http://dx.doi.org/10.5516/NET.03.2011.078

THE FUKUSHIMA DISASTER – SYSTEMIC FAILURES AS THE LACK OF RESILIENCE

ERIK HOLLNAGEL[1*] and YUSHI FUJITA[2]
[1]University of Southern Denmark, Odense, Denmark
[2]Technova Incorporation, Tokyo, Japan
[*]Corresponding author. E-mail : hollnagel.erik@gmail.com

Received December 22, 2011
Accepted for Publication May 09, 2012

This paper looks at the Fukushima disaster from the perspective of resilience engineering, which replaces a search for causes with an understanding of how the system failed in its performance. Referring to the four resilience abilities of responding, monitoring, learning, and anticipating, the paper focuses on how inadequate engineering anticipation or risk assessment during the design, in combination with inadequate response capabilities, precipitated the disaster. One lesson is that systems such as nuclear power plants are complicated, not only in how they function during everyday or exceptional conditions, but also during their whole life cycle. System functions are intrinsically coupled synchronically and diachronically in ways that may affect the ability to respond to extreme conditions.

KEYWORDS : Fukushima Disaster; Systemic Failure; Resilience Engineering; Risk Assessment; Requisite Imagination; Functional Resonance Analysis Method

1. INTRODUCTION

The Fukushima disaster, for accident is too mild a term, has already been the subject of many discussions and analyses, and will continue to be so for the foreseeable future. Some of these will try to establish the root causes, both out of a belief that such root causes can be found, and also to satisfy the feeling of justice – or responsibility – among the public, among experts, and among politicians (and probably among other groups as well). Others will debate the issues of the social responsibility of nuclear power, something that is already seen in, e.g., Germany and Italy. Yet others will re-analyse and reinterpret the events in order to find salient explanations that somehow transcend the simple notion of root causes, and point to system-wide and momentous factors that hitherto have been neglected. Here the continued scrutiny and repeated post mortems, of Tenerife, Three Mile Island, Challenger, Chernobyl, Columbia, Deepwater Horizon, etc., are good examples.

In this paper we will try something a little different, namely to look at the Fukushima disaster from the perspective of Resilience Engineering. Resilience Engineering represents a new way of thinking about safety, that, since the first Symposium in 2004, has become widely recognised as a valuable complement to the established safety view [1]. Both industry and academia have recognised that Resilience Engineering offers novel ways to confront the puzzles of complexity, interconnectedness, system of systems, and ultra high reliability. The concepts and principles of Resilience Engineering have been tested and refined by applications in such fields as air traffic management, offshore production, health care, and commercial fishing. Continued work has also made it clear that resilience is neither limited to handling threats and disturbances, nor confined to situations where something can go wrong. Today, resilience is understood as 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 [2]. This definition emphasises two characteristics of a resilient system. First, that when something happens, a resilient system does not simply try to react and recover; it also looks for how it can continue to achieve its purpose by adjusting and changing its functioning. Second, that it is prepared to deal not only with disturbances and disruptions, but with diverse conditions of functioning, expected as well as unexpected. Resilience engineering is not just important as a way to improve safety, but can also help to ensure the economic survival of an organisation when challenged by internal and external changes [3].

In order to be resilient, a system must be able to do four things: (1) To respond quickly and effectively to expected and unexpected conditions (disturbances as well as opportunities) either by implementing a prepared set of responses, or by adjusting everyday functioning, and also be able to sustain the response until control of the situation has been regained. (2) To monitor that which is, or can become, a change or disturbance in the near term, covering both what happens in the environment and what happens in the system itself, i.e., its own performance. (3) To learn from experience, in particular knowing how to learn the right lessons from the right experience – successes as well as failures. And finally, (4) to anticipate developments, threats, and opportunities further into the future, such as potential changes, novel needs, increased demands, and tighter constraints.

