Citation: Alderson, D.L., 2019, Overcoming Barriers to Greater Scientific Understanding of Critical Infrastructure Resilience, in M. Ruth & S. G. Reisemann (Eds), Handbook on Resilience of Socio-technical Systems, Edward Elgar Publishing, Northampton, MA.

5 Overcoming barriers to greater scientific understanding of critical infrastructure resilience

David L. Alderson

INTRODUCTION

There is a longstanding connection between 'infrastructure' and societal prosperity. The success of the ancient Roman Empire is commonly attributed in part to roads and bridges that enabled the movement of armies and commerce (Kaszynski 2000, pp. 9–11), while Roman dams and aqueducts supported the capture, movement and consumption of water for drinking, bathing and agriculture. Zeihan (2014) identifies access to navigable waterways as a key source of geopolitical advantage, tracing this influence in historical societies dating back to ancient Egypt, and noting that the continental United States (US) has a natural and tremendous advantage as an economic superpower because it happens to have more than half of all navigable waterways worldwide.

As documented by Brown (2006), many of the early infrastructure systems in the US in the late 1700s to 1800s were transportation systems developed initially to support economic growth (that is, roads, canals and railroads). However, in the early 1900s, the development and planning for these systems became more closely tied to war preparation, specifically mobilization for World War I and World War II. Much of the large spending for infrastructure by the US Public Works Administration (PWA) as part of President Roosevelt's New Deal had the dual purpose of providing jobs for economic growth as well as supporting wartime preparation (Brown 2006, pp. 15–16).

The modern study of infrastructure systems in the US was initiated by Presidential Executive Order 13010 on 15 July 1996, which established the President's Commission for Critical Infrastructure Protection (PCCIP) to investigate 'the scope and nature of the vulnerabilities of, and threats to, critical infrastructures' from both 'physical and cyber threats' (Executive Order 13010). This Executive Order also introduced the term 'critical infrastructure' as systems that are 'so vital that their incapacity or destruction would have a debilitating impact on the defense or economic security of the United States'. More recent definitions (for example, in the Homeland Security Act of 2002) have followed this basic concept. A key recognition in the charter of the PCCIP is the rapid growth of the Internet as a common information backbone connecting formerly disparate systems. As noted in its final report:

> More than any other country, we rely on a set of increasingly accessible and technologically reliable infrastructures, which in turn have a growing collective dependence on domestic and global networks. This provides great opportunity, but it also presents new vulnerabilities that can be exploited. It heightens risk of cascading technological failure, and therefore of cascading disruption in the flow of essential goods and services. (PCCIP 1997, pp. 4–5)

In the years following the attacks of 11 September 2001, US attention on critical infrastructure (CI) evolved from an initial focus on prevention (primarily against

terrorist attacks) to one of resilience against all hazards – to include storms such as Hurricane Katrina (2005) and Superstorm Sandy (2012) as well as engineering failures such as the Northeast Blackout (2003) and the Deepwater Horizon oil spill (2010). Accordingly, policy documents over the past decade have made increasing use of the term 'resilience' as the guiding objective for understanding and designing CI systems (Moteff 2015). The National Academies report, *Disaster Resilience: A National Imperative*, provides a review of the broad problem and also documents many of the federal initiatives to strengthen resilience (NRC 2012). Many of these are ongoing efforts by the US Departments of Homeland Security (DHS), Energy (DOE), and Defense (DOD) – often executed through the national laboratories and other federally funded research centers – to study the resilience of regional and national infrastructure systems. As a result, there are now considerable resources being devoted by federal funding agencies to the topic of resilience for critical infrastructure systems, communities, and society in general (for example, CIRI 2017; NIST 2017; NSF 2017; Rockefeller Foundation 2017).

Despite more than two decades of heightened policy attention on 'the problem of critical infrastructure', the general consensus is that infrastructure in the US is in poor condition and not getting better. The 2017 Infrastructure Report Card issued by the American Society of Civil Engineers (ASCE) scored a grade of 'D+' overall, with railroads as the only sector scoring better than 'C' (ASCE 2017). In the meantime, the impacts of disaster events in the US and worldwide – measured in terms of damage and the size of affected populations – seem only to be increasing (see Figure 5.1). Recent events such as Superstorm Sandy in 2012 demonstrate how the consequences of a disaster event can be exacerbated by the interruption of critical lifeline infrastructures, such as electricity, water and transportation (Flynn 2015).

