

PREPRINT of paper (Please refer to the paper for changes):

Back, J., Furniss, D., Hildebrandt, M., & Blandford, A.

Resilience markers for safer systems and organisations.

SAFE COMP 2008.

27th International Conference on Computer Safety, Reliability and Security.

Resilience Markers for Safer Systems and Organisations

Jonathan Back¹, Dominic Furniss¹, Michael Hildebrandt², and Ann Blandford¹.

¹ University College London Interaction Centre
{j.back, d.furniss, a.blandford}@ucl.ac.uk

² OECD Halden Reactor Project, Industrial Psychology Division
michael.hildebrandt@hrp.no

Abstract. If computer systems are to be designed to foster resilient performance it is important to be able to identify contributors to resilience. The emerging practice of Resilience Engineering has identified that people are still a primary source of resilience, and that the design of distributed systems should provide ways of helping people and organisations to cope with complexity. Although resilience has been identified as a desired property, researchers and practitioners do not have a clear understanding of what manifestations of resilience look like. This paper discusses some examples of strategies that people can adopt that improve the resilience of a system. Critically, analysis reveals that the generation of these strategies is only possible if the system facilitates them. As an example, this paper discusses practices, such as reflection, that are known to encourage resilient behavior in people. Reflection allows systems to better prepare for oncoming demands. We show that contributors to the practice of reflection manifest themselves at different levels of abstraction: from individual strategies to practices in, for example, control room environments. The analysis of interaction at these levels enables resilient properties of a system to be ‘seen’, so that systems can be designed to explicitly support them. We then present an analysis of resilience at an organisational level within the nuclear domain. This highlights some of the challenges facing the Resilience Engineering approach and the need for using a collective language to articulate knowledge of resilient practices across domains.

Keywords: human error, distributed cognition, control rooms, nuclear domain.

1 Introduction

In this paper we analyse manifestations of resilient practice at different levels of abstraction from the individual working with simple artefacts to more complex team working situations. Resilience markers can be any system feature or procedure that enables resilient practice to manifest. Identifying these markers may provide useful performance indicators, and allow the resilient characteristics of a system to be communicated, so that existing features or procedures can be augmented in a way that increases the capacity for resilience beyond that which is already present.

Resilience markers specify the conditions that need to hold for a system to perform resiliently. In addition to enabling the detection of error-prone or non-resilient computer systems, our approach provides a means of reasoning about resilience. This allows us to look at distributed systems from a new perspective. Resilience engineering takes the view that resilience is a characteristic of a system. This implies that a holistic perspective is required to develop an understanding. We are aware that the levels of granularity presented here are interrelated and so they should be considered collectively. However, much more work is needed to integrate these different levels. Indeed it could be argued that the nature of resilience goes against a level-based composition, however, our central focus is on finding evidence for resilience in the behaviour we observe, and identifying what type of behaviour we would classify as resilient. The aim of this approach is to develop an understanding of the system attributes that encourage people to engage in resilient activities (*see* Sections 3 and 4). We also discuss the difficulties of understanding resilience issues at an organisational level by presenting a case study from the nuclear domain (*see* Section 5). The examples presented in this paper should not be considered a full set of resilient behaviours that need to be supported: they have been selected as being representative of different levels of granularity that researchers and practitioners need to consider when designing systems that foster resilient performance (*see* Table 1).

Table 1. Levels of Granularity.

Granularity	Examples of Vulnerabilities	Resilient Manifestations	Resilient Markers
Individual Level (<i>see</i> Section 3)	Errors in procedural routine	1. Reflection 2. Cue creation	Providing an opportunity for meta-cognitive activities.
Small Team Level (<i>see</i> Section 4)	Coping with increased demand	1. Buffering 2. Work shadowing 3. Artefact use	Optimised flow of information and physical layout. An understanding of artefact use, social conditions.
Operational Level (<i>see</i> Section 5)	High complexity	Error recovery	Symptom-based emergency procedures, automatic safety systems, strategic crew leadership.
Plant Level (<i>see</i> Section 5)	Plant shut downs or failures to start up, major accidents	1. Plant safety record 2. Response to major disturbances	Maintenance regime, plant upgrades, risk analysis, training programs.
Industry Level (<i>see</i> Section 5)	Political and regulatory intervention	Performance necessity and availability of alternatives	Regulatory compliance, public/political perception, cost-benefit ratio, competitiveness.