"The system" is, of course, a loose concept. It may refer to designers, users, regulators, stakeholders, and the public. It may also refer to the Fukushima nuclear power plant, to the owner and operator, the Tokyo Electrical Power Company (TEPCO), to the regulator (the Nuclear and Industry Safety Agency), etc. Since this paper clearly cannot look at all possible systems, we have chosen to focus on two, namely the world of engineering experts (safety experts) and the operating company (TEPCO). One represents the design and pre-analyses of the power plant, while the other represents the operations and management of both the expected and the unexpected.

2. THE EVENT ITSELF

On March 11th 2011, a tsunami caused by a gigantic earthquake hit the Tohoku Region pacific coast in Northern Japan. The seismic centre was estimated to be about 130 kilometres east of the Oshika Peninsula of Tohoku, and 24 km underneath the seabed. It extended 500 km along the coastline with a width of 200 km. The intensity of the earthquake was 9.0 on the Richter scale, making it the fourth-largest earthquake recorded since 1900.

The earthquake created a gigantic tsunami wave that was about 10 m high at maximum. Once it reached land, it ran up to 40 m above the sea level and intruded 6 km inland, causing catastrophic damage to many people and towns along the coastline. About twenty thousand people lost their lives or are still missing, the major cause of their death being by drowning.

The earthquake totally disrupted the infrastructure of the region; electricity, gas, water, and railway. And while it is commonly known that several nuclear power plants were lost, the reduction in electrical generating capacity due to the loss of fossil power plants was actually larger. All major roads in the region were damaged, making it difficult for the rescuers to reach the affected areas.

Immediately after the earthquake occurred, all the operating nuclear reactors at the Fukushima #1 plant, three out of the six, were successfully shut down. But soon after that the external power was lost because:

- The electrical line shorted out.
- The electrical switchgear and transformer went out of order.
- A power transmission tower was toppled by the earthquake, rather than by the Tsunami.

Following the loss of the external electricity supply, the emergency backup diesel generators started successfully. But approximately fifty minutes after the earthquake, the Tsunami hit the plant with the wave running up to 14 to 15 m at the perimeter of the plant. Since the emergency backup generators were located under the ground, they were flooded with seawater. Electrical equipment, pumps, and fuel tanks were washed away or damaged. As a result, the plant suffered a total loss of electrical power (i.e. Station Blackout) in what can aptly be described as a complete common mode failure.

3. THE IMMEDIATE CONSEQUENCES

The immediate consequence of the loss of electrical power was core melt at Reactors #1, #2, and #3, which in turn caused the massive release of radioactive materials into the environment. Within a few days, the reactor buildings of Reactors #1, #3, and #4 blew up, because the hydrogen that was produced inside the reactor pressure vessels leaked into the buildings and exploded. Fallen building walls and ceilings further damaged the major equipment and piping systems. This made it difficult to cool the core and the spent fuel pits via regular cooling lines, and also severely hampered visual inspection.

At the time of the explosions, a large quantity of radioactive materials was released. An increased radiation level was measured in wide areas, including Tokyo, and the contamination by caesium was recorded even further away, in the Kanto Region west of Tokyo. The quantity of radioactive materials released to the environment was enormous, allegedly several hundreds Peta-Becquerel. The rise in radiation level and the concentration of radioactive materials in the environment, mostly caesium, were however, not significant outside a 20-km evacuation area, except for the so-called hot spots. In these, the radiation level was relatively higher because of fallout caused by rainfall and special geological conditions. For some of the hot spots, the government ordered people to evacuate even if the locations were outside the 20-km forced evacuation area. In addition, the government asked people in some areas to be ready for evacuation, if it became necessary due to long-term exposure.