There are many reasons why the US has been slow to improve the resilience of its infrastructures. As documented in Flynn (2015), among these reasons are a lack of economic incentives for investment together with general organizational and governance impediments. More fundamentally, however, we still lack a systematic understanding of how to assess and improve the resilience of our infrastructure systems.

In this chapter we argue that progress in our scientific understanding of resilience for critical infrastructure systems has been limited because of several fundamental barriers that are restricting the research community as a whole. These barriers include (1) the interdisciplinary nature of CI systems, (2) an overemphasis on threat scenarios, (3) an inability to share information about real systems and (4) a lack of understanding about resilience itself. We elaborate on each barrier and offer recommendations for potential solutions, identifying ongoing efforts at progress for each.

BARRIER 1: THE INTERDISCIPLINARY NATURE OF CI

Fundamentally, infrastructures lie at the intersection of social, technological and environmental systems. Sometimes these are labelled 'human', 'built' and 'natural', respectively. Therefore, the study of CI is an inherently interdisciplinary endeavor, for at least five reasons:

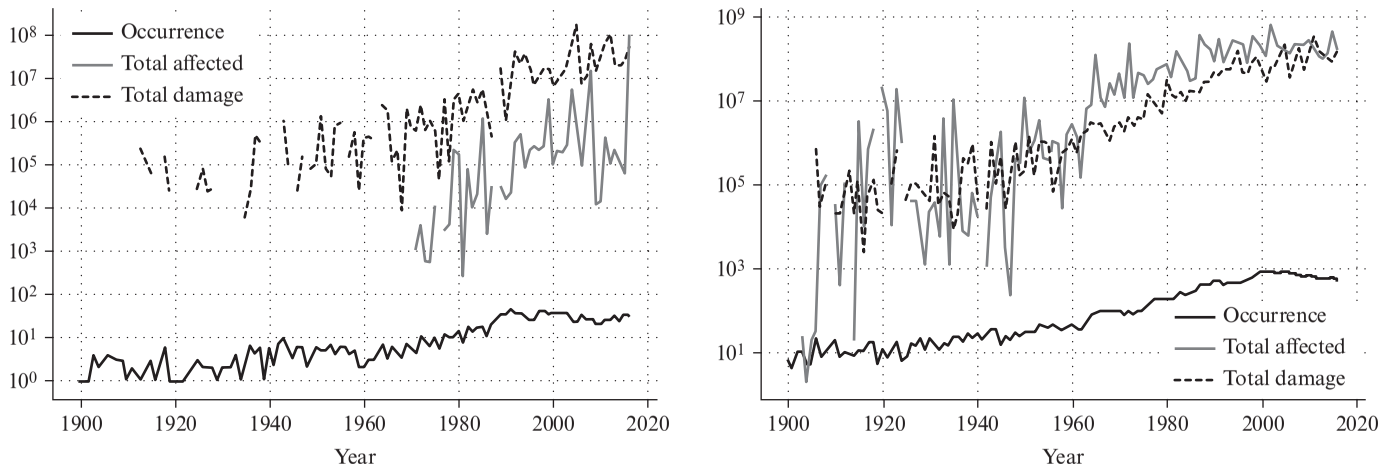

Note: Here, a 'disaster event' is a disaster that satisfies at least one of the following criteria: ten or more people dead, 100 or more people affected, the declaration of a state of emergency or a call for international assistance. Note the logarithmic scale on the vertical axis.

Source: EM-DAT: The Emergency Events Database, Université catholique de Louvain (UCL), CRED, D. Guha-Sapir, www.emdat.be (accessed 24 July 2018), Brussels.

Figure 5.1 Incidence of natural, technological and complex disasters for the United States (left) and the world (right) from 1900 to 2016, measured in terms of the number of disaster events (occurrence), total population affected, and total damage (000s US dollars)

