34

Design and Engineering of Resilience for Networked Computer Systems

David Hutchison, Mark Rouncefield, Antonios Gouglidis, and Tom Anderson

Introduction

Networked computer systems that are designed for resilience form the bedrock of many enterprises and activities in the modern world, from telecommunications (telephone, broadband) through utility networks (electricity, water, gas, etc.) to banking, commerce, government, and all sorts of organizations including those in areas of healthcare and transportation. These systems are composed of nodes (computers) and links (communication paths, which may be wired or wireless), interconnected in some topology or arrangement of links (e.g., mesh, star, or tree). It is convenient to represent networked systems as a number of services running on top of the communications topology (see Figure 34.1). Each is a combination of software and hardware. One reason for this representation is that it allows designers to separate the concerns of the communication topology from those of the services.

Modern networked systems definitely need to be reliable and trustworthy. In other words, the operators and, ultimately, the users need to know that the service they receive will be what they expect and also what they have paid for. Put simply, networked systems need to show resilience when strained. The subject of QoS has been a highly active research topic for many years and is still perhaps the most important aspect of any system because the service the user receives is its essential purpose.

In recent years, it has become evident that modern networked systems are critical infrastructures (and services), because of the reliance that users put on them. Not only that, if some of these systems fail to provide their expected service (perhaps a prolonged downtime),

David Hutchison, Mark Rouncefield, Antonios Gouglidis, and Tom Anderson, *Design and Engineering of Resilience for Networked Computer Systems* In: *Multisystemic Resilience*. Edited by: Michael Ungar, Oxford University Press (2021).
DOI: 10.1093/oso/9780190095888.003.0035

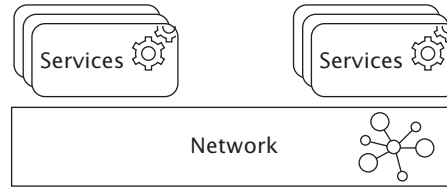

FIGURE 34.1 Networked system: topology and services.

then losses will occur in terms of time and money, and in extreme cases there may be damage and even loss of life. Critical infrastructures comprise of assets and systems that maintain societal functions, including health, safety, security, and the economic and social well-being of people. Supervisory control and data acquisition (SCADA) and industrial control systems (ICS) are particular examples of critical infrastructures for the monitoring, control, and automation of operational plants of various sorts, such as utility networks. SCADA systems monitor and control infrastructures including power plants, water utility, energy networks, and gas pipelines, which makes them highly critical. Providing protection in terms of security, safety, and resilience in such networks is inherently considered to be of vital importance. Traditionally, most of these systems were air-gapped (physically isolated) from other unsecured networks, but in several cases, access to these devices may still be available over a public network (e.g., the Internet) as a requirement to improve usability via providing operators with the potential to remotely access devices (Shirazi, Gouglidis, Farshad, & Hutchison, 2017).

While automation and interconnectivity increase the efficiency of these computer systems and reduce operational costs, they expose these systems to new threats. For instance, the existence of a vulnerability in a system on the top layers of the Purdue model, a way of modeling multiple layers and stages of the architectural life cycle (Obregon, 2015), may allow attackers to exploit them and to gradually take control of systems or devices that operate at the lower levels, such as SCADA systems; this could cause failure and hence serious disruptions. In recent times, there has been a significant increase in the functional demands upon utilities, for example, resulting in an increased rate of automation in networked controls and interconnections, as well as an increase in dependencies between diverse infrastructures. Consequently, utility networks are now more susceptible to sophisticated attacks including advanced persistent threats (König, Gouglidis, Green, & Solar, 2018). Additionally, new challenges arising from system complexity, overloading, unanticipated human behavior, and vulnerabilities from third-party sources must also be considered. Needless to say, providing protection in terms of security, safety, and resilience in such networks is vitally important. Research on the emerging area of security in critical infrastructures has resulted in rules, legislation, and good-practice guidelines that we will outline later in this chapter.

The sources of challenges for networked systems can include natural disasters such as flooding, weather events leading to failure of electrical power, overdemand for the services of the system, software bugs and consequent failures, hardware component faults, complexity leading to errors by a human operator, and cybersecurity attacks (Esposito et al., 2018;

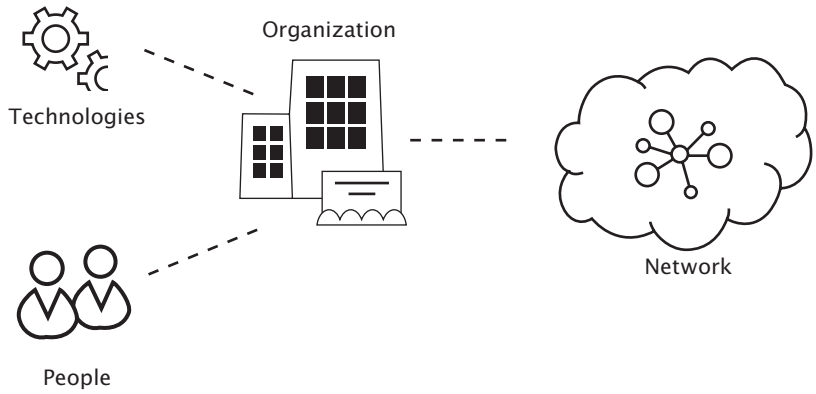

FIGURE 34.2 Technology, organization, and people in networked systems.

Machuca et al., 2016). Networked systems need to be able to continue to offer a satisfactory QoS no matter what challenge they experience—this is our definition of *resilience*. In this chapter we explain our approach to engineering resilience into such systems (Hutchison & Sterbenz, 2018).

Networked systems are generally complex, and they have three aspects that need to be considered in combination when building resilience into them: these are technology, organizations, and people, as illustrated in Figure 34.2.

We start by looking at the technology aspect, which is where we started in our own research. In later sections we consider organizations and people, by means of a case study based on work we did with utility networks. Originally, our work on resilience was in the context of future telecommunication systems, and we wanted to explore the extension of traditional QoS concerns (performance—throughput and delay in particular) to make sure these systems could be relied on, not only at the level of recovering from the failure of a node or link but also at the services level.

