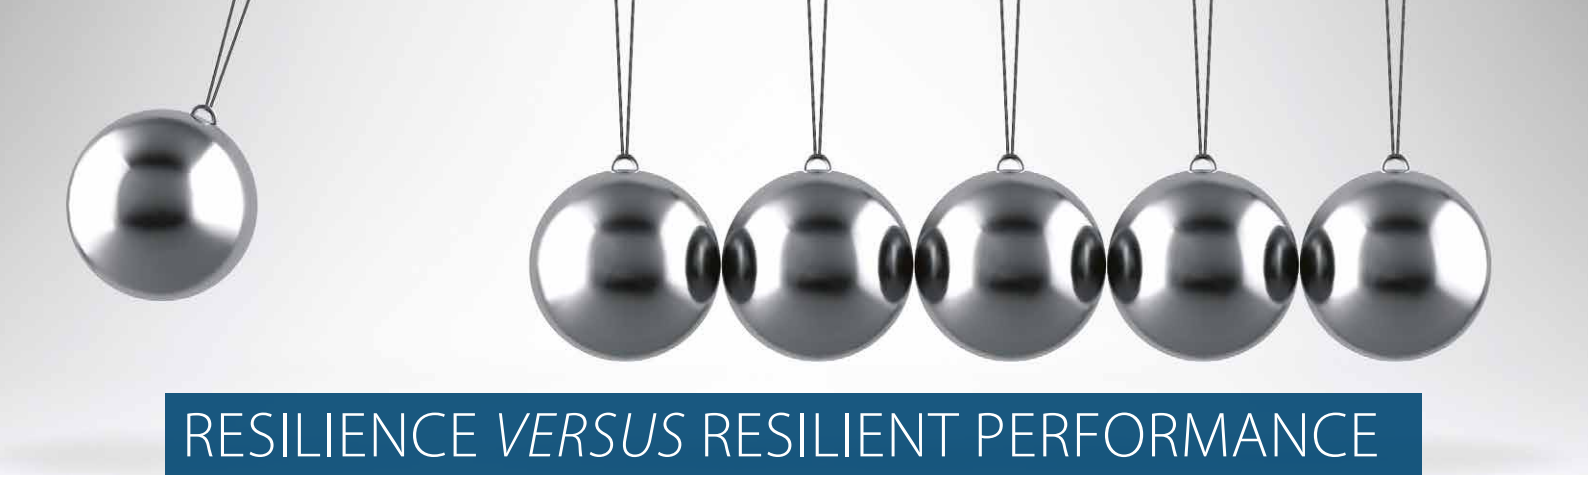

RESILIENCE *VERSUS* RESILIENT PERFORMANCE

The notion of a "resilience level" of a system or an organisation is derived from the idea that resilience is an identifiable – and measurable – property or quality of a system or an organisation. The structural level analogy is widely used for other concepts, such as level of safety culture, level of competence, etc.

When resilience became part of safety discussions, it was defined as "the intrinsic ability of an organisation (system) to maintain or regain a dynamically stable state which allows it to continue operations after a major mishap and/or in the presence of a continuous stress" (Hollnagel, 2006). This reflected the tradition of juxtaposing two states: one of stable functioning and one where a system has failed. Following the legacy of this thinking, the definition was also limited to consider situations of threat, risk or stress.

Five years after the first book, the definition of resilience had changed to "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" (Hollnagel, 2011). In this definition the emphasis on risks and threats had been reduced, and instead became how systems perform in "expected and unexpected conditions," including how such conditions could be anticipated. The focus had also moved from issues such as "safety criticality" and "responding to unplanned and unexpected sequences" to the ability to perform or function as required, not only in the face of adversity, but more importantly during normal conditions as well. Today, ten years later, a working definition of resilience might be "the ability to succeed under varying conditions, so that the number of intended and acceptable outcomes (in other words, everyday activities that go well) is as high as possible".

The definition has changed from being about resilience *per se* to being about how a system performs. Resilience is not a mystical or mythical system property or quality as such, hence not something that can be measured or managed on its own. Already the *Epilogue* of the first book argued that safety, and therefore *a fortiori* resilience, was something that a system *did* rather than something it *had* (Hollnagel & Woods, 2006, p. 347). The *Epilogue* continued to point out that "(we) can only measure the potential for resilience but not resilience itself".

The changes in the definition since 2006 have served to broaden the scope of resilient performance. It is no longer just the ability to recover from threats and stresses or the opposite of brittleness, but rather the ability to perform as required under a variety of conditions – which includes being able to respond appropriately to both disturbances and opportunities. Resilience is about how well systems perform in general, not just about how they remain safe. In addition, since something cannot go well and fail at the same time, it makes better sense to strengthen the former than to weaken the latter. The inclusion of opportunities also means that the resources spent should no longer be seen as a cost to prevent things from failing but as an investment to facilitate that things go well. Ultimately, resilience will have to be completely dissociated from safety, thereby leaving the increasingly sterile discussions and stereotypes of the past behind.

The need of a unified approach

In all sizeable companies, the management is organised by a number of departments with different purposes and ways of working. Safety is usually treated as a problem on its own that can be solved by specialised tools and methods. The same goes for quality and for other priorities or issues that a company may see as important – productivity, reliability, and so on. Yet if a company's performance is considered more broadly, any particular issue is just one side of it. It is, for instance, of little value for a company to be safe if it is not at the same time productive. Thus doubling or tripling the separation between aircraft from the usual 1000 feet vertical and 5 nautical miles horizontal would clearly reduce the likelihood of collisions, hence make flying safer, but it would have undesirable and unacceptable consequences for the productivity of both ANSPs and airlines, to say nothing of the secondary effects on societies. As an illustration, think of the responses to the eruption of the Eyejafjällajökull in 2010; here it was "safety first" until it soon became too expensive. In order to perform as required under increasingly unpredictable conditions, it is necessary to adopt a unified perspective that takes into account how traditionally separate issues are interrelated. The management of change must always consider multiple objectives and priorities, and can only be effective if these are pursued together rather than separately.

MANAGEMENT AND CONTROL

Management can be defined as the process of preparing, organising, and controlling the resources of an organisation to ensure that it can *perform as required*. The latter is important because the goals for management, hence the requirements for performance and the criteria for acceptable performance, are external to the management process itself. Management can in principle also serve three different purposes: to *maintain* a current state or position by compensating for potentially destabilising internal or external influences, to *approach* a more attractive future state or position, or to *avoid* a potentially unfavourable or destabilising future state or position. Management thus requires knowledge about the *goals*, knowledge about the current state or *position* (relative to the goals), and knowledge about effective ways or *means* to achieve the goal.