1. The purpose of infrastructure systems is primarily to support human activity and/or needs. Humans are the ultimate users of these systems. As articulated by Hart et al. (2014): 'infrastructure problems are, at their root, social problems because infrastructures deliver goods and services that both people individually and society collectively need to live and function'. This motivates participation by the social sciences.
2. Infrastructure systems are geospatial. Since humans themselves are geographically dispersed, the primary purpose of infrastructure in many cases is to support the movement of resources, goods, people or information. This motivates participation by the sciences of geography and spatial economics, among others.
3. Infrastructure systems interact with and shape the natural environment. Much of the current debate about large public works projects often centers on the potential impact on environmental systems, and there has been considerable progress over the past several decades to put environmental issues on par with those in favor of economic development. This motivates participation by the environmental sciences, which themselves are interdisciplinary.
4. Infrastructure systems use varying levels of technology. However, construction of such systems typically requires a large upfront cost, with the subsequent expectation that these systems will persist for a long time. The fixed nature of these systems requires that they are designed to last, and much of traditional engineering has based its design on the assumption that past experience is representative of the future (for example, design to withstand 100-year events). However, growing evidence that 'stationarity is dead' (Milly et al. 2008) is pushing the need for designs that are upgradable and/or modifiable so that we have flexibility and extensibility when the demands on infrastructure outlive its designed use. This motivates participation by the engineering sciences.
5. Infrastructure systems typically involve a mix of public and private systems. It is generally believed that more than 85 percent of the critical infrastructure in the US is privately owned and operated (US GAO 2006). Yet, the services provided by these infrastructures are often considered a public good and vital to the welfare of society at local, regional and national levels. This motivates participation by the political sciences.

Collectively, we must accept that there is no single vantage point from which we can see all aspects of this problem. Thus, it is easy to assert a potential advantage to a multidisciplinary approach to critical infrastructure.

However, interdisciplinary research is hard, for a variety of reasons. First and foremost, there is no shared lexicon, so there can be different terms for the same concept and similar terms can have different meanings. This creates the need for continuous translation. Perhaps more importantly, attempts to look exclusively through the lens of a single discipline typically rely on key assumptions – enabling specific analysis techniques used for gaining knowledge and understanding – that in turn are often incompatible with other perspectives. This is sometimes known as the 'N-cultures problem' (Rouse 1982). Therefore, trying to integrate results that rely on mutually compatible assumptions is often difficult, if not impossible. Integration between engineering and social science is known to be particularly hard (NRC 2006). Worse yet, relaxing key assumptions in a discipline often has the effect of creating impotence, which can be threatening. Thus, a

topic of study that simultaneously spans multiple disciplines but violates the assumptions of each is particularly problematic.

Moreover, the practice of interdisciplinary research itself is still evolving. Krishnan (2009) considers the merits and incentives for interdisciplinary studies from the diverse perspectives of anthropology, sociology, philosophy, management and education. Krishnan contrasts the forces favoring the formation of interdisciplinary studies with the factors that encourage traditional disciplinary views, noting on balance a tendency by universities and the academy more broadly in favor of the persistence of disciplines. More specifically, he argues that the disciplinary view is strongly supported for anthropological reasons (that is, disciplines help to protect knowledge and identity from outside infringement), sociological reasons (that is, disciplines are hard to overcome because of the self-interest of power groups), educational reasons (that is, disciplines provide coherence and stability in curricula which is helpful to students), as well as philosophical reasons (that is, disciplines are needed to validate claims about what constitutes truth). The process of tenure in academic departments can be viewed as a type of tribalism that is self-perpetuating. In contrast, Krishnan (2009, p. 45) notes that only the management perspective (that is, the need 'to prepare pupils and students for economic participation or for the job market') really argues in favor of an interdisciplinary view, as rarely do messy problems in the real world fall naturally along disciplinary lines.

Many funding agencies (for example, NSF 2017; US DOD MURI 2017) have recognized the need for interdisciplinary research on complex networks, critical infrastructure, and so on, and in many cases multidisciplinary teams and techniques are now a requirement for funding. Although researchers have become more skilled at writing grant proposals that represent interdisciplinary research teams, it remains unclear whether or not these teams themselves are producing the right work. This is in part because researchers, ever adept at recognizing the signals in calls for proposals, are getting more skilled at 'how to write about an interdisciplinary team', with the research itself all too often partitioned into disparate pieces executed by standalone researchers who retreat into their individual disciplines. Cynically, sometimes the most interdisciplinary part of the research occurs during grant writing.

While the recognition that interdisciplinary research is hard is not particularly insightful, an acceptance that the future of infrastructure research must embrace a multidisciplinary view can be helpful (Seager et al. 2017). It also sets the stage for understanding some of the other barriers to greater scientific progress in this problem domain.

BARRIER 2: OVEREMPHASIS ON PREDEFINED THREAT SCENARIOS

The size, scope, and interconnected nature of critical infrastructures makes it nontrivial to understand, much less predict, their behavior. This is particularly true when these systems are stressed outside their normal operating regimes.