Our early research (Sterbenz et al., 2010) reviewed the related terminology (including fault, error, failure, fault tolerance, trustworthiness, etc.) and we described in some detail the relationship of our definition of resilience with prior and related work; this is often understood differently in disciplines other than our own area of information and communication technologies (ICT).

Resilience and Related Terminology for Engineered Systems

The term *resilience* has been used in the past several decades in different ways to describe the ability of materials, engineered artefacts, ecosystems, communities, and other built and biological systems to adapt to changes and is also adopted by diverse sciences (e.g., in the discipline of psychology) and organizations (e.g., as a description of business continuity lifecycles; Hollnagel, Woods, & Leveson, 2006). Although the etymology of resilience clearly refers to the capacity to recover from difficulties, a single agreed, precise,

definition is currently elusive. This is mostly because of the complexity and diversity of contemporary sociotechnical systems, which eventually resulted in the many definitions of resilience. For instance, resilience engineering views resilience as an alternative or complement to the safety of systems (Hollnagel, Paries, Woods, & Wreathall, 2010); resilience may also be defined as the capability of a system to self-organize, learn, and adapt (Adger, Hughes, Folke, Carpenter, & Rockström, 2005); another definition describes resilience as the capability of a system to maintain its functions and structure in the presence of changes and to degrade when it must (Allenby & Fink, 2005). The lack of a standard definition for resilience implies the absence of agreed measures of resilience (Moteff, 2012).

For engineered systems, there is a debate about the validity of different opinions in the communities interested in quantifying resilience. In the context of networked computer systems, which arguably forms the basis of an increasing number of critical infrastructures, we define resilience as "the ability of a network or system to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation" (Sterbenz et al., 2010). The overall resilience strategy, which we have labeled as D^2R^2+DR, is depicted in Figure 34.3. This definition resulted from research conducted in ResumeNet (Bruncak et al., 2011), a Seventh Framework Programme European Union–funded Future Internet project and was subsequently adopted by the European Union Agency for Network and Information Security (ENISA; Górniak et al., 2011). Based on the previous references, it is clear that there exists no single, agreed definition of resilience, and current definitions rely on the specific area of application. However, there is clearly a common thread in all of the definitions. We propose to use the above, broad, ENISA definition, as it is sufficiently general and encompasses the elements that apply to the resilience of critical infrastructures.

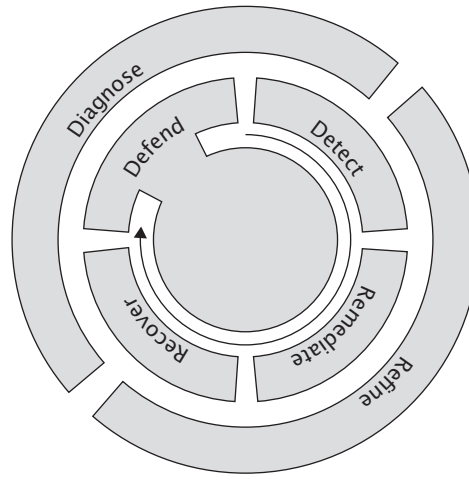

FIGURE 34.3 The D^2R^2+DR resilience strategy.

Engineering Resilience Using the Resilience Strategy

To design and build or adapt networked systems to be resilient, we adopt the D^2R^2+DR strategy, which is essentially two sets of steps organized in two "loops" as shown in Figure 34.3. The inner loop, D^2R, is intended to operate in real time (or as fast as possible) to detect and correct anomalies, whereas the outer loop, DR, can act more sedately (initially offline, mediated by a human expert, but ideally in the future it will function autonomously with the help of a machine expert; Smith et al., 2011). Each component of the model requires its own explanation if one is to grasp the complexity of their interactions.

Defend

Initially, a thorough system analysis needs to be carried out to decide how best to build defensively against perceived threats and vulnerabilities; this includes a risk assessment to prioritize the assets in the system—which of them needs to be protected and which of these, most urgently. Building resilience into a system inevitably incurs costs, and these need to be carefully weighed. As a result of the system analysis, the system designer will propose a range of actions including: building defensive walls (e.g., firewalls to defend against cyberattacks); adding some redundant links and nodes into the communications infrastructure; and at runtime, making appropriate adjustments such as firewall rules and resources.

Detect

The detect phase requires a monitoring system. Essentially, the network and/or networked system needs to be "instrumented" so that the effects or symptoms of any challenge to the system's normal operation can be rapidly observed. This is sometimes called "anomaly detection" or "intrusion detection," and it has been the subject of much research in past decades (Chandola, Banerjee, & Kumar, 2009). Nevertheless, it is difficult to distinguish the root cause of a challenge, and the detection may have to proceed without actually knowing for sure what is causing the problem. Typically, detected anomalies are classified and using this classification allows the next phase to be carried out.

Instrumenting the system implies knowing what (and where) to measure some artefact of the system that will indicate there is a threat. In a network it is usual to measure network traffic (i.e., the packets of information that are passing across it) to assess whether some variation indicates abnormal behavior. What is measured is often referred to as a "metric"; deciding which metrics to observe to estimate the resilience of a system remains an important topic of research.

Remediate

Remediation (or "mitigation" as used by some resilience researchers; Sedgewick, 2014) is the phase whereby some action is carried out to remove or improve the symptoms of a challenge or threat. In networked systems, it is typical to use traffic engineering to improve the situation—for example, to remove or redirect a particular stream of information packets that come from a suspicious source in the network and that is adversely affecting a destination

in the network such as a server that may be saturated with this traffic. Ideally, remediation should be done in real time, and it should be done autonomously—that is, the resilience management mechanism makes the decision what to remediate and how, and carries this out without human intervention. This is still a sensitive topic, and in existing systems the remediation will usually be carried out under the supervision of a human expert.

To make autonomous operation more feasible and trustworthy, it is important to get as much context as possible about the source and nature of the anomaly or challenge. Given that root cause analysis is likely to take too much time, a situational awareness (SA) subsystem could be employed to gather and assess contextual data about the environment or conditions surrounding the networked system. This can potentially provide enough information to assist the appropriate remediation decision to be made. For example, context data may be able to tell whether a web server is being saturated because of some malicious activity or, by contrast, if it is a national holiday or there is a surge of bookings for a new event and therefore not a denial of service cyberattack. SA is still a key research topic.