Managing something is often described using a travel or voyage metaphor. It is common to talk about how to keep or improve a position or how to get closer to or reach a target – or even to have a roadmap for the way ahead. The voyage metaphor is convenient since it clearly is essential to be able to control how something moves and changes position, whether the movement is physical or abstract and whether the subject is tangible or intangible. The metaphor is also useful because it points to the need of the three different types of knowledge introduced above and described further in the following. The metaphor is certainly appropriate for tangible systems since here it makes sense to describe movements in a physical or material sense. As this is more difficult for intangible systems, the metaphor is less appropriate here, even though it remains widely used.

- *Knowing what the position is:* Before beginning to make a change, it is obviously necessary to establish the initial position, regardless of whether the change is a movement in physical space or in a more abstract space. In addition to knowing the initial position, it is also necessary to know how the position changes, i.e., to know the current position at any particular time. Only by knowing the position at different times is it actually possible to determine whether the change (or movement) goes in the right direction and at the right rate of progress.

- *Knowing what the goal is:* In order to know whether a change goes in the right direction and at the appropriate rate, it is necessary to know what the goal or target is. This is also needed to determine when and whether the change or movement has been completed. The goal should therefore be described in practical or operational terms, preferably absolute and concrete rather than relative. While this is straightforward in the case of tangible systems and material processes (such as the movement of aircraft or the production of goods) it is less easy in the case of intangible systems and more abstract movements, such as a higher level of safety or an improved safety / learning / reporting culture, etc.

- *Knowing what the means are:* The third type of knowledge is about the effective means – how concretely to make a change, how to move from the current position to get closer to the goal. In the case of tangible systems – guiding vessels, production of goods or energy, transmission of information, energy, or materials – we usually know how to make the change because the process takes place in or by a physical system that has been designed and provided with the necessary means of measurement and control. But, few such means exist in the case of changes that refer to intangible systems, to concepts or to abstractions. What means are accessible for changing safety? For improving quality? For increasing precision or reducing delays? For enhancing the culture? The list could go on.

Another important part of controlling a change is knowing how long it will take, or at least to have a reliable estimate of it. This, of course, requires a good understanding of what actually goes on. Knowledge about how much time a change will require is essential both when detailed plans are made, when means of intervention are chosen, and when resources are set aside. It is also necessary to know about possible side-effects, about unplanned outcomes, in particular if they are detrimental. The better the system or process is known and understood, the fewer unexpected side effects there will be, and vice versa.

Understanding how something works

The effective management of a process – whether at the sharp or the blunt end – requires that it is possible to respond to whatever happens in an appropriate manner. Controllers that guide aircraft through a sector manage a mostly tangible process where the things that can happen are almost completely known, and where specific and effective interventions or responses have been prepared and are available. Managers that try to ensure the acceptable performance of an activity or of an organisation are in a comparable situation, except that they are managing an intangible process where the unexpected may happen and where it sometimes is uncertain what an effective intervention might be. Process management is also what general management does to ensure the long term existence – or even survival – of an organisation. In this case, the uncertainty of what could happen is even greater, and the availability of effective ready-made interventions therefore even smaller.

In order to manage an organisation, and indeed in order to manage or control anything, it is necessary to understand how it works – to understand why things happen the way they do, particularly because of deliberate changes and interventions. Apart from being good common sense, the need to understand how something works is also "one of the really fundamental laws of cybernetics" (Beer, 1966, p. 279). It was formulated in cybernetics in the 1940s and 1950s (Ashby, 1956), and is known as the Law of Requisite Variety. This simply states that the variety of the outcomes (of a system or a process) can only be decreased by increasing the variety in the controller of that system. Another way of expressing that is the *Good Regulator Theorem* that states that "every good regulator of a system must be a model of that system" (Conant & Ashby, 1970). In other words, if something happens in a system that can either not be recognised by the system management or for which management cannot provide a response, then control will be lost. The essence of this law is illustrated in Figure 6, where the icons are used to illustrate the (imperfect) match between environmental variety and control system variety.

Figure 6: The Law of Requisite Variety.

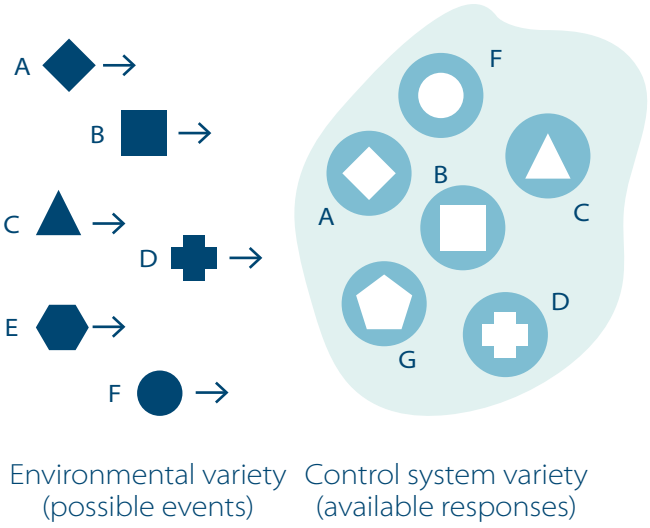

- Understanding how something happens is necessary in order to know what to do and when to do it – which means knowing how to *respond*. It is evidently impossible to control something unless an appropriate response can be made to whatever happens. This either requires that a response already exists (together with the necessary resources) – or that one can be developed before it is too late. The inability to respond, whether to two aircraft on a potential collision course, to a rapidly spreading bushfire, or to a pandemic, will eventually lead to a loss of control. Quite apart from that, no one likes to find themselves in a situation where they cannot respond to what happens.
- Understanding how something happens is necessary in order to know what to look for (signals and trends) and how often – which means knowing how to *monitor*. Although it sometimes may be possible to respond to unexpected events, "fire-fighting" is not a mode that can be sustained in the long run.

"In business organization, there are invariably more problems than people have the time to deal with. At best, this leads to situations where minor problems are ignored. At worst, chronic fire-fighting consumes an operation's resources. . . . Managers and engineers rush from task to task, not completing one before another interrupts them. Serious problem-solving efforts degenerate into quick-and-dirty patching. Productivity suffers." Bohn (2000).

It is therefore an advantage if the number of unexpected events is as small as possible. If something is expected to happen – even though the exact timing may be uncertain – then it is possible to prepare in advance or even to act pre-emptively. In order to be prepared for what may happen – at the next moment or in the immediate future – it is necessary to keep an eye on what happens in the system as well as in its surroundings.