Yet, it is exactly during extreme events when we are most concerned about infrastructure behavior. We want our infrastructure systems to be resilient in the presence of extreme weather, component failures, deliberate attacks, and more. As a result, much of the analysis to date on critical infrastructures has focused explicitly on specific threat scenarios. A

common motivating question for such analyses is, 'What will be the consequence if [event X] occurs?'

The problem with this as a starting point is that it puts the emphasis on the extreme event, often resulting in a type of 'tunnel vision' on details and circumstances that can be secondary to an understanding of how an infrastructure systems works (or could work) in such situations. It also creates a tendency to ignore situations outside the specified scenario (that is, 'That is not of concern because it wouldn't happen in event X'), at the expense of broader understanding of how infrastructure systems respond or adapt to disruptive events. Therefore, it becomes difficult to leverage learning outcomes beyond the stated scenario, requiring new (and often repetitive) analyses when stakeholders next become concerned about event Y, then event *Z*, and so on.

This perspective also misses a key point related to infrastructure systems: rarely is any system working in an ideal, nominal way. Instead, most systems are continuously facing a variety of (typically small) disruptions – resulting from unexpected failures, scheduled maintenance or environmental shocks – and they are continuously adapting to them with various formal protocols and/or ad hoc workarounds. That is, the normal operation of critical infrastructures includes the means to adapt to disruptive events, and it is the explicit job of the 'operator' of an infrastructure system to manage the resources and activities to make sure that 'everything keeps working' even when things do not go according to plan. Thus, a threat-based view of infrastructure that focuses on extreme events tends to ignore or discount this 'operational view' of infrastructure that is fundamental to understanding why infrastructures are built as they are and how such systems routinely adapt, respond and recover in the presence of disruptive events.

In contrast, if we start with an understanding of the objectives, constraints and decisions governing the operation of an infrastructure system, then it becomes natural to assess not only the consequences associated with event X, but any event. Moreover, we can ask 'What are the events that result in the worst possible consequences?', as well as 'What sets of components are most critical to the operation of this system?' and 'What investments will most improve the system's ability to respond to such events?' See Alderson et al. (2013, 2014, 2015) for detailed discussions about how to answer these questions.

The key point is that a focus on threat-based analyses tends to fall into the trap of being reactive, not proactive, based on the most recent disaster events. For example, following a hurricane, we tend to observe an increase in the number of hurricane-related studies. However, the potential impact of a hurricane on, say, an electric power system, could be largely the same as that following wildfire or a flood, if the same system components are affected. Focusing on the system and its operation, rather than the disruptive event, has the potential to create analysis that is more insightful and broadly applicable to surprise events outside our recent experience.

BARRIER 3: INABILITY TO SHARE INFORMATION ABOUT REAL SYSTEMS

The study of real infrastructure systems almost always results in the identification of weaknesses, vulnerabilities and/or areas for improvement. This simple fact creates strong

disincentives among infrastructure stakeholders for sharing such information, for a variety of reasons.

The owners and operators of such infrastructures fear that weaknesses:

- might be exploited by competitors;
- might make them the target of regulators imposing additional requirements that constrain their ability to operate effectively or profitably; or
- might make them a target of lawsuits by customers.

Government officials fear that such weaknesses:

- might be exploited by terrorists or other adversaries, thereby creating national security vulnerabilities; or
- might make them vulnerable politically.

Despite these strong disincentives, DHS policy on critical infrastructure is centered on information sharing among public–private partnerships, and it has had to invent a new type of classification to support restricted sharing of information: protected critical infrastructure information. Protected critical infrastructure information (PCII) provides assurances for collecting and maintaining information – for example, it creates exemptions from Freedom of Information Act (FOIA) requests – but it fundamentally prevents sharing of information beyond those who have been trained, certified, and have a demonstrated 'need to know'. As noted by Moteff (2015, p. 28): 'While the federal government is trying to increase the amount of information shared among appropriate stakeholders, it is also trying to maintain a tight control (short of classification) on who gets to see what information.'

Moreover, government-sponsored studies on critical infrastructure systems are commonly focused on three basic questions related to national security (see Box 5.1). Hesitation on the part of other stakeholders, especially infrastructure owners and operators, can make it difficult to obtain the underlying data necessary to conduct these studies. In some cases, there can be restrictions beyond PCII on sharing data and analysis.