4. A RESILIENCE ENGINEERING VIEW OF THE FUKUSHIMA DISASTER

Before we try to apply a resilience engineering per-

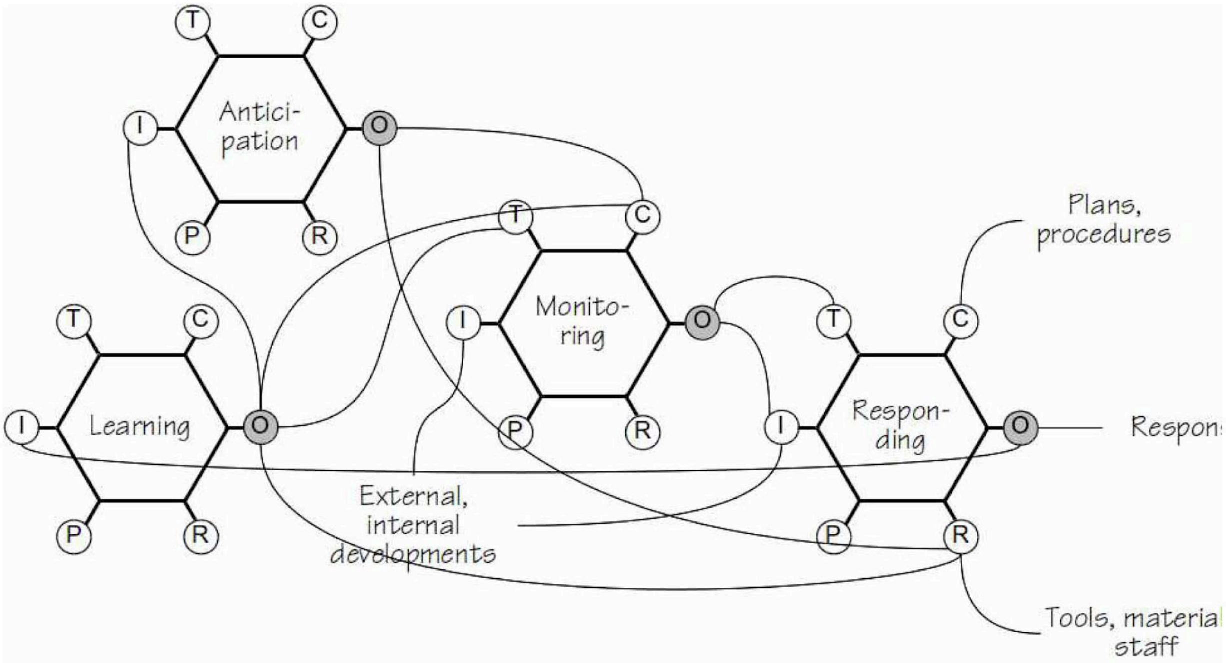

Fig. 1. Main Dependencies Among the four Capabilities of Resilience

spective to what the safety experts and the TEPCO did, it is necessary to look at the four resilience abilities a little more closely. As mentioned above, the four are the ability to respond, to monitor, to learn, and to anticipate.

- Responding means knowing what to do, or being able to respond to regular and irregular variability, disturbances, and opportunities, either by adjusting the way things are done, or by activating ready-made responses.
- Monitoring means knowing what to look for, or being able to monitor that which in the near term changes, or could change, so much that it would require a response. The monitoring must cover the system's own performance, as well as changes in the environment.
- Learning means knowing what has happened, or being able to learn from experience, in particular to learn the right lessons from the right experience.
- Anticipating means knowing what to expect, or being able to anticipate developments, threats, and opportunities further into the future, such as potential disruptions or changing operating conditions.

The four abilities offer a convenient way to characterise resilient performance, but they are of course not independent. In order to understand how a system or an organisation can be resilient, it is therefore necessary to describe the ways in which they are functionally coupled, or depend on each other. A detailed description of the functional couplings is beyond the scope of this paper, but the following should be sufficient to explain what this might mean in the case of the Fukushima disaster. The major dependencies are also shown in Figure 1, using the principles of the FRAM (Functional Resonance Analysis Method) [4]. The purpose of the FRAM is to identify and model the functions that are required to provide a specific capability, or to carry out a specific activity. Each function is described in terms of six aspects, namely Input (I), Output (O), Precondition (P), Resource (R), Control (C), and Time (T). Whenever any of the <I, P, R, C, T> aspects are specified for a function, they must also be represented as an <O> aspect of one or more other functions. Similarly, the <O> aspect of a function must occur as an <I, P, R, C, T> aspect of another function. These simple rules guarantee that the resulting model is complete.