- Understanding how something happens is necessary in order to know what the relevant experiences are and where they can be found – which means knowing how to *learn*. It is through learning that a system (or an organism) is able to change how it responds to something, by improving effective responses and by supressing ineffective responses. The essence of learning is therefore not the acquisition or accumulation of knowledge (or data or information) as such, but the effect this has on the ability to respond – as well as the ability to monitor.

  Without the ability to learn, responses would be limited to a fixed and pre-defined set. However, always responding in the same way is only feasible if conditions never change, i.e., if the world is fixed and stable. That is, however, rarely if ever the case. The same argument, of course, goes for monitoring. Learning can furthermore not be limited to what ***not*** to do (avoidance learning). It is equally important, if not more so, to learn from what works well, in order to reinforce appropriate responses. Learning is necessary to make performance more efficient, as in the progression from knowledge via rules to skills. The gradual automation of responses, the development of patterns, habits, and routines, and the recognition of weak signals are essential parts of how performance becomes attuned to the conditions – with all the risks that increased effectiveness involves.

- Understanding how something happens is finally necessary in order to know what to take into consideration and be concerned about before it happens – which means knowing how to *anticipate*. Where monitoring is concerned with keeping track of what happens here and now, and what may happen in the immediate future based on current trends and temporal patterns, anticipation is concerned with that which has not yet happened but which may. One form of anticipation deals with the intended (and expected) consequences of decisions made and of actions taken. When something is done to achieve a specific outcome, anticipation is like a "what-if" game or scenario: If we do so and so, then X, Y or Z will happen – depending, of course, on the conditions and on what others may do. Another, and perhaps more essential form of anticipation, is trying to foresee what may happen in the future, especially outside the system. Any such prediction must of course be based on a current understanding of what the situation is, including what information (indicators and trends) are the most important. It also includes trying to guess what other actors or players possibly might do: what competitors are in a market (with an active strategy), the development of a cancer or a pandemic, how a bushfire may spread, etc. In each case, anticipation is about potential scenarios and developments, hence often about unexampled events (Westrum, 2006) – something that could happen but which has never happened before. This makes anticipation both the most important of the four potentials and the most difficult to nurture and assess.

It can altogether be argued that a system, to perform resiliently (to perform as required under expected and unexpected conditions alike), needs the potentials to respond, to monitor, to learn, and to anticipate.[4] This provides the analytical basis for defining the weak signals – relatively speaking – that can be used to manage organisational change.

Scale invariance

A practical feature of the four potentials is their scale invariance. This means that they can be used to characterise performance on all system levels – from the bottom to the top and vice versa. The potentials to respond, to monitor, to learn, and to anticipate are needed for operational work as well as for strategic work – for the controllers at their positions, for the management of the control centre, as well as for the regulators.

The relative importance of the four potentials does, of course, depends on the nature of the activity or focus. For operational work such as guiding aircraft through a sector to maintain an optimal flow, the potentials to respond and to monitor may be more important than the potentials to learn and anticipate. For long-term traffic planning and air traffic transportation management, learning and anticipation may be more important than responding and monitoring. Nevertheless, at either level the potentials can be used to keep track of how well a system is able to perform, regardless of the scope and the temporal characteristics – whether it is something that happens now or something that evolves over an extended period.

It is not possible to prescribe a balance or proportion among the four potentials that is independent of the domain. For a fire brigade, for instance, it is more important to be able to respond than to anticipate. Whereas for a sales organisation, the potential to anticipate may be just as important as the potential to respond. Resilience engineering does make it clear that it is necessary for a

4 The four potentials have previously been described as the four cornerstones of resilience engineering (Hollnagel, 2009).

system to possess each of these potentials to some extent, in order to have the potential for resilient performance as defined above. All systems traditionally put some effort into the potential to respond. Many also put some effort into the potential to learn, although it often is in a very stereotyped manner. Fewer systems make a sustained effort to monitor, particularly if there has been a long period of stability. In addition, very few systems put any serious effort into the potential to anticipate.

The four potentials as weak signals

As explained above, the weak signals include the spatial and temporal patterns that people recognise and use to manage their work. In the case of intangible systems – as in change management in organisations – temporal patterns are imperceptible because they extend over time spans too long for humans to fathom. A possible solution is to apply something similar to a time-lapse approach as it is used in photography and film, except that in this case the interval between measurements must be in the order of months rather than minutes, hours, or days. If meaningful signals could be defined, this would make it possible (and sensible) to look for patterns in the ways in which an organisation performs.[5]

What constitutes a meaningful signal will of course depend on the nature and characteristics of the process being managed. In line with the above arguments for a unified approach, the signals should be relevant for multiple foci or priorities, rather than particular to just one of them. Given that it can be practically difficult to find specific signals in a system's performance, it may be better to look at that which lies behind, at what determines or shapes performance. The patterns in performance are first of all determined by the regularity (patterns) of the context, hence by the nature of the process(es) being controlled. But performance patterns are also determined by characteristics of how a system responds, monitors, learns and anticipates, similar to how observable organisational performance (artefacts and behaviours) depends on unobservable espoused values and basic underlying assumptions (Schein, 2004). In the case of slowly developing processes, such as safety, quality, etc., the potentials – or rather the assessed status of the potentials – can therefore serve as a source of information (or signals) that can be used in the management of changes, if only it is possible to assess them systematically.

Assessing the potentials

Accepting the argument that the four potentials can be used as meaningful indications for how well a system or organisation is likely to perform, the obvious next question is how they can be assessed or measured. Lord Kelvin, although speaking about electricity, is often quoted for the following view:

> *"I often say that when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind."*

Or, even simpler as "If you cannot measure it, you cannot improve it." The logic is straightforward. If you want to change something then you must be able to determine whether and when a change has taken place, and whether the intervention resulted in the expected outcomes. If the aim therefore is to manage – which means change – the potentials, then it is necessary somehow to assess or measure them.[6]

A first suggestion may be to assess the level of each potential as such, similar to assessing the level of safety culture, etc. But it is as futile to ask questions about the level of a potential as it is to ask questions about the "level of resilience". A question such as "what is the potential for monitoring" or "how high is your potential for learning", may well elicit an answer but it will hardly be meaningful – and most certainly not one that can be easily used to choose a relevant intervention. Each potential can, however, be seen as comprising a number of more specific facets or functions that are common to many types of activity. Instead of assessing each potential as a self-contained element, the potentials can therefore be characterised in terms of the several facets that a potential represents.