A result is that when government entities or their agents (for example, national laboratory researchers working on behalf of government clients) produce specific analysis and/or recommendations for real systems on the questions in Box 5.1, very rarely is that

BOX 5.1 THREE PRIMARY QUESTIONS ABOUT CRITICAL INFRASTRUCTURE FROM A NATIONAL SECURITY PERSPECTIVE

- *Consequence estimation*. Given some imagined scenario, what will be the consequence (for example, lives lost, damage, loss of service)? How bad could things be?
- *Critical components*. What are the most critical (sets of) components in the system? Where should we focus our attention?
- *Mitigation*. What could be done to mitigate the potential consequences of the imagined scenario? What is it going to cost? What is the potential return on investment?

analysis shared with anyone beyond the study team and the sponsoring agency. Although there are mechanisms for industry insiders to share information about vulnerabilities and best practices with other industry insiders (for example, Information Sharing and Analysis Centers, or ISACs; see https://www.nationalisacs.org, accessed 6 August 2018), very rarely do other researchers benefit from this analysis and insight. Therefore, the infrastructure analysis community as a whole is not moving forward.

The same tensions exist for models. The sensitive nature of infrastructure data sometimes makes detailed infrastructure models (built and executed to obtain insight on specific systems) unsuitable for publication or sharing. As a result, there has often been duplication of effort across the national laboratory and university system. We could argue that some diversity and competition is healthy and desirable, except that in many cases an individual laboratory or research center might have very little situational awareness about what others are doing. Again, the outcome is limited progress by the research community as a whole.

Within the research enterprise, both individual investigators as well as research program managers suffer because there are no real datasets available for use by modelers or algorithm developers, for purposes of testing, validation and/or benchmarking. Therefore, we often find sophisticated analysis techniques being applied to notional (and often unrealistic) data, with the result that these tools and techniques either look inappropriate or are of little practical use. Pure modeling abstractions (for example, based on graph topology) can sometimes yield insight into basic mechanisms, but these are often superficial when viewed in the domain-specific context of the infrastructure systems.

The basic foundation for our academic tradition is to learn by building on top of what those before us have learned. While that works for methods and procedures, it is extremely limited in the context of critical infrastructure systems because so much of what we need to learn is hiding in the specific data of specific infrastructures. If that is always kept secret, then comprehensive learning becomes impossible, and we can expect to make the same mistakes over and over again.

In summary, the inability to share information about real infrastructure systems is inhibiting progress within the scientific research community because it imposes significant friction for peer review, as well as the dissemination of state of the art techniques, best practices and lessons learned.

BARRIER 4: A LACK OF UNDERSTANDING ABOUT RESILIENCE ITSELF

The notion of resilience has become hyper-popular in discussions of disaster preparedness, critical infrastructure, and homeland security more broadly. Yet, there is no consensus on a definition for what resilience is or how to assess it.

The US government has put forward official definitions in policy documents, the most recent being Presidential Policy Directive 21 (PPD-21) which focuses on critical infrastructure and defines resilience as 'the ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions' (The White House 2013, p. 1). While there is a recognized need for resilience metrics by the National Academies (NRC 2012) and others, it remains unclear how to do this effectively; see Cutter (2016) for a recent survey of measures of resilience.

A comprehensive review of the relevant literature that attempts to define resilience is beyond the scope of this chapter. However, the first use of resilience to describe adaptation in systems is often attributed to Holling (1973) who used it in the context of ecology; for a more recent socio-ecological perspective, see Folke (2006). Importantly, there is a large and growing body of literature on resilience engineering that has its roots in systems safety and organizational theory (see Hollnagel et al. 2006, 2008, 2011; Nemeth et al. 2009; Nemeth and Hollnagel 2014).

The topic of resilience in engineering has received more than a decade of discussion and debate; see Woods (2006), Hale and Heijer (2006), Haimes (2006), Madni and Jackson (2009), Park et al. (2013), Alderson et al. (2015), Thoma et al. (2016), and references therein for discussion. Westrum (2006), Martin-Breen and Anderies (2011), and Zolli and Healy (2012) contrast different views of resilience in a diversity of contexts, while Alexander (2013) traces the etymological roots of resilience across disciplines. Rose (2004, 2007, 2009, 2017) has led the discussion on economic resilience. Longstaff et al. (2010, 2013) consider community resilience explicitly, both in terms of its definition and a framework for evaluation. Meerow and Newell (2015) and Meerow et al. (2016) review resilience in the context of urban planning and design.