2 Background

Making a system safer involves coupling the capabilities of humans with the technology they work with so that they can stay in control. A resilient system is able to recognise, adapt to and absorb disturbances so that it remains safe by being flexible to new demands [1]. We report on work using experimental microworlds that enable cognitive strategies to be understood, as well as studies of team working situations using distributed cognition modelling. We also look at how the design of computer systems in control room environments explicitly supports resilient practice.

Historically, there has been much more focus on why things go wrong than on why they work well. Conventional engineering approaches to ensuring safety attribute failure to a system component (human or technological) rather than the system as a whole. When systems fail, the cause is often attributed to ‘human error’ or to a technical problem associated with a control process. Attributing blame to a faulty component offers a pragmatic solution; the component can simply be replaced, fixed, or retrained. The traditional view of managing safety involves attempts to reduce the complexity of a system so that humans can maintain control under stress [2]. For example, one technique is to try and design systems that minimise the number of procedures by automating subsidiary interactions and leaving only the main parameters for the operators to worry about. Ostensibly, this decreases the system complexity from a human-computer interaction (HCI) perspective. However, Perrow’s account of high-risk technologies highlights that it is not complexity per se that causes accidents [3]. The existence of many system components is not a problem for either system designers or operators if their interactions are expected. Based on the analysis of case studies and foundational empirical work, we found that dealing with unexpected or hidden events is facilitated by: designs that provide operators with an opportunity to engage in reflection [4]; expanding the variability of actions operators can take [2]; supporting the use of artifacts (such as dynamically generated checklists) that augment the capabilities of human cognition [5, 6]. These types of interactions allow a system to maintain control by anticipating new demands. We classify them as being resilient interactions.

The performance of cognitive systems, ranging from the individual to a team, has been found to be sensitive to external factors such as time constraints and workload which erode control [1, 7]. However, experts are able to generate strategies that support resilient practice (e.g., [8]). Understanding how these strategies are generated will enable the development of computer systems that explicitly support resilient activities. Our approach is about understanding how systems can support the cognitive and communicative capabilities of humans. This enables the socio-technical system as a whole to adapt to oncoming demands. Work suggests that the process of managing demands is influenced by task structures and team roles [9], external cognitive artefacts and computer system design [10]. We suggest that these factors shape the potential for resilient interactions rather than simply attributing resilience to the capabilities of individuals themselves.

An opportunity to think about oncoming demands is essential for individuals, teams, and organisations to reason about ways that performance can be better supported, enabling future strategies to be formulated. For example, an opportunity to reflect can enable an individual to offload workload, allowing them to maintain levels of performance under stress or high load situations. For example when anticipating being in a rush to leave home for work, positioning your bag by the door reduces the likelihood of forgetting to take it with you. Reflecting in a team setting can allow for interruption management [11], task collaboration and temporal coordination [12]. Foundational work suggests reflection at an organisational level is unlikely to take place during routine operation. Nathanael and Marmas’s Repetitions-Distinctions-Descriptions model [13] suggests that encountering abnormal or different scenarios forces ‘distinctions’ from the normal routine to be made. These ‘distinctions’ trigger reflection-in-action to alter practice; this altered practice can then be absorbed back

into normal routine if appropriate. The ability of a socio-technical system, in which computer systems are an integral part, to prepare for oncoming demands is an important aspect of resilience. However, it is by no means the only one. Other aspects are discussed in Hollnagel and Woods [14].

3 Cognitive Resilience at the Individual Level

The first level of granularity to be considered is cognitive resilience. In safety-critical domains operators frequently perform routine tasks. Research on procedural routine has demonstrated that under increased workload individuals are more prone to slips [15]. Although the consequences of a slip do not necessarily move a system towards failure, the ability of an operator to perform effectively is influenced, since some control over the processes they are trying to manage has been lost. While most day-to-day slips result in minor annoyances, those that occur in safety-critical situations (such as in the aviation domain) can be catastrophic. Slip errors can occur systematically even when individuals have the required 'expert' procedural knowledge to perform a task correctly. Manifestations include omission errors (e.g. forgetting to collect the original document after making photocopies), and mode errors (e.g. typing with the Caps Lock mode activated). Slips cannot be eliminated through practice or increased motivation [16] but they can be reduced by adopting a resilient strategy (such as leaving your bag by the door). We hypothesised that reflection can support performance during HCI, allowing slip errors to be mitigated. To test this hypothesis, an understanding of under what conditions individuals are able to engage in reflection was needed. In order to address the question of how an individual's resilient cognitive activities emerge, a 'Fire Engine Dispatch Centre' microworld was developed [6]. The development of a microworld to study how individuals avoid slips improves understanding of what factors shape performance.