- Responding can be triggered by external and/or internal events. Responding requires that the system is in a state of readiness, and that the necessary resources (e.g., tools, materials, and people) are available. The actual response is controlled by plans and procedures, predefined or *ad hoc*, and may require a rescheduling or cessation of ongoing actions. The set of predefined responses represents both what has happened in the past (learning), and what may happen in the future (anticipation). Responding can obviously be primed and facilitated by monitoring.

- The input to monitoring comes from external and internal developments, while anticipation and learning provide the background for selecting indicators, and prioritising and interpreting data. Effective monitoring requires that time and resources are available, that there is a monitoring strategy, and that the people involved have the requisite skills and knowledge.
- Learning should use past events and responses, either own experiences or those of others, even if these have not resulted in something requiring a response. Learning is 'controlled' by the organisation's accident model, that in practice determines which data and events are considered [5].
- Anticipation depends on what has been learned from the past (lessons learned), and is guided by the organisation's 'model of the future,' in particular the types of threats or opportunities that this model can describe. The main resource is competent people, while a precondition is the organisational culture or mindfulness [6], or a 'constant sense of unease' [7].

To illustrate how the rendering in Figure 1 should be 'read,' consider the relation between the ability to learn, shown as the hexagon labelled 'Learning,' and the three other abilities. As the figure shows, the Output from Learning is one of the Inputs to Anticipation (and the only one shown here). It is also used by Monitoring, both as a Control of what is monitored, and to define the time characteristics of monitoring. The Output is finally used as a resource for Responding. Similarly, the Output from Responding is a primary Input to Learning.

Using these considerations, it seems reasonable to take a closer look at the dependency between the ability to anticipate and the ability to respond. It should first of all be noted that it clearly is impossible to prepare a response to something that has not been considered in advance. (It may be possible to give a response, of course, but a spontaneous response will in most cases be less efficient than one that has been prepared.) And the basis for considering something in advance is either experience (learning) or anticipation, where the latter depends on the former as described above.

In the case of Fukushima, we shall focus on the engineers' ability to anticipate, and how that, in combination with the TEPCO's ability to respond (or lack thereof), laid the grounds for the calamitous situation in March 2011.

5. ANTICIPATING FUKUSHIMA

In the case of Fukushima, two uses of anticipation are important. The first is the anticipation that is – or should be – part of the building and siting of the Fukushima nuclear power plant itself (design related anticipation). The second is the anticipation that deals with the safe operation of the plant, once it has been built and commissioned (operation related anticipation).

5.1 Design Related Anticipation

The design related anticipation refers to the ability to consider what may possibly happen, especially whether there is anything that may seriously jeopardise the plant's structural and functional integrity. Westrum [8] has argued that three types of threats should be considered. Regular threats are events that occur so often that the system learns how to respond, and so often that it is cost-effective to establish a standard response. Irregular threats are one-off events. While such events may be recognised, there are so many different varieties, that it is neither practically possible, nor cost-effective to provide a standard response. Finally, unexampled events represent those occurrences that are virtually impossible to imagine, and which exceed the organisation's collective experience. What happened at Fukushima on March 11, 2011 clearly falls into that category. Design related anticipation should look at least to the irregular threats, but should also acknowledge the possibility of unexampled events.