Most relevant to the discussion here, Woods (2015) characterizes four ways in which resilience is being used in the literature.

1. *Resilience as rebound* from traumatic events, often measured in terms of the cumulative deviation in performance over the recovery period.
2. *Resilience as robustness*, that is, the ability to adapt from well-modeled disruptions.
3. *Resilience as graceful extensibility*, that is, the capacity of a system to stretch and/or extend its performance near and beyond system boundaries.
4. *Resilience as sustained adaptability*, that is, the ability to manage adaptive capacities (via architectures and rules of governance) near hard limits in tradeoff spaces.

Woods argues that the first two uses (despite being the most prevalent) are not really resilience at all, because they fail to capture the way in which systems sense, anticipate and adapt to a constantly changing environment, especially when surprise circumstances challenge the system to operate near its performance boundaries. The basic idea is that resilience can be understood only as an expression of system behavior in response to an event, rather than something inherent to the system itself. This suggests that attempts to measure system resilience as a system attribute will ultimately fail because resilience is something that a system does, not something that a system has (Hollnagel et al. 2006).

If this is correct, and resilience is not something that a system has, then it simply might not be possible to increase the quantity of resilience in the way that might enhance other attributes. Instead, we perhaps need to evaluate and improve the means of interaction (that is, the protocols) that govern how elements of the system work together (Hollnagel et al. 2006; Alderson and Doyle 2010). See Eisenberg (2018) for a detailed discussion about how thinking about critical infrastructure resilience needs to change.

The reality here is that there is a lot that we still do not understand about what creates resilience in systems, particularly when these systems involve a mix of humans and automated technologies. Also, to the extent that we do not have a shared vision, it is difficult to create a resilient future.

RECOMMENDATION 1: THE NEED FOR CASE STUDIES AT SCALE

A key to understanding when and how systems adapt (or fail) in the presence of surprise events is the context in which these events occur.

Traditional safety engineering has focused on detailing the events leading up to an accident (for example, the story of the *Columbia* space shuttle accident; see Woods 2005), and there is growing interest within the resilience engineering community for documenting not only the things that go wrong but the things that go right in stories of resilience. As articulated by Woods and Branlat (2011, p. 129, original emphases):

> the data to measure resilience comes from observing/analysing how the system has adapted to disrupting events and changes *in the past* (Woods 2009, p. 500). Past incidents provide information about how a system was both *brittle*, by revealing how it was unable to adapt in a particular evolving situation, and *resilient*, by revealing aspects of how it routinely adapted to disruptions (Woods and Cook 2006). Analysis of data about *how* the system adapted and to *what*, can provide a characterisation of how well operational systems are prepared in advance to handle different kinds of challenge events and surprises. (Hollnagel et al. 2006)

There is a large and growing literature of case studies focused on these 'adaptive cycles', primarily within the safety research community – for example, in medicine (Cook 2006), in firefighting (Woods and Branlat 2011), in aviation (Dekker 2014) – with considerable attention paid to the human factors that provide graceful extensibility in the presence of surprise.

Context is essential also for understanding how critical infrastructure systems adapt and/or fail. However, as noted by Thoma et al. (2016), the results of the aforementioned case studies in safety science and the emerging field of resilience engineering are perhaps of limited value to critical infrastructure systems because their insights tend to be generic and focus on the human factors at work. Although resilience engineering, as a comparably interdisciplinary field, has been successful at facilitating discussion and creating consensus on what resilience means in system safety via these detailed case studies, it tends to focus on smaller groups or systems not comparable to the size and geographic extent of regional infrastructure systems.

For critical infrastructure, questions about resilience typically exist at the scale of a city or an entire region, and the system as a whole often spans multiple organizations, both public and private. It is often difficult to obtain a comprehensive systems-level view of the physical components (for example, the pipes and wires, along with the assets that they connect), and even harder to understand the organizational and human factors involved in the daily operation of these systems. Thus, it is often in the context of disaster events that resilience (or the lack of it) is documented in detail. Each additional disaster sheds light on a new infrastructure sector or geographic region. For example, Flynn (2015) presents a comprehensive study about the impact of Superstorm Sandy on the transportation, energy and healthcare systems in the New York Metropolitan region. Efforts such as these focus temporary attention on the problem, but attention spans and memories are short. As noted by Flynn (2007), there is a need to create living case studies during and after these events, with the idea to let no disaster be wasted.