The overall objective of the microworld experiment was to send navigational information to fire engines enabling the fastest possible incident response times. When a call was processed the location of the nearest fire engine and the location of the incident were displayed automatically as waypoints on a map. Participants had three minutes to identify the best route based on information displayed on a traffic information ticker. Training trials were used to ensure that participants became familiar with the sequence of actions. After performing two 'error free' training trials consecutively, a participant was allowed to move on to twelve experimental trials. Two error-prone task steps, outlined below, were built into the design: an initialisation step and a mode selection step. The emergence of resilient strategies associated with these steps provides concrete examples of cognitive resilience.

Initialization Step. When commencing a new trial an individual had to decide which call to prioritize before clicking on the 'Start next call' button (*see* Figure 1).

For each trial there was only one correct call prioritization selection. Participants were trained to know that incidents in poor fire engine coverage areas should be selected before incidents in good coverage areas. They also knew that high priority calls took precedence over normal priority calls irrespective of fire engine coverage.

The first step in the process of setting call priority involved clicking on the radio button that was located alongside the required call ID. For example, in Figure 1 a participant is required to select ID 4. Clicking on ‘Confirm priority change’ is the second procedural step. Participants were instructed that the ‘Start next call’ button should only be clicked when both the new call ID has been selected and the ‘Confirm priority change’ button has been clicked.

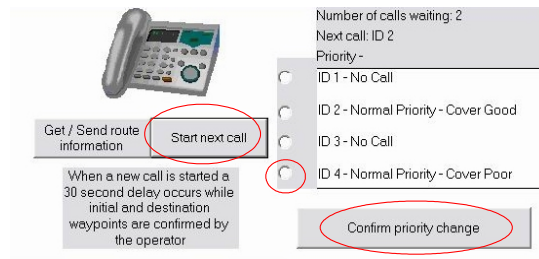


Fig. 1.Initialisation Step

When a routine task is learned task, steps become associatively linked, i.e. action x (e.g. inserting a DVD) becomes a procedural cue for action y (e.g. locating the remote control). The initialization step could not be procedurally cued, since there was no preceding step, making it highly error-prone. The error occurred when participants omitted the initialisation step, which involved prioritizing calls to the dispatch centre, and instead clicked on the start next call button. The start button captured attention away from the correct procedure since it moved a participant towards starting the primary task of routing fire engines. Experimentation revealed that initialization errors were more avoidable if participants were given the opportunity to reflect on task requirements. The number of initialisation errors made by participants in Condition A, where the system encouraged reflection by displaying the control interface during a trial resumption delay, was compared with the number of errors made by participants in Conditions B, where participants were presented with a blank screen. The mean error rate when display cues were present was 6.09% compared to 23.12% when cues were absent (*Mann-Whitney U* = 40.2, *Wilcoxon W* = 158.5, *Z* = -2.605, $p < .01$, across 24 participants). Providing users with an opportunity to rehearse procedural steps allows for reflection. System designers can modify the task environment to ensure that rehearsal is possible and in some cases, where problematic interactions have been identified in the past, is actively encouraged (by enforcing delays). Providing a window-of-opportunity as a means of facilitating reflection is a useful marker for resilient design.

Mode Selection Step. After identifying a route, a participant had to select the required route construction mode.

When a participant commenced the route construction procedure (after clicking on the start button) the first requirement was to identify the most appropriate route on the map. Participants had to select the best route based on traffic information (i.e. they had to ensure a proposed route did not run through an accident or heavy traffic area).

The device provided a signal that informed participants of the required method of route construction (located above the telephone image, *see* Figure 2). This signal was available after 35-45 seconds from pressing the start button. Participants were required to attend to this signal so that they could determine what type of route information was needed. If GPS was available then the centrally located menu could be used. Clicking on this menu enabled one of the automatically generated routes to be selected. The drop-down menu located below and to the left of the automatic route selection menu was used for manual route construction. A mode error occurred when a participant used the wrong route construction method.