During the initial design review of Fukushima, a scientific study was obviously made to assess the likelihood of major earthquakes. In 2004, an earthquake that was larger than the design basis hit the Kashiwazaki Nuclear Power Plant located on the opposite side of the main land, i.e., at the Sea of Japan. In this case some geological faults had apparently been overlooked in the investigation that was part of the initial site appraisal. This demonstrates that the initial appraisal cannot be complete, and that new information can become available. The case of the Fukushima disaster is another example of this.

The scale of the March 2011 Tohoku earthquake was much larger than the design basis, and the height of the Tsunami was twice what had been assumed. The tsunami wall was designed to withstand a 5.7 meter wave, which was far too little in the actual case. Does this mean that the design assumptions were wrong? In hindsight we can, of course, say that they were inadequate, but to learn anything from this disaster we need to find out why that was so.

The probability of a large earthquake hitting the affected areas was known to be very high, well before the earthquake actually hit the region. But the initial investigation apparently did not assume that more than few faults would be activated simultaneously. With the Tsunami, the assumptions were also based on a historical review, but tsunamis are few and far between. The tsunami wall was designed with a height of 5.7 meters, although the reason for that is not clear. In 2002 the Tsunami Evaluation Subcommittee of the Nuclear Civil Engineering Committee of the Japanese Society of Civil Engineers published a report on "Tsunami Assessment Method for Nuclear Power Plants in Japan". In 2008, TEPCO used this method to confirm the safety of the nuclear plants at Daiichi. But there is, of course, no way this could have influenced the design decision when the plant was built in the 1960s. After the Tsunami had happened, it became clear that a historical study had revealed that a much larger Tsunami occurred in the middle of the ninth century (estimated to be in AD 869), and that a researcher had made a strong recommendation for refurbishment of the plant in 2006. But, the recommendation was reportedly turned down for the reason that the tsunami was hypothetical, and because the claimed evidence was not accepted by specialists in the nuclear sector.

The case for a nuclear power plant design usually looks perfect – as indeed, it has to. But people tend to forget that the actual design may include assumptions that are only justified by a strong belief that the overall framework is perfect. This belief may override the scientific value of PRA (Probabilistic Risk Analysis). PRA is widely used in the nuclear industry to assess the possibility of rare events that are beyond the design basis. But a PRA involves subjective judgements that may be subject to the influence of faulty belief.

One of the authors found meeting minutes in which an expert representing an electrical company said that it might be generally worthwhile to show that nuclear power plants would be able to withstand a Tsunami. The connotation for this person was that the result of assessment was already given, before the assessment was actually conducted during the design process. In other words, the assessment was conducted to support a belief that the plant would be safe. It cannot be ruled out that such an atmosphere prevails in the nuclear industry – in Japan and elsewhere.

As this brief discussion demonstrates, design related anticipation was constrained by other concerns. Quite apart from the 'imperious immediacy of interest' [9], even a preliminary risk assessment will be constrained by time and resources. It is not very difficult to find a very large number of potential risks or threats, but there may be insufficient time and resources – or even motivation – to do so, and to evaluate them thoroughly. The anticipation is therefore constrained, often by referring to shared assumptions about what is likely and what is not. But thoroughness is a prerequisite for efficiency, not least in the perspective of a system's life cycle [10].

5.2 Operation Related Anticipation

Risk assessment is the formalised way of anticipation that has become the *de facto* standard across most industries. Institutionally, it could be claimed that the design of the Fukushima reactors met the regulatory requirements. The fact that the accidents happened could therefore mean either that the regulations were faulty, or that the design was not good enough. But more important than that was the complacency that came from a strong belief among experts that the plant design was perfect. It was therefore, in practice, unthinkable for them to imagine a situation where the plant would totally lose electricity, and consequently totally lose its cooling capability. In Westrum's terminology (op. cit.), this was not just an irregular threat, it was an unexampled event. We thus see that the neglect of scientific advice was most grave; scientists did point out that a much larger Tsunami might hit the nuclear power plant, and that the plant needed to be modified for it to be able to survive. But this was one of the conditions that were discarded, because they were seen as being too improbable.