More commonly, we find case studies on infrastructure systems in which the objective of using real infrastructure data is to demonstrate the utility of some particular analysis

technique. For example, within engineering we typically find studies that are focused on an individual infrastructure sector, such as electric power or water (for example, Adachi and Ellingwood 2008; Ouyang and Dueñas-Osorio 2014), or a particular disruption type, such as seismic events (for example, Dueñas-Osorio et al. 2007; Poljanšek et al. 2012), and/or a particular geographic location (for example, Gómez-Baggethun et al. 2012).

Not surprisingly, these case studies are conducted from different disciplinary views and therefore are distributed across disparate journals and academic communities, making a comprehensive view of them difficult (barrier 1). This creates an opportunity for periodic literature reviews, but more broadly suggests the need for a systematic means of cataloging and/or disseminating them. Moreover, case studies provide an opportunity to build on the success in resilience engineering to create a more cohesive understanding of resilience (barrier 4) across interrelated disciplines.

Case studies are important for peer review and as a means of disseminating results within the research community, but they are potentially even more important for education and training. When case studies become a common touch point for peer review and training, they start to break down knowledge barriers (barrier 1) by creating common modes of communication across disparate disciplines. There is a recognized need to prepare the next generation of engineers, designers and policy makers to handle the challenges facing critical infrastructure systems, and developing appropriate curricula is a focus of several university-led research centers (CIRI 2017; GMU 2017; NEU 2017; NPS 2017). Although we do not yet know how to create resilience in an integrated manner, case studies provide a means to highlight individual success stories.

Case studies can also provide the opportunity to learn about the tensions and tradeoffs faced in the management and operation of critical infrastructure systems. For example, consider the possibility of education programs patterned after the 'case study method' championed at Harvard Business School (HBS 2017) and elsewhere. In this method, each individual class session is organized around a case that not only confronts students with a key challenge faced by an organization, but also places the student in the role of a decision maker who has to take action. Each case contains context-rich details (sometimes anonymized) that attempt to simulate the tensions at play, and through facilitated discussion the objective is to help the student understand the problem from multiple perspectives. Since case study learning is proven to work without need for specific knowledge of the original case (barrier 3), there exists an opportunity to create case studies that are based on real systems but sharable with broad stakeholders.

The ongoing development of modeling and simulation tools to support analysis and decision making for critical infrastructure systems further creates opportunity to supplement traditional case studies with software-based tools that could enhance student learning through experimentation. As discussed by Seager et al. (2017), experimentation is an important part of the learning process, and the incorporation of 'simulation games' into the classroom creates a potentially powerful vehicle for demonstrating infrastructure complexity and the conflicts that arise naturally between technological, social, and economic forces.

Current repositories of case studies are unavailable to the public (a consequence of barrier 3), but some concerted efforts demonstrate that when they become available for comparison across contexts, they can produce generalizable knowledge across infrastructure sectors, disruptions and geographic regions (barrier 2). For example, Bowman

(2016) considers 33 'case studies' in the form of reports generated by the Department of Homeland Security's Regional Resiliency Assessment Program (RRAP; see US DHS 2017c). He collects 3683 individual coded blocks of information from 4466 pages of material, and then uses qualitative analysis techniques to identify 'recurring empirical resilience gaps that may exist within and across lifeline critical infrastructure sectors and geographic regions' (Bowman 2016, p. 12) as well as to assess the extent to which these are caused by systematic barriers. At the time of his research, Bowman was one of only a handful of individuals (and the only true researcher) who had read all of the RRAP reports, in part owing to the sensitive nature of the reports themselves (often marked PCII). The lesson is that it is hard to generalize knowledge across such studies if nobody has access to them.

On balance, there is much more that can be achieved through the systematic use of case studies on critical infrastructure systems.

RECOMMENDATION 2: A REPOSITORY OF 'VIRTUAL, YET REALISTIC' INFRASTRUCTURE SYSTEM DATA

As discussed previously, data on real infrastructure systems is typically sensitive because it reveals information about customers (exposing proprietary data associated with competitive advantage and/or customer privacy) as well as vulnerabilities (potentially creating security threats). Nonetheless, researchers need data to support the development of models and analysis techniques. How do we proceed in the face of barrier 3?