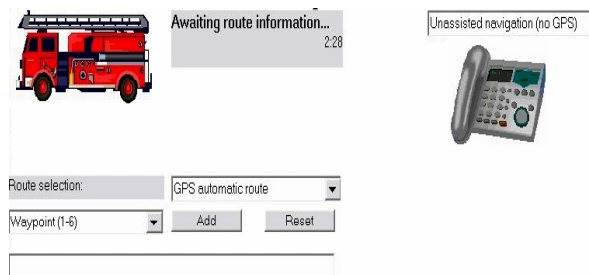


Fig. 2. Mode Selection Step

Attending to the mode selection step indicator required an attentional shift away from the main problem-solving task, making it highly error prone. A post-hoc analysis revealed the generation of a tractable resilient cognitive strategy. If participants placed the mouse cursor close to the signal status display (above the telephone in Figure 2) before the signal status appeared, they were less likely to forget to attend to the display before selecting the appropriate route construction method. When the mouse cursor was placed $< 2\text{cm}$ from the display, participants were significantly less likely to make a mode error (Wilcoxon $Z = -1.870$, $p < .05$, two related samples test, across forty-eight participants). Positioning the mouse cursor enables the creation of a sensory cue. If the cursor is attended to then it may indicate that the display should be attended to when route identification is complete.

Further experimentation revealed that the generation of this resilient strategy was significantly more likely under a mixed workload condition. The complexity of the routing task was manipulated so that half of the participants only performed difficult routing tasks while the other half performed both easy and difficult tasks (mixed workload). In the mixed workload condition 64% of participants adopted the cursor strategy. In the high workload only condition only 27% of participants used the mouse cursor as a candidate cue. Critically, participants in the mixed workload condition who adopted this strategy were able to apply it during easy and difficult tasks (Wilcoxon signed rank test, related samples, *Wilcoxon* $Z = -1.039$, $p < .05$). Analysis of these findings enables us to identify a further marker for resilience that has implications for the design of computer systems. Personalised cue creation is spontaneous and can be used to minimize the likelihood of error. Allowing users to position markers (like 'Post-it' notes) provides support for attentional control.

However, the use of such cues is only likely in situations where distinctions to the normal routine can be made. Mixed workload participants had: the cognitive resources available to think of an appropriate cue to guide attention (when workload was low) and the motivation for doing so, i.e. to support performance during high workload trials. Systems designers need to design scenarios that encourage metacognition during routine performance. It is generally agreed that the metacognitive activity consists of two basic processes occurring simultaneously: monitoring progress, and selecting or generating strategies to support performance [17]. Individuals need to be encouraged by the system to engage in metacognition so that they can develop a repertoire of resilient strategies. Reflection encourages the development of appropriate strategies and so enables levels of performance to be maintained under stress.

4 Resilience at the Small Team Level

As illustrated in the previous section, markers for resilient performance can be ‘seen’ in the laboratory. However, as previously discussed, manifestations of resilient practice occur at different levels of abstraction; next, we consider more complex team working situations. There are many different things to ‘see’ in socio-technical contexts, often too many, and so it is helpful to have approaches that can facilitate our perception in the ‘noise’ of real world contexts. DiCoT (Distributed Cognition in Teamwork) has been developed as an approach to applying distributed cognition to teamwork contexts [18]. Distributed cognition is a theoretical area which maintains the computational vocabulary associated with cognitive psychology but expands its unit of analysis. Hollan et al. [19] suggests three ways in which this expansion occurs:

- “Cognitive processes may be distributed across members of a social group”;
- “Cognitive processes may involve coordination between internal and external (material or environmental) structure”;
- “Processes may be distributed through time in such a way that the products of earlier events can transform the nature of later events.”

This expansion has important implications for reflecting on and preparing for oncoming demands. For example: What system are we considering to be receiving these demands e.g. an individual, a team, a department, a company? Who is passing on the information and how? What timeframe and what sort of demands are we talking about e.g. restructuring the company over years or preparing for the next five minutes? How is this information structured internally within individuals? How is it represented externally in procedures? DiCoT encourages a system description which helps engage with these issues. Hollan et al. [19] indicate that what functionally influences the computation of the system is the concern of DC. DiCoT encourages analysts to look at these functional influences through five interdependent models. These look at the structure of information flows in the system, the artefacts which are used, the physical layout of the system, the social structures and factors in the system, and how the system has changed over time. These models, and the way they can be used to reflect on oncoming demands, are introduced below with reference to a London Ambulance Service control room study.