Recover

In the recovery phase, the aim is to return the networked system to normal behavior if possible, and to try to make sure that the system takes account of the conditions that caused the anomalies. This implies some form of machine or human learning to improve the system's resilience. The recover activity should, of course, be carried out once the source of the challenge has been removed. Policies for high-level guidance may be used in this phase (Gouglidis, Hu, Busby, & Hutchison, 2017).

Diagnose and Refine

The outer loop of the resilience strategy is an underexplored research area. The idea is that in future there will be a machine learning phase that steadily learns from previous experiences and builds up a body of expert knowledge on which to draw to improve the remediation and recovery activities and the resilience model that underlies them both. This requires providing real historical data for a DR prototype and, in turn, the development of resilience subsystems that are subsequently deployed in the field. This raises an important ethical question—whether, for networked systems that operate critical infrastructures and services, there will or should always be a human in the loop.

System Risk

Risk is defined by ENISA as "the chance of something happening that will have an impact upon objectives. It is measured in terms of impact and likelihood." (ENISA Glossary, 2019). Therefore, a cyber risk can be conceived as a risk in the context of ICS and/or ICT systems. In addition, an operational cyber security risk can be defined as "operational risks to information and technology assets that have consequences affecting the confidentiality, availability, or integrity of information or information systems" (Cebula & Young, 2010, p. 1). They classify the taxonomy of operational cyber security risks into four main groups: (a) *actions of people* is considered with actions taken or not taken by individuals in a given situation; (b) *systems and*

technology failures refers to technology assets and specifically in their problematic, abnormal or unexpected functioning; (c) *failed internal processes* refers to needed or expected performance of internal processes and associations with problematic failures; and (d) *external events* refers to external events that might affect an organization's control. Therefore, to consider how a system's risk is affected when various type of changes apply to a system, it is required to examine the system under organization, technology, and individual (OTI) viewpoints (Gouglidis, Green, et al., 2016). The organization viewpoint is concerned with the groups of people who work together in an organized way for a shared purpose as well as any type of policies, processes, and procedures in the organization. The technology viewpoint references the implemented technologies in a system including the software, hardware, and network components, as well as any type of communication among them. The individual viewpoint brings awareness of a single person or entity and how it acts or behaves in a particular situation or under particular conditions. We have already covered the technology viewpoint sufficiently in previous sections of this chapter, so we move directly now to address the organizations and individual (people) aspects. It is also noteworthy that the last of these three viewpoints is able to enhance the awareness of the state of a system. In subsequent sections of this chapter, we consider the organization and individual (people) aspects, having already addressed the technology parts.

More specifically, the application of OTI (three viewpoints) may provide awareness of all the previously discussed four categories since system risks in external events may be identified by the organizational viewpoint. Likewise, system risks due to system and technology failures, or failed internal processes may be identified by the technological viewpoint. Similarly, system risks regarding actions of people might be identified by the individual viewpoint. Therefore, the application of OTI as a first point of contact toward an architecture capable of protecting ICS is capable of identifying in a timely manner various type of threads, and simultaneously considering current, evolving, or potential system risks due to a feedback process.

The components of the D^2R^2+DR resilience strategy can be used as an overarching process in the context of a wider risk management framework to provide the indicators and measurements to ensure an ongoing and effective monitoring of the networked systems. In the context of ISO 31000 (2009), a resilience framework may operate as part of the "monitoring and review" component (Schauer, 2018). The latter is responsible to provide indicators, progress measurement of conducting the risk management plan, risk reports, and reviews of design and effectiveness of the applied risk management measures implemented as an ongoing effectiveness monitoring of the complete framework (Austrian Standards Institute, 2010, Section 19). This component includes a constant feedback loop, taking the main and partial results from each step and evaluating their effectiveness. Risk-related information may be provided by other components of the framework, which could include the general organizational structure coming from "establishing the context" up to the estimation of the consequences and likelihood for identified threats under "risk analysis."

Situation Awareness and Resilience

SA is defined by the Committee on National Security Systems (2010) as "within a volume of time and space, the perception of an enterprise's security posture and its threat environment; the comprehension/meaning of both taken together (risk); and the projection of their status

into the near future" (p. 69). In addition, cyber SA can be defined as the part of SA that is concerned with the cyber environments (Franke & Brynielsson, 2014). Here we present related work with regards to the application of SA in ICS and elaborate on cyber SA in utility networks.

Utility networks are complex organizations where interactions take place among the assets of the network, the participating people and the ICS (Gouglidis, Shirazi, Simpson, Smith, & Hutchison, 2016). Any of these might be vulnerable to various types of threats, and therefore, become a risk for the network. Annual reports from agencies (e.g., ENISA) and major consultancy firms elaborate a list of threats to critical infrastructures. Nevertheless, considering the wide variety of ICS systems and their continually evolving environment, operational SA should be considered. Therefore, in the context of providing a holistic approach toward protecting utility networks, we propose the application of the OTI viewpoint-based approach as a first step toward gaining cyber SA on utility networks. Cyber SA is crucial to apply in networks to safeguard sensitive data, sustain fundamental operations, and protect infrastructures (all aspects of making the network more resilient).

Linking Technology, Organizations and People

A common approach toward conceptually understanding networks is to divide them into levels based on their function. Considering a utility network, for example, a simple three-level approach is adopted: field site, control center, and corporate (Wei, Lu, Jafari, Skare, & Rohde, 2011). Specific devices, boundaries, processes, etc. are then associated with each level, depending on the industry and network topology in question. More detailed layering approaches, such as the Purdue model, are able to provide further granularity by introducing a six-level view approach. Nevertheless, in all cases there is a clear indication of the complexity and interconnections between the levels. The application of the OTI viewpoints enables a broader view of a system (e.g., a utility network) and its levels as it can provide a representation of the whole system from the perspective of a related set of concerns—as stated before, this may help in increasing the level of threat awareness by identifying potential vulnerability-creating behaviors.