Another problem for anticipation is the way that domain experts perceive the role and function of PRA. We all know that PRA has both advantages and disadvantages. The formal framework of PRA looks well balanced, and provides a systematic way of evaluating the potential risks associated with situations beyond the design basis. However, a PRA as such is static, and therefore unable properly to represent the dynamics of severe conditions. Another problem is that the availability of data is questionable and that the quality does not correspond to the 'beauty' of the formal framework. In particular, since 1990, domain experts have argued over the availability and the reliability of data on erroneous human behaviours [11]. Even the staunchest believers in Human Reliability Assessment (HRA), today accept that the modelling of erroneous human behaviours leaves much room for improvement. Yet despite these misgivings, PRA is still used as if it was perfectly reliable.

6. RESPONDING TO THE DISASTER

The ability to respond refers both to the day-to-day operations, and to exceptional circumstances. In this case we shall consider only the latter.

As explained above, domain experts seem to believe that PRA gives a sufficient foundation for risk management, hence for the ability to respond. We beg to disagree. What is required in really severe situations is far beyond what the plant design can provide. Consider the following example. In case of the Fukushima disaster, the electricity was totally lost. Plant personnel therefore had to find other ways of securing electricity, even if the methods were unusual. Electricity trucks (i.e., large vehicles equipped with generators, and used as a movable electrical source) were thought of as one of the last resorts. Many electricity trucks did in fact try to reach the plant, but failed to do so in time. TEPCO tried to use two trucks available at the plant, but it took too long before they could start supplying the electricity, and soon after they did the electricity-hydrogen explosion happened, which damaged the cable. At that point the electricity was irrevocably lost. During the course of these efforts, the governmental emergency management team tried to help TEPCO by transporting an electricity truck by helicopter, but found that the truck was too heavy for airborne transportation. In the end, the core was severely damaged because electrical power could not be provided in time.

It is not unreasonable to assume that the availability of electricity trucks was not thoroughly investigated in the PRA. If the PRA analysts seriously studied the total loss of electricity and countermeasures against it, they should have realized that the available electricity trucks were too heavy, and that remedial measures therefore had to be considered. In practice, however, the purpose of PRA is to assess the failure probabilities and improve on important risk factors that are found in the plant design. The purpose is not to find all the possible failures, or to consider remedial measures for severe conditions that go beyond the use of regular plant equipment. (Nevertheless, there are some cases in which supplemental systems are added to the original plant system, in consequence of a PRA, so that severe situations with relatively lower probabilities, e.g., 10E-9, can be mitigated.)

The flexibility of organisational or human characteristics can make the system well tuned to regularly recurring situations by trading-off thoroughness for efficiency. But at the same time, this trade-off will deprive the system of the resilient properties that are needed to face real world problems. One lesson learnt from the accident of the Fukushima nuclear power plant is that we need to think not only about how the plant should work and respond to regular threats, but also how it possibly could have or acquire the properties that would make it resilient. In the case of Fukushima, it appears that the assumed technical perfection was a hindrance to even thinking of acquiring these properties. In other words, overconfidence in the experts' anticipation of what might go wrong limited the ability to monitor and respond – and, to some extent, also to learn – hence impeded the development of resilience.

7. DISCUSSION

It is hardly a surprise that people in the nuclear industry generally believe in the accuracy of technical analyses and expert judgements, at least until an accident actually happens. For this reason an accident is, for most of them, really a bolt out of the blue – an unexpected and unexampled event. It is all too easy to find statements from experts who once were full of confidence in the theoretical bases of the design, as well as the rationality of their practical assumptions, but who later regretted their short-sightedness. It has repeatedly been reported by the Japanese media that the commissioner of the Nuclear Safety Commission (NSC) of Japan apologetically admitted that their judgement to exclude the consideration of a Station Blackout (SBO) was a mistake. It is known that a working group of NSC in 1993 pointed out that an SBO was possible, but that the NSC and the government did not listen to the advice.