Information flow model. The information flow model concerns itself with the propagation and transformation of information within the system. This model underlies the other models. Firstly, the overall computational function of the system is represented in an input-process-output diagram. For example, the input-process-output diagram of an ambulance dispatch system is shown in Figure 3. After this the make-up of the computational system can be explored. Figure 4 shows the abstract computational structure of an ambulance dispatch team. From this we notice that the structure of the system is designed to cope with the oncoming demands of the system. The raw material from the External Callers is filtered into critical information for the decision hub. The buffers control information to the decision hub considering its workload and the criticality of the information. The filter does not hold up information in this way: it just changes its form for computational purposes. If the flow of information around a system is designed in a way that enables critical performance to be maintained during variability in workload, this can be considered a marker for resilience.




Fig. 3. The input-process-output diagram of an ambulance dispatch system.




Fig. 4. Overview of main information flow properties the ambulance dispatch system. The Call Taker (C) filters the raw information from the External Caller (Ex C). This structured information is passed to the Allocator (A) who decides which ambulance should attend based on priority, availability and location. Depending on the status of the ambulance the Allocator (A) will channel information to the Telephone Dispatcher (T) or the Radio Operator (R), who will contact an ambulance crew at a station (Crew St) or one which is mobile (Crew Mob). Feedback from the ambulance crews (Crew St and Crew Mob) goes back through the Telephone Dispatcher (T) and the Radio Operator (R) who act as buffers for the decision hub i.e. holding up information when the hub is busy, if it is non-critical and would be disruptive.

Physical Model. The physical model concerns itself with functional influence of the physical layout of the system. For example, at the time of the study, the ambulance dispatch control room in London had seven desks, each of which is responsible for allocating ambulances to a different area of London. The arrangement of the seven desks reflects their geographical location, as adjacent areas will sometimes collaborate on the shared use of resources and attending incidents. This is particularly important with incidents near their shared border. This layout facilitates the oncoming demand of cross-boundary collaboration.

Figure 5 shows the seating arrangement of one of the allocating desks. The Allocator and Radio Operator work closely together, and so are adjacent. This facilitates their collaboration as the Radio Operator is implicitly aware of the Allocator's activities by shadowing them i.e. listening to their communication with others and watching their monitors. This allows the Radio Operator to prepare for oncoming activities before their receipt. This augmented awareness of work demands, through physical co-location, can be considered as a marker for resilience.




Fig. 5. An allocating desk. The information flow concerned with this sector desk is described in Figure 2.

Artefact Model. The artefact model concerns itself with the influence of the use of artefacts in the system.

Two brief examples of preparing for oncoming demand include: that the sector desks use a computer and card system which prepares them for the eventuality that the computer system might fail; and the computer system will manipulate the colour of incidents to indicate their criticality, facilitating the Allocator's prioritisation of incidents. Also, as soon as the Call Taker has established the location of the incident,

the Allocator will have access to the updating details so that they can prepare for the oncoming demand. Redundancy and the support of decision making are important resilience markers.

Social Model. The social model concerns itself with the functional influence of the social structure and factors within the system. An example of inbuilt resilience at this level is that people generally get promoted from Call Taker, to Telephone Dispatcher, to Radio Operator to Allocator: so the more responsibility they have, the more aware they are of the other functions in the system and the way they work. An example of such resilience is that the Allocator may contact the ambulance directly if the Telephone Dispatcher is busy. Effective knowledge and responsibility transfer is a marker for resilience.

Evolutionary Model. The evolutionary model concerns itself with how the computational structure and functions of the system have changed over time. An example of a major change in the ambulance dispatch scenario was the introduction of GPS mapping. This gives Allocators a dynamic visual display of where the incident is and where their ambulances are located. These changes typically happen as a result of a constant pressure to improve processes, to respond to increasing demand from the environment, and to respond to the potentials new technology can offer. Exploiting technological advances to better cope with demands from the environment is a marker for resilience.