Background and foreground questions

Whenever something is done to change a system, such as introducing new automation, knowledge about the reasons for the change – the rationale or design basis – is needed in order to evaluate how well the change works. In cases where this information is incomplete or lacking, it is necessary to retrieve and/or supplement it. This can be done by probing the details of the system design as they pertain to the four potentials. The answers are important for developing the questions that are the core of the SPM. Since information about the design rationale can be assumed to remain stable during the change – and presumably also for some time after – it can be considered

5 Accident statistics and trends in accidents – or other reported events – are in principle a kind of time lapse recording. These events are, however, usually treated individually or statistically rather than together.

6 Strictly speaking, what we need is a way to notice and integrate signals that are widely distributed in time so that patterns can be recognised. This is not the same as measurements in the sense that Lord Kelvin intended it.

as background information that only has to be acquired once.

Background questions for the potential to respond should be directed at facets such as the justification for the list of events that needed a response, ditto for the list of response, their relevance, the threshold for responses, and the verification of responses. For instance, whether the selection of events and response is based on tradition, regulator requirements, design basis, experience, expertise, risk assessment, industry standard, or something else. Background questions for the potential to monitor should be directed at facets such as how the indicators and measurements have been selected, and how their relevance is established. Other relevant aspects of monitoring would be the validity of indicators, delays in sampling, how measurements are combined and analysed, organisational role of and support of monitoring, etc. Background questions for the potential to learn should look at the role and importance of learning in the organisation, how learning is resourced and managed, the balance between reporting and learning, how "lessons learned" should be used – and maintained – by the organisation, and so on. Background questions for the potential to anticipate should examine the purpose and potential value of anticipation, how it fits into a long term strategy or vision, whether it is an internal or outsourced function, and how it aligns with the organisational culture and values.

Other facets that relate to how the system actually performs will not be stable but likely to change – and indeed supposed to change – as a project goes forward. Many of the changes are the predictable consequences that constitute the purpose or motivation for the project. However, there will always be some unanticipated changes that could be detrimental to the purpose (Merton, 1936). The SPM uses four sets of questions – called foreground questions – to assess the changes in these facets as a means to manage the change. In contrast to the background questions, the foreground questions should be used repeatedly throughout the project.

Foreground questions should refer to what a potential means in practice. Examples of foreground questions for the four potentials are provided in Annex I. The questions are generic because they do not refer to a specific process or application. In order to assess the potential to respond it is, for instance, clearly useful to know how serious a condition must be before a response is made, or in other words what the threshold is. If it is set too low there will be many "false starts"; if it is set too high, a response may come too late. Yet another question could be to determine whether the required resources are permanently available, or whether there will be a delay before a response can begin. By continuing in this way it is possible to develop a set of questions that together provide a basis for assessing the potential to respond.

The same approach can be used to develop sets of questions for the three other potentials. In the case of the potential to monitor, it is useful to know: the relevance of the indicators or "signals" for the purpose, how often measurements or observations are made, whether there is any delay in interpreting or analysing them, how they are used (as leading or lagging indications, for instance), and how meaningful they are for the sharp end and the blunt end.

For the potential to learn, important issues are whether learning has the right focus, whether it is reactive or continuous, how the "lessons learned" are shared and used within an organisation, which priority – and resources – learning is given, etc. Finally, for the potential to anticipate, foreground questions can query the relevance of the strategy, whether it is broad or constrained, whether anticipation takes place often enough, and how the results are used – and by whom. The potential to anticipate differs from the three others – and especially from responding and monitoring – by offering fewer opportunities actually to observe how it is done. Anticipation is rarely a routine undertaking (which itself might be something to ask about), and questions about the potential to anticipate must therefore be based more on theory than on practice.

Altogether, the four sets of questions provide the core of a method called Systemic Potentials Management (SPM). The systemic potentials effectively describe or define what an organisation ought to do – the performance criteria – and answers to the questions (the potential performance profile) – consequently represents how well these criteria have been achieved.

Developing Detailed Questions

The SPM is intended as a tool for managing how well an organisation performs and how changes are implemented. Since the generic questions do not refer to any particular domain or process, a first step is to develop questions that are *specific* for the organisation or change being considered. The questions should be about operational aspects, about how something is being done rather than about how someone thinks about it or likes/dislikes it. The generic foreground questions can serve as a starting point, but only questions that are considered relevant for the actual purpose should be used. The generic questions can and should obviously be supplemented by additional questions based on knowledge about the organisation and the nature of the intended change. The questions should be *diagnostic*, i.e., concerned with problems or issues that are known to be relevant and therefore meaningful to assess. If the generic questions seem to miss known important issues, other diagnostic questions should be added.

When a question is asked it is important that it is meaningful but also that the answers are concrete

and practical. In the context of the four potentials the questions should be *formative,* so that the answers can be used directly to choose an appropriate intervention. Questions should be as concrete and operational as possible, and preferably refer to recognisable activities or functions, which can be assessed in terms of whether they are adequately performed – as operations rather than as subjective impressions. The answers can then be used as a starting point for proposing specific actions or changes to improve – or maintain – the conditions, keeping in mind that the potentials are not independent of each other.

To illustrate that, consider for instance the following statement:

"The organisation learns from reported events."

The answers would typically be either assenting or dissenting. This would indicate the general attitude to organisational learning among the responders, but it would not be directly useful to develop an effective intervention.[7] The statement could, however, also be expressed as follows:

"There are sufficient resources to write reports."

or

"Submitted reports are being investigated sufficiently."

In this case it is far easier to think of what to do, in the case of an assenting answer to maintain conditions and in the case of a dissenting answer to improve them.

Yet another example is the second question in Table 6 (Annex II), which is

"Are the set of indicators regularly revised?"

If the answers show this not to be the case, i.e., that the respondents disagree that indicators are revised regularly and properly, then that is a good basis for deciding what to do, even though the choice of means requires detailed knowledge about how the organisation in question works rather than off the shelf solutions.

How to formulate the questions

The foreground questions shown in Annex I have all been formulated in the same manner, but it is obviously possible to formulate questions in many different ways, for instance as statements that respondents can agree or disagree with:

"We revisit and revise our list of events and action plans on a systematic basis."

or

"The period covered by the lagging indicators is appropriate."

or

"The employees are being motivated to write reports."

In either case the answers can be binary ("Yes" or "No") or graduated by using a Likert scale. Questions can also be open-ended or closed-ended, negative ("Indicators are not revised regularly and properly") or positive ("Submitted reports are being investigated sufficiently"). Questions should generally be descriptive (referring to what goes on or happens) rather than relational (speculating about how something relates to something else) or causal (speculating about what-if relations).