Organizations and People

The investigation of organizational aspects of networks may increase our understanding with respect to their resilience against vulnerabilities that arise from working conditions, technology affordances and social context. As Randell (2000) writes:

> how important it is to accept the reality of human fallibility and frailty, both in the design and the use of computer systems . . . all too often, the latest information technology research and development ideas and plans are described in a style which would not seem out of place in an advertisement for hair restorer. (p. 105)

In the context of organizations, issues of resilience are not simply technical issues resolved by technical means. Specifically, the investigation of organizational aspects will help in

understanding how these create vulnerabilities in the technology (e.g., networked computer systems), how organizational aspects may help mitigate vulnerabilities in the technology, and eventually how organizational functioning becomes vulnerable to utility failures. Randell's comment about the issues of designing dependable systems illustrates that making critical infrastructure systems inherently more resilient and safer is more than a simple technical problem. Instead, what a range of studies of critical infrastructure failure have illustrated is that such complex systems also have important organizational and human components that need to be understood and integrated into design to make such systems more resilient (Clarke, Hardstone, Rouncefield, & Sommerville, 2005; Dewsbury & Dobson, 2007). Consequently, we see the prevalence of what are termed "human and organizational factors" and a range of interdisciplinary approaches as a means of developing a more nuanced understanding—in the same way as we have striven to develop more nuanced understandings of resilience and its relationship to other very similar or related (and perhaps more researched) topics like risk, trust, dependability, and sensemaking.

Resilience and the Mental Models Used in Reasoning About Risk and the Importance of Trust

As part of our attempt to understand some of the human and organizational factors involved in resilience, we conducted ethnographic observations and interviews in various utility and information organizations in different parts of Europe (Gouglidis, Green, et al., 2016). In our analysis of the ethnographic data and in trying to understand the components of individual and organizational resilience, we were interested in unpacking people's ideas about risk and how these might relate to other notions such as trust or organizational resilience. Our focus was on how organizational members modeled risk as part of their organizational roles; how different models of risk might interact or impact on each other; how the models changed in response to organizational events; and how these models might be interrelated with notions of trust—in individuals and the organization—and thereby might affect individual and organizational resilience in the face of change and the possibility of failure.

The elicitation of mental models in risk studies was aimed at uncovering deficiencies in individuals' understanding of complex risks (e.g., Bostrom, Morgan, Fischhoff, & Read, 1994)—how they understand exactly what is going on in terms of various kinds of risk. Our work used unstructured interviewing and ethnography to get a contextualized understanding of how organizational members use particular interpretive schemes, heuristics, and other forms of discursive reasoning to deal with organizational risks. We therefore developed an analysis that is closer to notions of the social construction of risk—and Hilgartner's (1992) approach where risk objects come into prominence, or recede out of prominence, in a process termed *emplacement* and *displacement*. Emplacement occurs when the consequences of a risk become magnified, or the causes of risk seem to be less manageable and more likely. Displacement occurs when risk seems to come under greater control. Our primary concern was with how people's risk models perform this process of emplacement and displacement.

From our analysis, it was clear that organizational members have a wide variety of risk models that are not generally integrated, uniform, or self-consistent representations. These risk models tend to be, or at least emerge as, fragmentary and partial, and serve as discursive resources to justify a claim as much as resources for reasoning toward a claim. The function of risk models in one utility organization, for example, appeared much more often displacement than emplacement, but more emplacement than displacement for the information systems organization. Sometimes emplacement and displacement went together. For example, a risk may be emplaced to show how the organization has taken it seriously enough to displace it with strong controls—for example, by having a clear and monitored set of processes, such as having a clear reporting structure for email phishing attacks. One of the fieldwork sites had a clear notion of actual and potential risks, meaning that the organization could acknowledge that some risk existed, but had good grounds for not devoting resources to managing it—because, for example, it was argued that even if someone obtained access to the system they would not be able to do much, such as switching people's electricity off, since this depended on a different set of controls. Potential risks were in some sense theoretical and general—decontextualized and offering no reason for acting on them in this particular organization at this particular time. The fieldworker identified what appeared to be two registers of risk—the actual and the potential. At both field sites, cyber security risks were displaced by other risks—safety, usability, customer satisfaction—which were seen as substantially more important.

Different kinds of risk models were found in the analysis. Failure path models represented sequences of action that were required to bring about some kind of failure state. These enabled people to reason about how plausible different failure or cyber security scenarios were. For example, one respondent reasoned that risks were low because of the fact that an attacker would need one kind of expertise to gain access to computing devices, but a different kind of expertise to actuate physical devices. Technical boundary models involved representations of the technical system as a collection of devices that were strongly partitioned, and typically supplied by different providers. The boundaries represented boundaries of responsibility for risk and boundaries of competence. Sometimes people would say, "We can only do something about X but not Y" to indicate a residual uncertainty about a risk that was partially the responsibility of someone else and beyond their control. Experiential narrative models were stories of incidents or materialized risks of some sort. Narratives provided structured accounts of some issue or problem (in this case security risk) that had come into discourse. Often the narrative involved emplacing a risk, explaining an event in the recent past and then displacing it by reasoning about how controls had subsequently been brought in. The narrative sequence of some experienced event followed by some remedial action seemed to help people reason about security. Ordering models placed the explanation for a lack of interest in certain kinds of risk on priority—an ordering that put security risk well below other risks and other demands on resources. In the utility organization, the main risks were seen by some as being commercial, displacing cyber security risks; in the system's organization, the main risks were said to be seen as being physical. For an organizational actor, it may be less important to have a descriptive representation of risk than to have a list of actions and associated priorities. As our fieldworker remarked,

> more importance is given to safety at work due the deaths of some employees . . . in the past. The fear of court cases and also bad media coverage means that more

> money is invested in this area rather than in cyber-security, which is seen to be less important . . . it is easy to cover up cyber-security attacks. The repercussions of these breaches were also deemed to be less serious.

Cost–benefit models involved reasoning that risks were low because the costs to an attacker were high and benefits low. This was typically a risk displacement strategy. They also argued that possible risk controls were unnecessary as their cost exceeded their benefit. Abstract, or global attribute models were simple characterizations of the entire organization or some particular situation. Those of the utility organization were more optimistic. For example, some people had a simple model of sufficiency, a general belief that there were enough appropriate controls to nullify risk. In the systems organization these were more pessimistic, characterizing the organization as having a culture inappropriate to security in a number of ways.