DiCoT can be used to understand the computation of the socio-technical system within these five interdependent models. This analysis notices how the system is coordinated to cope with oncoming demands.

5 Identifying Resilience in the Nuclear Domain

Two factors make the nuclear industry a particularly interesting context in which to discuss resilience. First, in a high-revenue, high-consequence socio-technical system such as the nuclear industry, significant safety and productivity gains can be expected if the promises of the Resilience Engineering approach can be delivered. Second, the nuclear domain presents an ideal environment for developing, operationalising and testing models of resilience. One of the reasons for this is that analysis of events in the nuclear industry requires a systemic approach. It is virtually impossible to discuss issues at one level of abstraction (operational, plant, industry and regulatory) without recourse to other levels. The degree of interconnectedness becomes clear when elaborating some of the defining characteristics of nuclear operations: information-rich operational environment; stable operations; possibility for severe disturbances; highly trained crews of operators; operational support network; highly proceduralised emergency operations reflecting thorough analysis of design-base accident scenarios; possibility of beyond-design-base incidents (e.g. fire); tight regulatory oversight and reporting regime; high investments and operations cost; high revenue; a variety of stakeholders, including operators, utilities, vendors, politicians and the public.

This section aims to identify manifestations of resilience at different organisational levels in the nuclear industry. The analysis is based on a number of information sources, including results from full-scale simulator experiments and training; incident and event reports; observational, ethnographic and interview studies (e.g. [20]); as well as Performance Shaping Factors that have been found to affect mission success over a range of scenarios in the context of Human Reliability Assessment.

Operational level. Nuclear operations are characterised by a high level of proceduralisation (especially during emergencies), and by a set of automatic safety functions designed to prevent the most severe consequences of accidents (core damage, release of radiation). There is crew-to-crew variability in procedure adherence, but crews are expected to follow the procedures as closely as possible. This system of operators, procedures, control room equipment and automation is expected to perform reliably for design-base incidents, i.e. those scenarios that have been considered during system design and in Probabilistic Risk Assessment. It is the successful interaction between these system components that creates resilience for design-base scenarios. Beyond these systemic properties, a number of factors have been recognized to improve the ability of the system to respond to disturbances. For instance, the move from event-based to symptom-based emergency procedures has allowed a wider range of plant states to be addressed, and provides operators with a simpler and more unified way of responding to complex events [21].

Even with these well-designed and well-tested procedures, plant conditions can arise that challenge the procedures and require knowledge-based situation assessment [22]. To respond successfully to these unanticipated, beyond-design-base events, both instrumentation and crew responses play an important role. Instrumentation helps the crew maintain an overview of the situation and develop an appropriate response plan. Other industries (e.g. petroleum) have already gone further down this path, and the nuclear industry can benefit from developments such as large-screen and information-rich displays, trend displays and ecological interface design. When considering crew responses to beyond-design-base events, a number of characteristics for success have been identified in recent simulator studies [22], including shift supervisors' team leadership style and situation assessment. This suggests that success in nuclear control tasks at a mission level *may* not depend only on success or failure of low-level activities, such as slips, lapses or misidentifications. Given the operational context and time available, such erroneous actions *should* be recovered from without significantly affecting the overall mission. Instead it appears to be crew-level factors, work styles and orientations that are more likely to determine mission-level success or failure. Differences between domains in the significance of low-level failures may be accounted for by the role of time. In domains with acute time pressure such as aviation, it is more likely that low-level erroneous actions can have catastrophic consequences, whereas in the timeframe available to nuclear operators, recovery mechanisms are in place that can compensate for low-level failures. Therefore, available time, and the situational and systemic factors that compensate for failures of individual system components, can be considered resilience markers. Investigating differences between domains as to how these factors influence mission success may provide important insights into markers for resilience.

Plant level. Plants react to outside influences (safety requirements, economics, public opinion) through upgrade programs, training, perseverance, or closure. Several candidates for markers of resilience at this level are available, including performance measures, safety measures (incidents / accidents), and safety culture measures. If and how these indicators measure resilience, in the sense of the plant's ability to respond to and recover from major disturbances, and to adapt to long-term outside forces, is unclear. Analysis of cases where plants have been built but never started up, were shut down well before the end of the designed life cycle, or consistently produce below-expectation power outputs may significantly improve our understanding of resilience. Case studies suggest that the management of organisational change plays an important role, and may constitute a marker for resilience. Organisational factors include conflicts between professional groups within a plant (e.g. operations, maintenance, engineering, managerial), problems of staff recruitment and retention (especially with regards to an aging work force in a so-called 'sunset industry'), and the effects of organisational re-structuring (e.g. mergers, change of ownership). Each of these factors can generate disturbances that compromise the resilience of the plant. A better understanding of these factors is needed as plants prepare for upgrades that will see their lifetimes extend for several decades.