Questions should be formulated so that they are meaningful to respondents and easy to answer, either because they refer to something that is part of the respondents' competence or experience, or knowledgeable about in general. *But most importantly, a question should not be asked unless you know how to use the answer*.

Administering the Questions

The purpose of the systemic potentials is to support the management of organisational change rather than to be used as a benchmark. It is therefore essential that the questions are administered not just once, but at regular intervals. The time between surveys must depend on the nature of the process or change being managed and an understanding of how rapidly (or slowly) changes take place, just as the frame rate in time-lapse photography is determined by the speed of what is being recorded. The answers to the questions provide information about an organisation's "position", and that information must be updated as the situation changes.[8] How frequently the questions should be used depends on the nature of the change, but once every 3-6 months could be a reasonable interval for many organisational changes. The frequency should match the (assumed) rate of change being considered – either a change being deliberately made or the changes/fluctuations in external conditions that require management interventions. The number of times the SPM should be used or the overall duration also depends on the nature of the change, but it may easily be in the order of years rather than months.

7 If the answer is dissenting, it would clearly be necessary to do something. But the statement is not specific enough to help with that.

8 Consider, for instance, how long it will take to change the safety culture of an organisation.

Figure 7: A Profile as a Radar Chart

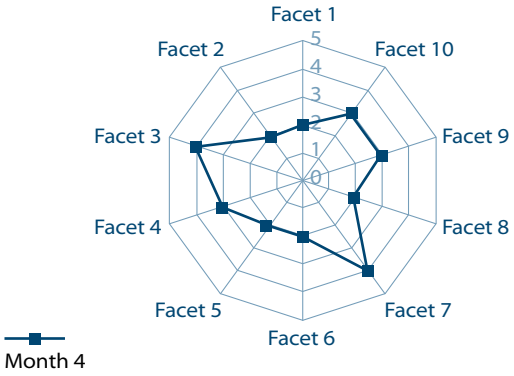

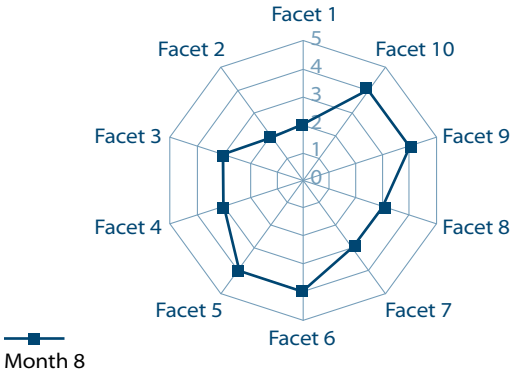

Since the assessments are intended to support the management of change, it is important that they are reliable. This can be ensured by having, as far as possible, a stable group of respondents. They must clearly be people who are directly involved with the functions that are being managed, but it can also be useful to have people in different roles and organisational positions – or even outside experts as a way of checking for potential biases in the answers. The number of respondents should not be too large and it should also be easy, quick, and convenient to answer the questions.

Interpreting the Answers

Because the SPM is a tool for change management, the answers cannot simply be compared to a standard or external reference, but only to previous answers from the same organisation, or between departments of the same organisation (e.g., between two ACCs or two TWRs).[9] By comparing answers from repeated applications of the question sets, the answer profiles will show how an organisation changes over time.

A convenient way to construct a profile is to use the so-called radar charts or nets that are part of most analytics packages. If the answers are given on a Likert scale it is simple to produce the corresponding graphics. Otherwise some kind of appropriate transformation may be needed. An example of what the answers may look like is shown in Figure 7, in this case for facets of the potential to respond. The values are for the purpose of illustration only and do not represent an actual sample. By comparing the profiles from month 4 and 8, it is easy to see which answers have become more acceptable, which have stayed the same, and which have become less acceptable. In this way, the two profiles show the change in "position" of the organisation with regard to the potential to respond. Similar profiles should be made for the other potentials.

The radar chart is a convenient way to visualise a potentially complicated relation but not actually the weak signal we are looking for. The radar chart is analogous to a frame in time lapse photography and the weak signal is the temporal pattern – i.e., the changes in the shapes of the polygons – that appears when multiple profiles are seen together. It is of course also useful to look more closely at each profile by itself, to determine whether the status of the facets are as expected and as intended. If they are not, it may be because the frequency was too high (effects have not manifested themselves yet), or because the chosen intervention did not work as intended.

The profiles of the systemic potentials not only provide a practical way to determine what the situation or position is, but they can also be used to represent what the goal or target should be. The goal or purpose of an organisational change is usually given verbally, or at best as a specific numerical value, such as zero accidents, an X% reduction in reported events, or an Y% increase in quality at a certain point in time. But if a goal is thought of in terms of specific performance targets, it may easily be transformed into a profile for any or all of the potentials. One way of showing that is illustrated by Figure 8. Here the outer profile (green) represents the goal, while the inner profile (red) shows the situation at a specific time.

Figure 8: A Goal shown as a Profile

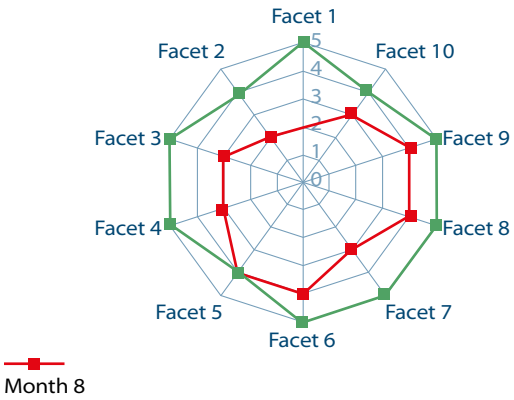

9 And also from the same respondents if possible.

The goal could, of course, also have been shown simply as a regular polygon with the maximum scoring of answers, in this case 5. That would correspond to a statement such as "the potential to respond should be as high as possible". But by differentiating between the various facets of the potential to respond, it becomes possible to express the goal in a more nuanced manner to signify that some facets are more important than others.

Couplings

A radar chart may reveal changes to facets due to underlying dependencies that have not been correctly understood – meaning that the underlying model is inadequate in the cybernetic sense discussed above.

When it comes to specifying the means, i.e., the detailed planning and implementation of interventions to bring about specific changes, it is essential to keep in mind both that the four potentials are mutually coupled but also that the facets – the more detailed functions – may depend on each other. While it may be tempting to address the potentials or their facets one by one, it is clearly not advisable. Even the brief description of the potentials given here makes it clear that they depend on each other – for instance in the way that *responding* depends on *monitoring*. Recognising and understanding these dependencies may not only be helpful in formulating the specific questions, but will also be indispensable in choosing what interventions or modifications to make. The same, of course, goes for the facets of the four potentials.