In terms of our ideas about resilience, these different risk models could be a source of vulnerability or of resilience. It is the specific context and specific manifestation that may prove decisive. But it is instructive how wide-ranging the types of model are. They are qualitatively quite different and point to the resourcefulness of organizational members in coping with a world that is complex. Some of these demands involve having an appropriate representation of a conventional system, but others involve having an appropriate representation of other people's expectations and capacities, of norms and conventions, and so on. This means it will always be insufficient to assess mental models of security or resilience merely in terms of their technical correctness, as it will sometimes be more important how well they represent prevailing social issues and requirements. What is important is that there is an awareness within the system of how those models contribute to, or detract from, its security.

Related to ideas about risk, the fieldwork interviews and observations in the different utilities also provide insight into how workers perceive trust, and how trust is an implicit and taken-for-granted feature in the accomplishment of work and therefore a key aspect of resilience. The extent to which risk is perceived and acted upon is linked to some degree to the extent to which people, technology, and organizations are trusted. This in turn impacts on resilience, on the ability of the organization to respond to sudden change or failure. Lack of trust acts as a contributor to a range of problems—be it the poor quality of work resulting from collaboration or a failure to complete tasks at all. Trust and the degree and quality of trust existing between collaborating parties shapes the possibilities for how parties undertake and complete work. Collaboration within and between organizations presupposes trust.

Trust is generally assumed to be organizationally important and a key contributor to the prevention of organizational failure, as a number of studies have suggested. For example, the U.S. Government's Baker Report (Baker et al., 2007) highlighted lack of trust as a precursor to the 2005 explosion at BP's Texas City refinery that killed 15 people, noting that the

> single most important factor in creating a good process safety culture is trust [and] that employees and contractors must trust that they can report incidents, near misses, and other concerns—even when it reflects poorly on their own knowledge, skills, or conduct—without fear of punishment or repercussion. (p. 75)

In a similar fashion, the high reliability organization (HRO) literature sees trust as an important organizational and cultural feature of reliability. Trust is, however, an elusive and difficult to define concept since there are multiple and diverse perspectives. Contemporary research areas include reciprocating trust among teams (Serva, Fuller, & Mayer, 2005), trust in leaders (Burke, Sims, Lazzara, & Salas, 2007), and trust as a heuristics in decision-making and its effects on attitudes, perceptions, behaviors, and performance within organizational settings (Dirks & Ferrin, 2001).

Resilience as Sensemaking and Mindfulness

In this last section we consider how the topic of resilience might relate to other theoretical and empirical concepts in the organizational literature. One possible way of thinking about resilience is to compare our findings and approach with the reliability and dependability literature—specifically that connected to the idea of the HRO—or what we might perhaps want to rephrase as the "high resilience organization." Of particular interest, as far as resilience is concerned, are the concepts of sensemaking and mindfulness that are invoked when considering high reliability organizations (Snook, 2000; Weick, 1987, 2001; Weick & Sutcliffe, 2007; Weick, Sutcliffe, & Obstfeld, 2005).

An HRO is depicted as an organization that has accurate, precise, and commonly held understandings about current operations and the relationship between those operations and potential accidents (Cook & Rasmussen, 2005). A basic assumption is that accidents can be prevented through good organizational design and management and that HROs organize themselves in such a way that they are better able to notice and stop unexpected events. If they cannot halt such a development, they are resilient and able to swiftly restore the functioning of the system (e.g., Rochlin, La Porte, & Roberts, 1987). This approach is commonly contrasted with what is termed "normal accident theory" (Perrow, 1999)—although the approach to resilience may well be similar in both cases.

In terms of sensemaking for resilience and for the HRO, there is a range of research on sensemaking, across the individual, group/organizational, multiorganizational, and societal levels. Weick describes sensemaking as "a developing set of ideas with explanatory possibilities, rather than as a body of knowledge" (Weick, 1995, p. xi). Weick views the concept of sensemaking as a collective, social activity—a cognitive process that can be described through seven properties that "involves turning circumstances into a situation that is comprehended explicitly in words and that serves as a springboard into action" (Weick et al., 2005, p. 409). The seven properties of sensemaking appear equally applicable to resilience. Adapted, these properties appear as follows:

- *Social*: People do not discover resilience, rather they create it. In other words, organizational resilience is interactive.
- *Identity*: Resilience unfolds from identities. People develop identities for themselves during inexplicable events (e.g., as victim, fighter), and this identity can lock them into particular options.

- *Retrospect*: Resilience is constructed by reference to the past. Faced with the inexplicable, people often act their way out of ambiguity by talking about the past and assessing what they have said before about similar events, to discover what they should do and how they should think in the present.
- *Cues*: Resilience is developed as people deal with the inexplicable by paying attention to small cues that enable them to construct a larger story. They look for cues that confirm their analysis and, in doing so, ignore other less relevant information.
- *Ongoing*: Resilience is dynamic and requires continuous updating and re-accomplishment. Resilience requires that people stay attuned to what is happening around them—if not, they lose context and information.
- *Plausibility*: Resilience depends on robust and plausible analyses rather than fixation on a single plausible explanation of an event.
- *Enactment*: In inexplicable times, people have to keep moving. Recovery lies not in thinking and then doing, but in thinking while doing something.

Sensemaking, then, like resilience, involves the ongoing retrospective development of plausible images that rationalize what people are doing (Weick et al., 2005) and points to the need for rapid assessment of a constantly changing environment and to the constant reinterpretation of perceived reality. Taken together, these properties suggest that increased skill at sensemaking—and resilience—should occur when people are socialized to make do, to treat constraints as self-imposed, strive for plausibility, keep showing up, use the past to get a sense of direction in the present, and articulate descriptions that energize.