A critical factor for resilience at both the operational and plant level, and a potential marker for resilience, is training. While regular training on well-known initiating events (e.g. steam generator tube rupture) improves response reliability on design-base scenario, training for beyond-design-base operations may require different approaches. More recently, training programs have started to place emphasis on scenarios that challenge procedure support, require knowledge-based diagnosis and planning, involve close crew interactions and communication, and are specifically designed to promote the shift supervisor's situation assessment. Debriefing of simulator training runs is moving from an instructional, failure-based approach towards a crew-guided, reflection-oriented approach.

Industry level. Many of the themes discussed in the previous section re-emerge when considering resilience at the industry level. Judging by the outcome, the nuclear industry possesses remarkable resilience. It recovered from severe accidents and the resulting hostile public opinion. While the survival of the industry was predicated on the organisational changes and safety improvements that followed in the wake of these events, the need for power output and lack of alternatives also played an important role. This suggests that resilience refers not only to the internal quality of a system to adapt to changes in its environment. Instead the environment itself (in this case: politics, the public) is in turn shaped by the perceived value of the products and services provided by the system. From this point of view, resilience markers at an industry level include pricing, demand and competition as well as safety records. Even the sheer size of the industry and the investments made in the infrastructure may contribute to its continued survival (resilience by inertia).

Finally, an important aspect of resilience in the nuclear industry is the role of the regulator. Many aspects of nuclear operations are subject to regulatory oversight. Regulatory practices such as risk-informed decision making have made safety assessment of highly complex systems feasible, while leaving plants some degree of

flexibility in implementing and managing their own safety programs. The effect of regulatory oversight on the ability of the industry to adapt and change, the model of performance variability embedded in regulatory practices, and the analysis of outside forces affecting the regulators themselves, are important fields for resilience research.

6 Discussion and Conclusions

The examples presented in this paper are representative of different levels of granularity that researchers and practitioners need to consider when designing computer systems that foster resilient performance. All these examples demonstrate that people are an important source of resilience in creating safety under performance pressure. Our findings are incompatible with the view that erratic people degrade an otherwise safe system, and align with the viewpoint of Cook and Woods [23], who argue that humans need to be supported in a way that helps them cope with complexity. As Rochlin [24] identified, when managing hazardous technical operations, a high level of performance does not flow from eliminating error but rather through anticipating and planning for events and surprises.

At the cognitive level (*see* Section 3) we demonstrated how computer systems can be designed to enable individuals to develop resilient strategies. By allowing individuals to reflect on task requirements, the generation of these strategies becomes spontaneous. The spontaneity of using artefacts in the environment (such as a mouse cursor) to support performance when task demands are increased results in resilient human performance. At the small team level (*see* Section 4) the use of a methodological approach such as DiCoT is able to reveal the hidden complexity of team interactions. DiCoT provides potential to be used as a tool to analyze the performance of the system and recommend improvement in processes, in layout, in technologies, and in social structures within a system's history of change. Being able to represent interactions at a team level is important for understanding resilience, as manifestations, such as the ability to buffer, need to be supported by the way a control room is designed. Computer systems play an increasingly influential role in control rooms so should be considered as an integral component during design. At operational, plant, and industry levels (*see* Section 5) manifestations of resilience are harder to observe. However, the examples presented illustrate that people are still an essential source of resilience, and that the design of complex distributed systems should provide ways of helping people cope with complexity. Computer systems need to support: symptom-based diagnosis of problems at the operational level; flexibility and extendibility at the plant level; and survivability at the industry level.

Resilience markers can aid analyses of simulated scenarios at the individual and team levels, which can be used to evaluate the performance of safety-critical systems. Resilience markers at operational, plant, and industry levels can be used retrospectively. However using markers to predict performance and survivability requires researchers and practitioners to consider the interrelations between all levels collectively. More work needs to be done on understanding their integration.