Understanding the dependencies among the four potentials – or even better, the dependencies among the facets of the potentials – requires a detailed model of how an organisation functions. While it is not possible here to explain how such a detailed model can be developed, one approach could be a functional analysis using the Functional Resonance Analysis Method (FRAM; Hollnagel, 2012). Figure 9 illustrates how this might look for a generic rather than a specific organisation. Suffice it to say, it is important to resist using a "one problem – one solution" approach in any kind of systemic management, whether the focus is safety, quality, productivity, resilience, or something else.

Figure 9: The four potentials as FRAM functions

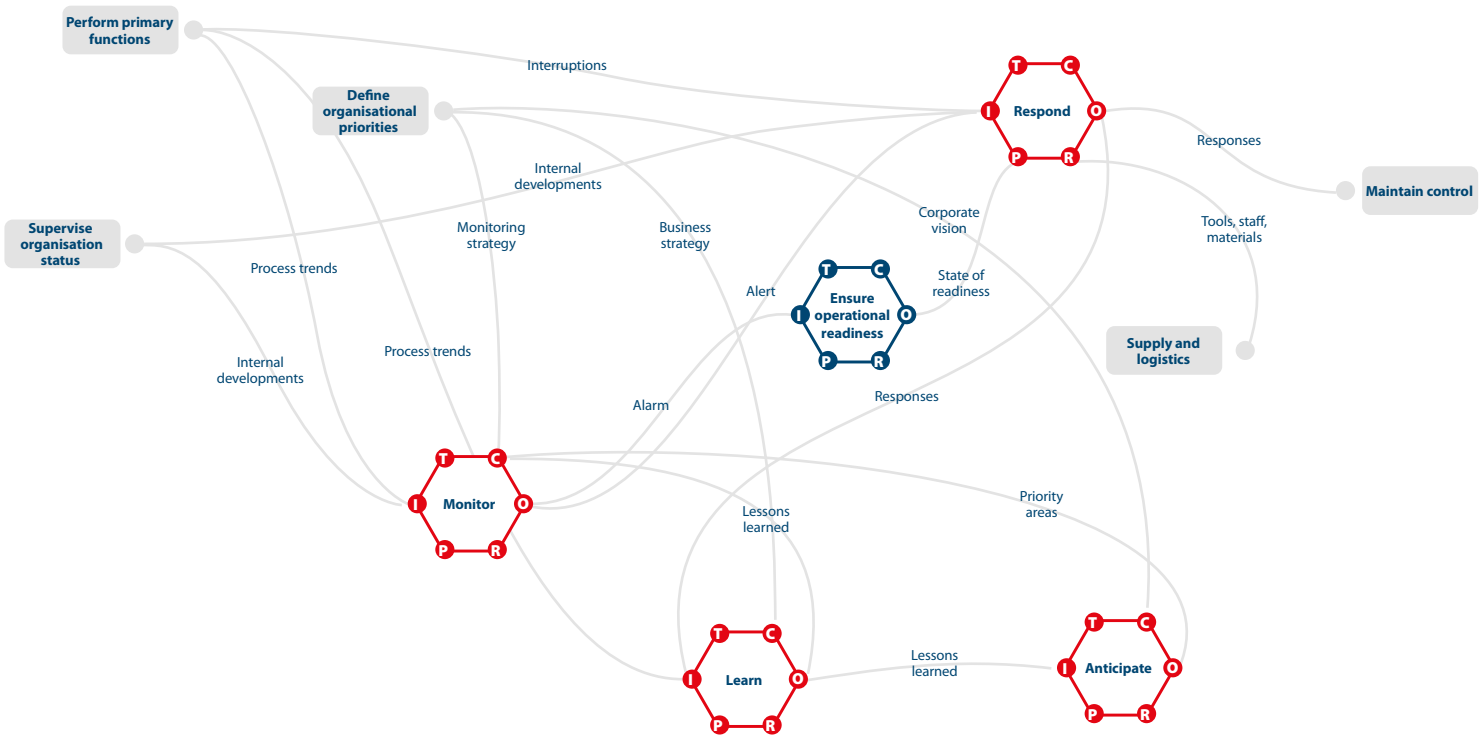

A PRACTICAL EXAMPLE

As an example of how the SPM can be used in practice, consider the general problem of automation in ATM, broadly defined as to "let ATCOs focus on the real, challenging work, to do what they are the best at, and leave the routine work to the machine". To accomplish this, an automation project must carefully consider how work should be distributed between controllers and automation, so that the overall performance of the system improves.

To follow the progress of introducing automation for this purpose, it is necessary to consider how it will affect the ATCO's overall performance. Take, for instance, an arrival manager (AMAN) defined as "a system that calculates a planned arrival flow with the goal to maintain an optimal throughput at the runway, reduce arrival queuing and distribute the information to various stakeholders" (EUROCONTROL, 2010). In this case, automation may be for "information acquisition", "information analysis", "decision and action selection", "action implementation" – or for several of these together. The automation will directly affect one part of the work, managing arrivals, but will also indirectly affect the rest. An expected positive outcome is that more time available for other parts of the work will lead to improvements to these in the long term. A possible undesirable outcome is that controllers gradually may lose competence in the parts that have been taken over by automation and that this may have second-order effects on what else they do. For any changes to people's work it is essential not only to focus on making the specific change itself, but also to make sure that people are kept informed about what is going on.

The rationale for the automation and the basis for design decisions should be known when specific SPM questions are developed. The rationale may also reflect considerations such as the internal policy on use and levels of automation, regulatory requirements, risk assessment, industry standard, recommendations from incident investigations, successful R&D programmes, capacity improvement needs, etc. These and other issues are important as a background for how to evaluate and keep track of the change, but are not themselves a focus on an SPM, because they are not expected to – not supposed to – change during the project.

For the introduction of an AMAN, the potentials to respond and to monitor should be a priority. How well are people aware of what is going to happen? How well have they been prepared and trained for that? How is their feedback and experience taken into account? How do they experience the changed working conditions and the impact of automation? Here the generic foreground questions shown in Annex II can serve as a starting point, but local expertise and experience is needed both to select and reformulate those that are appropriate and to propose more specific questions.

It should also be considered how facets of the potentials to learn and to anticipate can be addressed not least, because automation is supposed to "keep staff engaged in critical decision making". An immediate focus for learning could be for the controllers to understand how the automation works, since it is safe to assume that it will not function exactly as the designers have imagined. In a longer perspective, automation is bound to have consequences for the overall work situation, hence also for how the controllers learn and improve their competence.[10] One focus for anticipation could be what controllers assume that other stakeholders know and therefore how they will act.