Ultimately for Weick et al. (2005), the language of sensemaking "captures the realities of agency, flow, equivocality, transience, re-accomplishment, unfolding, and emergence" (p. 410). The means by which this is best achieved according to Weick (2009) is by using the processes of mindfulness. According to Langer and Moldoveanu (2000), mindfulness has been used as a basis for investigating a number of research areas, including decision-making and has also been associated with organizational learning (Levinthal & Rerup, 2006). Following Langer's work, the idea of mindfulness has been extended from analysis at the individual level to analysis at the organizational level (Weick, Sutcliffe, & Obstfeld, 2008). In doing this, Weick et al. (2005) shifted the focus from individual mindfulness to collective mindfulness and "heedful interrelating" (Weick & Roberts, 1993). Heedful interrelating arises when people "act like they are under the direction of a single organizing centre . . . have a visualized representation of a group's meshed contributions . . . and bring group facts into existence" (Weick, 2009, p. 218). By analyzing data from HROs, Weick showed that individuals within these organizations collectively used five cognitive processes related to mindfulness to overcome a broad range of unexpected events.

Preoccupation With Failure

HROs are preoccupied with failures. There is a constant concern in HROs that error is embedded in ongoing day-to-day activities and that unexpected failures and limitations of foresight may amplify small errors. HROs realize that if separate small errors occur simultaneously, then the result could potentially be disastrous. Worrying about failure gives HROs

much of their distinctive quality, and this distinctiveness arises from the simple fact that failures are a rare occurrence. This means that HROs are preoccupied with something that is seldom visible. To foster organization-wide concern with failure, HROs encourage personnel at all levels to report errors when they arise and make the most of any failure that is reported.

Reluctance to Simplify Operations

A common property of organizations is that their members simplify tasks—either in the way work is carried out, or in the way they perceive risk. For HROs, this simplification is potentially dangerous as it limits the precautions people take and the number of undesired consequences that they envision. Simplification increases the likelihood of eventual surprise and allows anomalies to accumulate, intuitions to be disregarded, and undesired consequences to grow more serious. To resist temptations to simplify, HROs cultivate requisite variety, which takes such forms as diverse checks and balances, including a proliferation of committees and meetings, selecting new employees with nontypical prior experience, frequent job rotation, and retraining. Redundancy also forms an important component of HROs, not only in the form of system standbys and backups, but also in the form of scrutiny of information and the inclusion of conceptual slack—defined as "a divergence in analytical perspectives among members of an organization over theories, models, or causal assumptions pertaining to its technology or production processes" (Vogus & Welbourne, 2003, p. 13).

Sensitivity to Operations

HROs are attentive to the frontline where the real work is being done. When people have well developed SA, they can make continuous adjustments that prevent errors from accumulating end enlarging. This is achieved through a combination of collective story building, shared mental representations, situation assessing with continual updates, and knowledge of physical realities of the organization's systems.

Commitment to Resilience

Weick and Sutcliffe (2001) define resilience as "the process of being mindful of errors that have already occurred and correcting them before they worsen or cause more serious harm" (p. 67). People in HROs are encouraged to make their system transparent and their operational practices widely known. This helps people to appreciate weaknesses and manage them better. People in HROs are committed to resilience and actively work to keep errors small and improvise workarounds to keep systems functioning. HROs see this "firefighting" as evidence that they are able to contain the unexpected. This is in contrast to other organizations, where managers may perceive successful firefighting as evidence that they are distracted and therefore unable to do their normal work (Weick & Sutcliffe, 2001). HROs need to have a broad repertoire of actions they can roll out when required, including informal skill and knowledge-based networks that organize themselves when potentially dangerous situations arise.

Underspecification of Structures

Weick, Sutcliffe, and Obstfeld (2008) argue that HROs are failure-free despite their orderliness, not because of it. An orderly hierarchy can amplify errors, and higher-level errors tend

to amalgamate with lower level errors. This combination of errors is harder to understand when more interactive and complex. It is the very reliability that HROs cultivate that makes it possible for small errors to spread, accumulate, interact, and trigger serious consequences. To prevent this, HROs allow for underspecification of structures (also referred to as "deference to expertise"). Decisions may come from the top during normal times but during times of potential danger, decision-making migrates, and a predefined emergency structure comes into force. Decision-making can be made on the frontline, and authority is given to people with the most expertise, regardless of their rank. The decision-making structure in HROs is a hybrid of hierarchy and specialization. Decision-making authority therefore is shifted down to the lowest possible levels and clear lines of responsibility are put into place.

Mindfulness, then, is the

> combination of ongoing scrutiny of existing expectations, continuous refinement and differentiation of expectations based on newer experiences, willingness and capability to invent new expectations that make sense of unprecedented events, a more nuanced appreciation of context and ways to deal with it, and identification of new dimensions of context that improve foresight and current functioning. (Weick & Sutcliffe, 2001, p. 42)

For Weick, the five processes of mindfulness mobilize the resources for sensemaking.

Conclusion

The advance of digital technologies has substantially improved the resilience and efficiency of networked computer systems. These technologies provide various processes, including monitoring, control, and automation, to help achieve resilience. Networked systems are widely used for communication purposes and thus are essential. Yet, they face growing and evolving cyber-physical and social risks, as well as other challenges including natural disasters. These risks result not only from growing direct threats, but also from interdependencies and associated cascading effects. Ambitious investment in innovation is required to increase the resilience of networked systems, especially in the context of sensitive industrial sites and plants when protection measures against impacting events fail. The critical functions that sensitive industrial sites and plants provide, including safety and security, need to be resilient when adverse conditions present themselves. A holistic approach to resilience should include both technical and nontechnical approaches to promptly cope with cyber-physical and social-related threats to networked systems.

Our current and future work is concerned with using the technical, human, and organizational insights we have obtained from our studies of resilience and applying them to understand and develop resilient and secure industrial systems in the European and indeed the global economy. Of particular concern is the impact of cyberattacks on sensitive industrial sites as digital technologies become increasingly vital for ICS that control and monitor safety, security, and production processes. Manufacturing and industrial sites constitute a

critical component for the sustainable development of economies and society. These sites constitute an interdependent network of plants and facilities. Sensitive industrial sites and other industrial plants such as nuclear facilities produce or handle hazardous materials (e.g., radioactive materials, toxic chemicals, explosive materials). An attack or challenge at one of these sites could lead to significant environmental damage including loss of life and disruption of global supply chains.