The same phenomenon of over-confidence applies to the importance given to PRA. In the nuclear industry, PRA is used as a tool that provides a systematic way of rank-ordering important risk-dominant events that are beyond the design basis. For events that turn out to be critical as expressed by, e.g., the commonly used risk matrix, design changes or remedial improvement are made. Low probability – and low consequence – events are simply left out. PRA is thus a plant design tool that leads to an economically justifiable basis. The result is a plant where risks are "As Low As Reasonable Practicable", except that 'practicable' usually means 'affordable.'

The height of the Tsunami should have been rigorously assessed, and the calculations formally approved in the initial design. But people forgot to be mindful, and to sustain a "constant sense of unease." In other words, they forgot that previously justified assumptions should be questioned every now and then.

It is, however, not acceptable to neglect severe events just because they have a low probability. It is important also to study the probability of survival from severe conditions – e.g., the success paths [12]. Instead of focusing on what can go wrong or how something can fail, we must focus on what should go right or how something should work. Since this cannot be done by PRA, due to its focus on hazards, risks, and failures, we need to find a perspective that looks at system performance as a whole, especially how it dynamically develops and changes over time. This is what resilience engineering tries to achieve. Instead of looking at how the lack of one or more of the four abilities can lead to failure, resilience engineering emphasises how the strengthening of these abilities can lead to success, or at least to survival. Because established risk assessment methods set a limit based on what is believed to be logical thinking, and do not care about regions beyond that, anticipation becomes very constrained. This has consequences for the abilities to monitor and respond, as we have seen in the Fukushima example.

A system cannot be resilient, i.e., be able to cope with unexpected severe situations, if it is limited to designed features alone. Over-confidence limits the engineers' imagination, and impedes the justification of additional remedial measures on both economical and philosophical grounds. Such over-confidence is widespread among experts in the nuclear industry, and in other industries as well – offshore production being the most spectacular example. The problem is thus unfortunately not new. But because it is old it has become dormant, and is therefore not easily visible.

The trivial lesson from the study of rare events is that they are very difficult to anticipate. Almost as trivial is the lesson that these events are due to a non-trivial combination of multiple factors, many of which are considered irrelevant to safety. In the traditional safety perspective these factors are described, e.g., as active failures – particularly 'human errors' – and latent conditions. Resilience engineering sees them differently, namely as the variability of everyday functioning – either on the general level of the four abilities, or on the concrete level of everyday performance. This variability is not confined to the operational state, but exists throughout the life cycle of the installation – from design to decommissioning. What we can learn by looking at an accident from the resilience engineering perspective is, thus, how differences in emphasis of (or priority assigned to) the four abilities can be used to understand how potential – and actual – weaknesses in the system arise, and where they can be found. The purpose is not to determine what failed, or look for who is responsible. The point is rather to try to understand, as best we may, how complicated such systems are, not only in their momentary functioning – whether it be during everyday or exceptional conditions – but through the extended existence or life cycle. Resilience can be seen in the way the organisation responded when the event happened. But resilience can also be seen in the way the organisation functioned in the preceding years. It is only in this context that we can begin to glimpse – but perhaps not yet fully see – how choices made at one stage influence what happens and what can be done later on. What is gained from such an exercise is not just an appreciation of how complicated things are, but rather an understanding of how functions are coupled, and how these couplings may affect the ability to respond to extreme conditions. Once this has been understood, even tentatively, it is possible to work constructively on how to improve matters.

8. CONCLUSIONS

The purpose of this paper has been to demonstrate how a resilience engineering perspective can supplement the traditional approaches to understanding industrial safety – both when it works and when it fails. The case in point has been the disaster at the Fukushima-Daiichi nuclear power plant in Japan, and resilience engineering principles have been used to describe how shortcomings in the ability to anticipate, both during design and during the response to the natural disaster, laid the grounds for the unfortunate outcomes.