In terms of how often the questions should be asked and at which intervals, the introduction of automation is usually a relatively brief and well-defined affair. The first assessment should, however, not be when the system is ready to be deployed, but when the practical organisational preparations, such as special training, commences. Given that the nature of the change is rather concrete – in contrast to a change in organisational culture, for instance – the intervals between assessments should be relatively short, possible in the order of a few months. The assessments should, however, be continued for a longer time, ideally for a couple of years. A pragmatic stop criterion may be when the profiles become stable or when another major change to the working conditions takes place – unless that change also makes use of the SPM. Examples of SPM questions for the potentials to respond and monitor are provided as Tables 5 and 6 in Annex II.

10 Automation ascribes to the substitution myth, the common assumption that artefacts are value neutral and only have intended and no unintended effects. The substitution myth is, however, invalid since any change to a system necessarily disturbs an often delicate balance.

CONCLUSION

Increasing degrees of automation are inevitable in ATM as well as in practically every other service that contemporary societies require. Automation projects must be carefully managed to ensure that the intended and desired outcomes are obtained. This typically relies on traditional techniques such as detailed individual work-packages, budget, time for deliveries, distinct milestones, etc. But simply replacing "project-as-done" with "project-as-imagined" may lead to a loss of information about where the system stands and where the project will go.

There is an obvious and persistent need for effective ways to manage systems and system changes. Whilst it is necessary to understand the actual expression of unwanted outcomes, such as uncontrolled releases of energy, a traditional "find-and-fix" strategy is no longer sufficient. Other strategies, that better cope with with the issues of the complex socio-technical systems are needed. Managing a change such as development and introduction of new automation is a challenging undertaking where influences may come from both internal and external sources, e.g., time pressure, financial issues, revised priorities, or changes in partners and collaborators. Projects inevitably need to make compromises in the early phases, and rely on adjustments that may solve the actual (current) problems in the short term, but often at the cost of a clear adjustment of the desired goal of the overall project and the long-term trajectory.

Preventing this requires a thorough understanding of the system's performance, as well as a practical method to keep track of that over time. Effective system management is impossible without a thorough appreciation of the intricate couplings among the functions that are needed for acceptable performance. Think, for instance, about how many systems had to work well for you to be able to read this White Paper! We are still in need of methods that provide an effective and reliable way to examine the otherwise hard-to-define and hard-to-quantify, but nevertheless important, weak signals.

Using the SPM to assess the systemic potentials and how they change over time can give an organisation a grasp of the available capacity to absorb disturbances in different areas of operations. It is also a good basis for developing strategies to unify the diverse approaches to productivity, quality, safety and reliability that exist in today's organisations.

REFERENCES

Ansoff, H. I. (1975). Managing Strategic Surprise by Response to Weak Signals. *California Management Review*, 18(2), 21-33.

Ashby, W. R. (1956). *An introduction to cybernetics*. London: Chapman & Hall, Ltd.

Beer, S. (1966). *Decision & control*. London: Wiley.

Bohn, R. (2000). Stop Fighting the Fires. *Harvard Business Review*, *78*(4), 83-92.

Conant, R. C. & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, *1*(2), 89-97.

Cowan, N. (2008). What are the differences between long-term, short-term, and working memory? *Progress in Brain Research*, *169*, 323-338.

Dekker, S. W. A. & Woods, D. D. (2002). MABA-MABA or Abracadabra? Progress on Human-Automation Co-ordination. *Cognition, Technology & Work*, *4*(4), 240-244.

Fitts, P. M. (Ed). (1951). *Human engineering for an effective air navigation and traffic-control system*. Ohio State University Research Foundation, Columbus, Ohio.

Hollnagel, E. (2006). *Resilience – the challenge of the unstable*. In E. Hollnagel, D. D. Woods & N. Leveson (Eds.) Resilience engineering: Concepts and precepts. Aldershot, UK: Ashgate.

Hollnagel, E. (2009). The four cornerstones of resilience engineering. In: Nemeth, C. P., Hollnagel, E. & Dekker, S. (Eds.), Preparation and restoration (p. 117-134). Aldershot, UK: Ashgate.

Hollnagel, E. (2011). *Prologue: The scope of resilience engineering*. In: E. Hollnagel, J. Pariès, D. D. Woods & J. Wreathall (Eds). Resilience Engineering in Practice. A Guidebook. Farnham, UK: Ashgate.

Hollnagel, E. (2012). *FRAM. The Functional Resonance Analysis Method for modelling complex sociotechnical systems*. Farnham, UK: Ashgate.

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

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

Klein, G. A. (1993). A recognition-primed decision (RPD) model of rapid decision making. In G. A. Klein, J. Orasanu, R. Calderwood & C. E. Zsambok (Eds.) *Decision making in action: Models and methods*, 138-147.

Lundberg, J., Rollenhagen, C. & Hollnagel, E. (2009). What-You-Look-For-Is-What-You-Find - The consequences of underlying accident models in eight accident investigation manuals. *Safety Science*, *47*, 1297-1311.

Merton, R. S. (1936). The Unanticipated Consequences of Purposive Social Action. *American Sociological Review*, Vol.1, No. 6. Harvard University, Cambridge, MA.

Schoemaker, P. J., & Day, G. S. (2009). How to make sense of weak signals. *Leading Organizations: Perspectives for a New Era*, 37. MIT Sloan Management Review. 50(3). Massachusetts.

Sheridan, T. B. (1992). *Telerobotics, automation and human supervisory control*. MIT Press, Cambridge, MA.

Simons, D. J., & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in cognitive sciences*, *9*(1), 16-20.

Schein, E. (2004). *Organizational Culture and Leadership*. Jossey Bass, San Francisco.

Weick, K. E. (1987). Organizational culture as a source of high reliability. *California Management Review*, *29*(2), 112-128.

Weick, K. E. (2011). Organizing for Transient Reliability: The Production of Dynamic Non-Events. *Journal of Contingencies and Crisis Management*, *19*(1), 21-27.

Westrum, R. (2006). A typology of resilient situations. In Hollnagel, E.; Woods, D. D.; Leveson, N. G. (Eds.), *Resilience engineering: Concepts and precepts*. Aldershot, UK: Ashgate.

Woods D.D, Licu A., Leonhardt J., Rayo M., Asher Balkin E., Cioponea R,(2021) Patterns in How People Think and Work: *Importance of Patterns Discovery for Understanding Complex Adaptive Systems : A White Paper.* BE:EUROCONTROL.