Therefore, investment in research and innovation is required to increase the resilience of sensitive industrial sites and plants when protection measures against impacting events fail. The critical functions that sensitive industrial sites and plants provide, including safety and security, need to be preserved when adverse conditions present themselves. To minimize the associated risks, measures are necessary to prevent major accidents and to ensure appropriate preparedness and response should such accidents happen. Future research needs to be concerned with enhancing the resilience of ICT systems, ICS, and associated processes. Special attention must be paid to communications and information-sharing regarding about incidents and possible precursor indicators of cascading impacts that result from neighboring events.

Key Messages

1. Modern networked systems are critical infrastructures.
2. Modern networked computer systems need to be designed and engineered to have resilience as a major property.
3. Resilience is "the ability of a network or system to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation" (Sterbenz, Hutchison, et al., 2010)
4. Modern networked systems are complex and have three aspects that need to be considered in combination when building resilience into them: these are technology, organization, and people.
5. Investment in research and innovation is required to increase the resilience of sensitive industrial sites and plants when protection measures against impacting events fail. The critical functions that sensitive industrial sites and plants provide, including safety and security, need to be preserved when adverse conditions present themselves.

References

Adger, W. N., Hughes, T. P., Folke, C., Carpenter, S. R., & Rockström, J. (2005). Social-ecological resilience to coastal disasters. *Science*, *309*(5737), 1036–1039. doi:10.1126/science.1112122

Allenby, B., & Fink, J. (2005). Toward inherently secure and resilient societies. *Science*, *309*(5737), 1034–1036. doi:10.1126/science.1111534

Austrian Standards Institute. (2010). *Hrsg., Risikomanagement—Grundsätze und Richtlinien*. Bd. 1. Retrieved from https://shop.austrian-standards.at/action/de/public/details/353917/OENORM_ISO_31000_2010_02_01

Baker, J. A., III., Bowman, F. L., Erwin, G., Gorton, S., Hendershot, D., Leveson, N., . . . Wilson, L. D. (2007). *The report of the BP US refineries independent safety review panel*. Retrieved from http://www.csb.gov.assets/1/19/Baker_panel_report1.pdf

Bostrom, A., Morgan, M. G., Fischhoff, B., & Read, D. (1994). What do people know about global climate change? 1. Mental models. *Risk Analysis*, *14*(6), 959–970. doi:10.1111/j.1539-6924.1994.tb00065.x

Bruncak, R., Bohra, N., Simpson, S. P., Rohrer, J. P., Sterbenz, J. P., Hutchison, D., . . . de Meer, H. (2011). *Resilience and survivability for future networking: Framework, mechanisms, and experimental evaluation*. ResumeNET. Retrieved from https://pdfs.semanticscholar.org/33a2/14185f74972c9b1b6f9f296a2de1af1ee8ad.pdf

Burke, C. S., Sims, D. E., Lazzara, E. H., & Salas, E. (2007). Trust in leadership: A multi-level review and integration. *The Leadership Quarterly*, *18*(6), 606–632. doi:10.1016/j.leaqua.2007.09.006

Cebula, J. J., & Young, L. R. (2010). *A taxonomy of operational cyber security risks* (No. CMU/SEI-2010-TN-028). Pittsburgh, PA: Carnegie Mellon University, Software Engineering Institute.

Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys (CSUR)*, *41*(3), 15. doi:10.1145/1541880.1541882

Clarke, K., Hardstone, G., Rouncefield, M., & Sommerville, I. (2005). *Trust in technology: A socio-technical perspective*. Dordrecht, The Netherlands: Springer.

Committee on National Security Systems. (2010). *National information assurance (IA) glossary* (CNSS Instruction No. 4009). Retrieved from https://www.cdse.edu/documents/toolkits-issm/cnssi4009.pdf

Cook, R., & Rasmussen, J. (2005). "Going solid": A model of system dynamics and consequences for patient safety. *Quality and Safety in Health Care*, *14*(2), 130–134. doi:10.1136/qshc.2003.009530

Dewsbury, G., & Dobson, J. (2007). *Responsibility and dependable systems*. Heidelberg, Germany: Springer Verlag.

Dirks, K. T., & Ferrin, D. L. (2001). The role of trust in organizational settings. *Organization Science*, *12*(4), 450–467. doi:10.1287/orsc.12.4.450.10640

ENISA Glossary. (2019). *Risk management*. Retrieved from https://www.enisa.europa.eu/topics/threat-risk-management/risk-management/current-risk/bcm-resilience/glossary/p-z

Esposito, C., Gouglidis, A., Hutchison, D., Gurtov, A. V., Helvik, B. E., Heegaard, P. E., . . . Rak, J. (2018). On the disaster resiliency within the context of 5G networks: The RECODIS experience. In *European Conference on Networks and Communications 2018.* New York, NY: IEEE.

Franke, U., & Brynielsson, J. (2014). Cyber situational awareness: A systematic review of the literature. *Computers & Security*, *46*, 18–31. doi:10.1016/j.cose.2014.06.008

Górniak, S., Tirtea, R., Ikonomou, D., Cadzow, S., Gierszal, H., Sutton, D., . . . Vishik, C. (2011). *Enabling and managing end-to-end resilience*. ENISA. Retrieved from https://www.enisa.europa.eu/publications/end-to-end-resilience

Gouglidis, A., Green, B., Busby, J., Rouncefield, M., Hutchison, D., & Schauer, S. (2016, September). Threat awareness for critical infrastructures resilience. In *2016 8th International Workshop on Resilient Networks Design and Modeling (RNDM)* (pp. 196–202). Halmstad, Sweden: IEEE. doi:10.1109/RNDM.2016.7608287

Gouglidis, A., Shirazi, S. N., Simpson, S., Smith, P., & Hutchison, D. (2016). A multi-level approach to resilience of critical infrastructures and services. In *2016 23rd International Conference on Telecommunications (ICT)* (pp. 1–5). Thessaloniki, Greece: IEEE. doi:10.1109/ICT.2016.7500410

Gouglidis, A., Hu, V. C., Busby, J. S., & Hutchison, D. (2017). Verification of resilience policies that assist attribute based access control. In *Proceedings of the 2nd ACM Workshop on Attribute-Based Access Control* (pp. 43–52). Scottsdale, AZ: ACM. doi:10.1145/3041048.3041049

Hilgartner, S. (1992). The social construction of risk objects: Or, how to pry open networks of risk. In J. F. Short & L. Clarke (Eds.), *Organizations, uncertainties, and risk* (pp. 39–53). Boulder, CO: Westview Press.