The main conclusion is that formal risk assessments trust established methods and models more than they should. Established methods and models have become accepted in practice, because they seem to offer an acceptable trade-off between thoroughness and efficiency. In other words, they seem to offer the necessary thoroughness of analysis, meaning that they identify all the risks that are ‘necessary,’ but without being unnecessarily costly in time or resources. This happens in every field of activity, and examples can easily be found in finance, in engineering, in medicine, and in offshore exploration. Because severe accidents are very rare, we easily become overconfident in the analysis methods. The reasoning seems to go something like this: We have analysed the possible risks; we have built the installation following the recommendations; and we have operated safely for *n* years – whatever *n* is. This reasoning is, however, fallacious, because the absence of a failure does not prove that the precautions were correct, or even sufficient. Resilience engineering advocates a constant sense of unease, that we should be mindful of what we do, to counteract the overconfidence that is a side effect of the relative safety of nuclear installations.

Resilience engineering provides a way to identify the capabilities that a complex socio-technical system must have to perform acceptably in everyday situations, as well as during accidents. In that sense, it is a complement to other engineering methods, rather than a replacement. Resilience engineering provides a contrast to classical risk assessment methods such as PRA, by showing how things can go right, and what is needed for this to happen, rather than just showing how things can go wrong. Safety is not the freedom from unacceptable risks, but the ability to succeed during expected and unexpected conditions alike. In order to bring this about, it is necessary to describe the system as a whole, and to understand how the various functions are coupled and depend on each other. Large-scale socio-technical systems, such as nuclear power plants, have become so complex that they seriously challenge established safety methods. While resilience engineering may not provide ready-made answers to the new problems, it does help us to see them more clearly, and also provides the basic principles from which such solutions can be developed.

REFERENCES

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

[2] Hollnagel, E., Paries, J., Woods, D. D. & Wreathall, J. (Eds.) *Resilience engineering in practice: A guidebook.* Farnham, UK: Ashgate (2011).

[3] Sundström, G. A. & Hollnagel, E. (Eds). *Governance and control of financial systems: A resilience engineering perspective*. Aldershot, UK: Ashgate (2011).

[4] Hollnagel, E. FRAM: *The functional resonance analysis method for modelling complex socio-technical systems*. Farnham, UK: Ashgate (2012).

[5] Lundberg, J., Rollenhagen, C. & Hollnagel, E. What-you-look-for-is-what-you-find – the consequences of underlying accident models in eight accident investigation manuals. *Safety Science, 47*(10), 1297–1311 (2009).

[6] Weick, K. E. & Sutcliffe K. M. *Managing the unexpected*. San Francisco: Jossey-Bass (2001).

[7] Hollnagel, E., Nemeth, C. P. & Dekker, S. (Eds). *Remaining sensitive to the possibility of failure*. Resilience Engineering Perspectives, 1, Ashgate, Aldershot, UK (2008)

[8] Westrum, R. A typology of Resilience Situations. In Hollnagel, E., Woods, D. D. & Leveson, N. (Eds.), *Resilience Engineering: Concepts and Precepts* (pp. 55-65). Aldershot, UK: Ashgate (2006).

[9] Merton, R. The Unanticipated Consequences of Purposive Social Action, *American Sociological Review*, 1, 894-904 (1936).

[10] Hollnagel, E. *The ETTO Principle: Efficiency-Thoroughness Trade-Off: Why things that go right sometimes go wrong*. Ashgate, Aldershot UK (2009).

[11] Dougherty, E. M. Jr. Human Reliability Analysis-where shouldst thou turn? *Reliability Engineering and System*

Safety, *29*, 283-299 (1990).

[12] Meijer, C. H., Callaghan, V. C. & Hollnagel, E. Improved man-machine system design for nuclear power plants. Proceedings of the ANS/ASME *Conference on Design. Construction Operation of Nuclear Power Plants*. Portland, Oregon. August 5-8 (1984).