ANNEX I

Generic Foreground Questions

Table 1: Possible foreground facets of the potential to respond

The set of events for which responses exist is adequate.
The set of events is regularly reviewed / revised. The frequency of such revisions is adequate.
It is clear when a response should be given. Responses are simple to make and do not require choices.
The prescribed responses are appropriate for the situations where they are needed.
Responses can be started / initiated fast enough. Responses be fully implemented fast enough.
Effective responses can be sustained for long enough.
It is clear when a response (or situation) is no longer needed.
There are sufficient resources (people, equipment, materials) to ensure response readiness. Enough resources are exclusive for the response potential. The resources are maintained at an acceptable level
The readiness to respond is regularly checked and maintained

Table 2: Possible foreground facets of the potential to monitor

The set of indicators is adequate for the purpose.
The set of indicators is regularly revised. The revisions make use of practical experience.
The ratio of 'leading' to 'lagging' indicators is appropriate. The basis for the indicators is simple to understand.
Indicators are checked with sufficient frequency (continuously, regularly, every now and then).
The indicators are directly meaningful and do not require additional interpretation.
Monitoring is recognised as an important part of the organisation's performance.

Table 3: Possible foreground facets of the potential to learn

The basis for learning (the selected events) is sufficiently broad.
We try to learn both from successes (things that go well) and from failures (things that go wrong).
Event reports are easy to understand. Event reports contain sufficient details and data.
There are well-defined procedures for data collection, analysis and learning.
There is adequate formal training and organisational support for data collection, analysis and learning.
Learning is continuous rather than event-driven.
The resources allocated to investigation and learning are adequate.
The delay between reporting and analysis results (lessons learned) is acceptable. The outcomes from learning are communicated effectively across the organisation.
The lessons learned are directed at the right level (for instance, individuals, groups, departments).
The 'lessons learned' are properly implemented (as regulations, procedures, norms, training, instructions, redesign, reorganisation, etc.).

Table 4: Possible foreground facets of the potential to anticipate

Future threat and opportunities are assessed with sufficient frequency.
Information about future events is communicated or shared within the organisation.
The organisation has a clearly formulated vision for the future.
The organisation is clearly concerned about what could happen far ahead.
Risk awareness is part of the organisational culture.

ANNEX II

Examples of AMAN Specific Foreground Questions

Table 5: AMAN Example – facets of the potential to respond

Are the prepared responses adequate for the chosen automation features?
Is there a list of automation "changes" and a "concept of operations" before, during and after the changes?
Does the concept of operations adequately describe the operational features that you experience?
Have the automation solutions been revised since the start? If yes, what triggered the review of the change?
Is it clear who is responsible for maintaining and managing the change in automation?
Are the triggering criterion or threshold for responses predictable?
Do the responses of the automation contribute to an optimal throughput at the runway?
Are the responses in case of total failure / partial failure / corruption of automation adequate?
Does the automation respond fast enough in all cases? How fast can an effective response be implemented for total failure, partial failure or corruption of the system?
Can the automation consistently sustain an effective response be sustained? Does it differ for different modes?
If the automation becomes disrupted, is it clear when it has returned to a "normal" state?
Are sufficient resources allocated to ensure response readiness in case of total failure / partial failure of the automation?
Who is responsible for maintaining the response potential? Is it documented anywhere?
Is the readiness to respond maintained well enough? Is sufficient training and instructions provided?

Table 6: AMAN Examples relating to the potential to monitor

Are the defined indicators adequate for all situations?
Are the set of indicators regularly revised?
On which basis are or will they be revised?
Who is responsible for maintaining the list?
Are enough of the indicators 'leading'?
How is the validity of an indicator established (regardless of whether it is 'leading' or 'lagging')?
Are indicators used individually or in conjunction with other indicators?
Is there any delay or lag in getting measurements/indicators acceptable?
Are the measurements made frequently enough?
Are the measurements specific to the system or part of a more systemic measurement of the ATM System at ABC Terminal Control?
Are the measurements / indicators directly meaningful and how many require analysis of some kind?
How are the results communicated and used?
Are the measured effects transient or permanent? Reliably up-to-date?
Is OPS survey schedule properly resourced? Is it a report available?
Are resources available in time for the analysis and interpretation of the indicators that are not directly meaningful?

ABBREVIATIONS

AMAN	Arrival Manager
ANSP	Air Navigation Service Provider
ATCO	Air Traffic Controller Officer
ASLSP	As Slow As Possible
ATM	Air Traffic Management
DFS	Deutsche Flugsicherung GmbH (German ANSP)
EUROCONTROL	European Organisation for the Safety of Air Navigation
FABEC	Functional Airspace Block Europe Central
FRAM	Functional Resonance Analysis Method
JND	Just Noticeable Difference
OPS	Operational
R&D	Research and Development
SPM	Systemic Potentials Management

AUTHORS

Erik Hollnagel is Professor Emeritus from universities in Sweden, France and Denmark, and Visiting Professional Fellow at Macquarie University (Australia). He has throughout his career worked at universities, research centres and with industries in many countries and with problems from a variety of domains. Erik is an internationally recognised expert in the fields of safety and complex system analysis and has helped create the fields of resilience engineering, resilient health care, and Safety-II. He has published widely and is the author/editor of 27 books as well as 500+ papers and book chapters. sensei@safetysynthesis.com

Jörg Leonhardt is Head of User Experience and Ergonomics Department at DFS- the German Air Navigation Service Provider. He holds a Master Degree in Human Factors and System Safety from Lund University, Sweden. He is project leader of the EUROCONTROL "Weak Signals" project and Co-organizer of the EUROCONTROL Human Factors and System Safety Conferences. His area of expertise is User Centred Design, Automation and the development of System Potentials Management. joerg.leonhardt@dfs.de

Tony Licu is Head of Operational Safety, SQS and Integrated Risk Management Unit within Network Management Directorate of EUROCONTROL. He leads the support of safety management and human factors deployment programmes of EUROCONTROL. He has extensive ATC operational and engineering background and holds a Master degree in avionics from Faculty of Aerospace Engineering in Bucharest. Tony co-chairs the EUROCONTROL Safety Team and EUROCONTROL Safety Human Performance Sub-group. antonio.licu@eurocontrol.int

ACKNOWLEDGEMENTS

The authors would like to thank all the members of Weak Signals II project (from DFS, NAVIAR, EUROCONTROL, and Sapienza University of Rome) for their valuable input in the realisation of this White Paper.

NOTES

EUROCONTROL

SUPPORTING EUROPEAN AVIATION

www.eurocontrol.int