Hollnagel, E., Paries, J., Woods, D. W., & Wreathall, J. (2010). *Resilience engineering in practice: A guidebook*. Boca Raton, FL: CRC Press.

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

Hutchison, D., & Sterbenz, J. P. (2018). Architecture and design for resilient networked systems. *Computer Communications*, *131*, 13–21. doi:10.1016/j.comcom.2018.07.028

International Organization for Standardization. (2009). *ISO 31000: 2009 Risk management—Principles and guidelines*. Geneva, Switzerland: Author.

König, S., Gouglidis, A., Green, B., & Solar, A. (2018). Assessing the impact of malware attacks in utility networks. In S. Rass & S. Schauer (Eds.), *Game theory for security and risk management* (pp. 335–351). Basel, Switzerland: Birkhäuser.

Langer, E. J., & Moldoveanu, M. (2000). The construct of mindfulness. *Journal of Social Issues*, *56*(1), 1–9. doi:10.1111/0022-4537.00148

Levinthal, D., & Rerup, C. (2006). Crossing an apparent chasm: Bridging mindful and less-mindful perspectives on organizational learning. *Organization Science*, *17*(4), 502–513. doi:10.1287/orsc.1060.0197

Machuca, C. M., Secci, S., Vizarreta, P., Kuipers, F., Gouglidis, A., Hutchison, D., . . . Ristov, S. (2016). Technology-related disasters: A survey towards disaster-resilient software defined networks. In *2016 8th International Workshop on Resilient Networks Design and Monitoring (RNDM)* (pp. 35–42). Halmstad, Sweden: IEEE. doi:10.1109/RNDM.2016.7608265

Moteff, J. D. (2012). *Critical infrastructure resilience: The evolution of policy and programs and issues for congress*. Congressional Research Service Reports. Washington, DC: Library of Congress.

Obregon, L. (2015). *Secure architecture for industrial control systems*. The SANS Institute. Retrieved from https://www.sans.org/reading-room/whitepapers/ICS/secure-architecture-industrial-control-systems-36327

Perrow, C. (1999). *Normal accidents: Living with high-risk technologies* (2nd ed.). Chichester, UK: Princeton University Press.

Randell, B. (2000). Turing memorial lecture facing up to faults. *The Computer Journal*, *43*(2), 95–106. doi:10.1093/comjnl/43.2.95

Rochlin, G. I., La Porte, T. R., & Roberts, K. H. (1987). The self-designing high-reliability organization: Aircraft carrier flight operations at sea. *Naval War College Review*, *40*(4), 76–90.

Schauer, S. (2018). A risk management approach for highly interconnected networks. In S. Rass & S. Schauer (Eds.), *Game theory for security and risk management* (pp. 285–311). Cham, Switzerland: Birkhäuser.

Sedgewick, A. (2014). *Framework for improving critical infrastructure cybersecurity* (Version 1.1). NIST Cybersecurity Framework. doi:10.6028/NIST.CSWP.02122014

Serva, M. A., Fuller, M. A., & Mayer, R. C. (2005). The reciprocal nature of trust: A longitudinal study of interacting teams. *Journal of Organizational Behavior*, *26*(6), 625–648. doi:10.1002/job.331

Shirazi, S. N., Gouglidis, A., Farshad, A., & Hutchison, D. (2017). The extended cloud: Review and analysis of mobile edge computing and fog from a security and resilience perspective. *IEEE Journal on Selected Areas in Communications*, *35*(11), 2586–2595. doi:10.1109/JSAC.2017.2760478

Smith, P., Hutchison, D., Sterbenz, J. P., Schöller, M., Fessi, A., Karaliopoulos, M., & Plattner, B. (2011). Network resilience: a systematic approach. *IEEE Communications Magazine*, *49*(7), 88–97. doi:10.1109/MCOM.2011.5936160

Snook, S. (2000). *Friendly fire*. Chichester, UK: Princeton University Press.

Sterbenz, J. P., Hutchison, D., Çetinkaya, E. K., Jabbar, A., Rohrer, J. P., Schöller, M., & Smith, P. (2010). Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Computer Networks*, *54*(8), 1245–1265. doi:10.1016/j.comnet.2010.03.005

Vogus, T. J., & Welbourne, T. M. (2003). Structuring for high reliability: HR practices and mindful processes in reliability-seeking organizations. *Journal of Organizational Behavior*, *24*(7), 877–903. doi:10.1002/job.221

Wei, D., Lu, Y., Jafari, M., Skare, P. M., & Rohde, K. (2011). Protecting smart grid automation systems against cyberattacks. *IEEE Transactions on Smart Grid*, *2*(4), 782–795. doi:10.1109/tsg.2011.2159999

Weick, K. E. (1987). Organizational culture as a source of high reliability. *California Management Review*, *29*(2), 112–127. doi:10.2307/41165243

Weick, K. E. (1995). *Sensemaking in organizations*. Thousand Oaks, CA: SAGE.

Weick, K. E. (2001). *Making sense of the organization*. Oxford, UK: Wiley-Blackwell.

Weick, K. E. (2009). *Making sense of the organization: Vol. 2: The impermanent organization*. Chichester, UK: John Wiley.

Weick, K. E., & Roberts, K. H. (1993). Collective mind in organizations: Heedful interrelating on flight decks. *Administrative Science Quarterly*, *38*(3), 357–381. doi:10.2307/2393372

Weick, K. E., & Sutcliffe, K. M. (2001). *Managing the unexpected: Assuring high performance in an age of complexity*. San Francisco, CA: Jossey-Bass.

Weick, K. E., & Sutcliffe, K. M. (2007). *Managing the unexpected: Resilient performance in an age of uncertainty.* San Francisco, CA: John Wiley.

Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. (2005). Organizing and the process of sensemaking. *Organization Science*, *16*(4), 409–421. doi:10.1287/orsc.1050.0133

Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. (2008). Organizing for high reliability: Processes of collective mindfulness. In A. Boin (Ed.), *Crisis management:* (Vol. III, pp. 31–66). London, UK: SAGE.