Acknowledgement. Back, Furniss, and Blandford were supported by EPSRC grant GR/S37494.

References

1. Hollnagel, E., Woods, D.D.: Joint cognitive systems: Foundations of cognitive systems engineering. Taylor & Francis, Boca Raton, FL (2005)
2. Dekker, S.: Failure to adapt or adaptations that fail: contrasting models on procedures and safety. *Applied Ergonomics* 34(3), 233-238 (2003)
3. Perrow, C.: *Normal Accidents: Living with High-Risk Technologies*. Basic Books (1999)
4. Back, J., Furniss, D., Blandford, A.: Cognitive Resilience: Reflection-in-action and on-action. In: *Proc. Resilience Workshop*, pp. 1-6. Linköping University (2007)
5. Masino, G., Zamarian, M.: Information technology artefacts as structuring devices in organizations. *Interacting with Computers* 15(5), 693-707 (2003)
6. Back, J., Blandford, A., Furniss, D., Curzon, P.: *Avoiding Slips*. Submitted for journal publication (2008)
7. Wright, P.: The harassed decision maker: Time pressures, distractions, and the use of evidence. *Journal of Applied Psychology* 59, 555-561 (1974)
8. Klein, G., Orasanu, J., Calderwood, R., Zsombok, C.E.: *Decision Making in Action: Models and Methods*. Ablex Publishing Co., Norwood, NJ (1993)
9. Kirsh, D.: Adapting the environment instead of oneself. *Adaptive Behaviour* 4 (3/4), 415-452 (1996)
10. Spillers, F., Loewus-Deitch, D.: Temporal attributes of shared artifacts in collaborative task environments. In *Proc: HCI 2003 workshop on temporal aspects of tasks* (2003)
11. Furniss, D., Blandford, A.: Understanding Emergency Medical Dispatch in terms of Distributed Cognition: a case study. *Ergonomics Journal* 49 (12/13), 1174-1203 (2006)
12. Bardram, J.E.: Temporal coordination: On time and coordination of collaborative activities at a surgical department. *Computer Supported Cooperated Work* 9, 157-187 (2000)
13. Nathanael, D., Marmas, N.: The interplay between work practices and prescription: a key issue for organisational resilience. In: *Proc. 2nd Resilience Eng. Symp.*, 229-237 (2006)
14. Hollnagel, E. & Woods, D.D.: Epilogue: Resilience engineering precepts. In: E. Hollnagel, D.D. Woods, N. Leveson (eds.). *Resilience engineering: Concepts and precepts*. Ashgate, 347-358 (2006)
15. Byrne, M.D., Bovair, S.: A working memory model of a common procedural error. *Cognitive Science* 21, 31-61 (1997)
16. Back, J., Cheng, W.L., Dann, R., Curzon, P., Blandford, A.: Does being motivated to avoid procedural errors influence their systematicity? *Proc. HCI 2006*, pp.151-157 (2006)
17. Ertmer, P.A., Newby, T.J.: The expert learner: Strategic, self-regulated, and reflective., *Instructional Science* 24, 1-24 (1996)
18. Blandford, A., Furniss, D.: DiCoT: a methodology for applying Distributed Cognition to the team working systems. *DSVIS 2005, LNCS*, vol. 3941, pp 26-38. Springer (2005).
19. Hollan, J., Hutchins, E., Kirsh, D: Distributed cognition: toward a new foundation for human-computer interaction. *ACM Trans. Comput.-Hum. Interact.* 7(2), 174-196 (2000)
20. Perin, C.: *Shouldering Risks*. Princeton University Press (2004)
21. Ujita, H., Kubota, R., Ikeda, K.: Development and Verification of a Plant Navigation System. *Cognition, Technology & Work* 3, 22-32 (2001)
22. Halden Work Report 844. The International HRA empirical study – Pilot phase report. OECD Halden Reactor Project. Halden, Norway (2008)
23. Cook, R.I, Woods, D.D.: Operating at the Sharp End: The Complexity of Human Error In: Bogner MS (ed.). *Human Error in Medicine*. Lawrence Erlbaum, 255-310 (1994)
24. Rochlin, G.: Safe operation as a social construct. *Ergonomics* 42, 1549-1560 (1999)