One possibility is to use historical data that is no longer operationally relevant to real systems. Alderson et al. (2011) study Euler's classic 'seven bridges of Konigsburg' problem from a modern traffic perspective, while Alderson et al. (2013) consider the Soviet rail system *circa* 1955 to illustrate key points involving the 'most vital' components of an infrastructure system. Historical examples such as these can provide geographic realism that is useful in understanding salient features of infrastructure systems. Unfortunately, the availability of suitable historical studies is often limited.

Another option is to develop small, canonical 'toy' models that vary in complexity and realism but are almost always small in scale. For example, Morlok and Chang (2004) consider a small container transport system, while Alderson et al. (2015) consider a small fuel system. Such examples are useful for developing, testing and explaining an analysis technique, typically within a single discipline, but they often have limited explanatory power for understanding the resilience of real systems, particularly those in other disciplines (barrier 1).

There is another option that remains largely unexplored, namely, to invest in a repository of 'realistic, yet fictitious' data suitable for infrastructure modeling, analysis, benchmarking, and validation. In general, there is an ongoing need for a modeling and simulation environment for critical infrastructures that has both functional realism (that is, it has to behave like a real system) and geographic realism (that is, it has to look like a real place). Although there is a growing catalog of real US infrastructure assets collected by the government (for example, the Homeland Security Infrastructure Program; see https://www.dhs.gov/infrastructure-information-partnerships, accessed 6 August 2018), these are typically restricted in their availability to the research community as a whole.

The Center for Homeland Defense and Security (CHDS) at the Naval Postgraduate School (NPS) has developed a 'virtual place' called Dystopia (because it has been designed for disasters and undesirable events) for the purposes of supporting training, exercises and analyses. Many of the exercises and analyses conducted at NPS for student theses and funded projects require deep contextual environments in which to operate if they are to produce usable results, but real environments tend to be either unavailable or classified. Figure 5.2 presents a screenshot of the Dystopia interface (CHDS 2011).

Dystopia is a collection of geo-specific data that includes spatial information, two-dimensional map products, and a database of metadata about the people and places located in a fictional world. In its current state, Dystopia has all the things we would expect a growing metropolis to have: an international airport, a large port, an army base, a university (with a nuclear program), multiple cities separated by an international border, a diverse population with its own economy, and many other features useful in a wide variety of exercises and analyses. Dystopia is built in data layers so that data for specific purposes can be turned on or off as needed. The Dystopia website includes maps of infrastructure systems, including energy, water, transportation and communication assets; however, Dystopia does not currently include functional models of these infrastructures, nor is the infrastructure data itself consistent with the engineering reality of operational systems. Although the current infrastructure data in Dystopia is not sufficient to assess quantitatively, for example, the impact of a lost electric power substation or the destruction of a bridge, Dystopia has been carefully designed to allow for the addition of this data and/or the integration of other infrastructure models. See Alderson and Darken (2018) for a detailed discussion of the motivation, design principles and implementation of Dystopia.

Dystopia contains two key design features relevant to this discussion: (1) standardized data formats that make it possible to swap out Dystopia data layers (possibly for real infrastructure) with minimal impact, and (2) use of a standard test environment that makes it possible to share models and data sets across a diversity of partners for demonstration, validation and further exploration. Dystopia is not only sharable, but also modifiable, similar to open source software, except that it supports a repository of data. Thus, as new users add content, it becomes available to all future users. Realizing this vision requires content creation and management tools for end users. However, reuse of this environment has several key advantages, among them the sharing of validated scenarios and the sharing of validated datasets for comparing results across methods, processes and algorithms.

Currently, there have been only preliminary efforts to supplement the original Dystopia data with functional models of infrastructure operation. Martin (2014) presents a means for integrating a telecommunications system in Dystopia, and Ruether (2015) presents a simplified electric power and fuel system.

Ultimately, the goal would be to develop a repository of canonical infrastructure system models, with corresponding validated simulations of infrastructure functions, and verified data that can be used for benchmarking analysis of infrastructure behavior, as well as for testing novel designs. These models need to have the ability to conduct systematic 'what if' analysis on system response to disruptive events. By developing simulated interdependent infrastructure systems that are geographically and functionally realistic, it also becomes possible to generate 'stress test' scenarios for evaluating potential design changes to enhance the resilience of these systems at the local and regional levels. Moreover, these systems and related data can be used to develop case studies to support exercises, training and

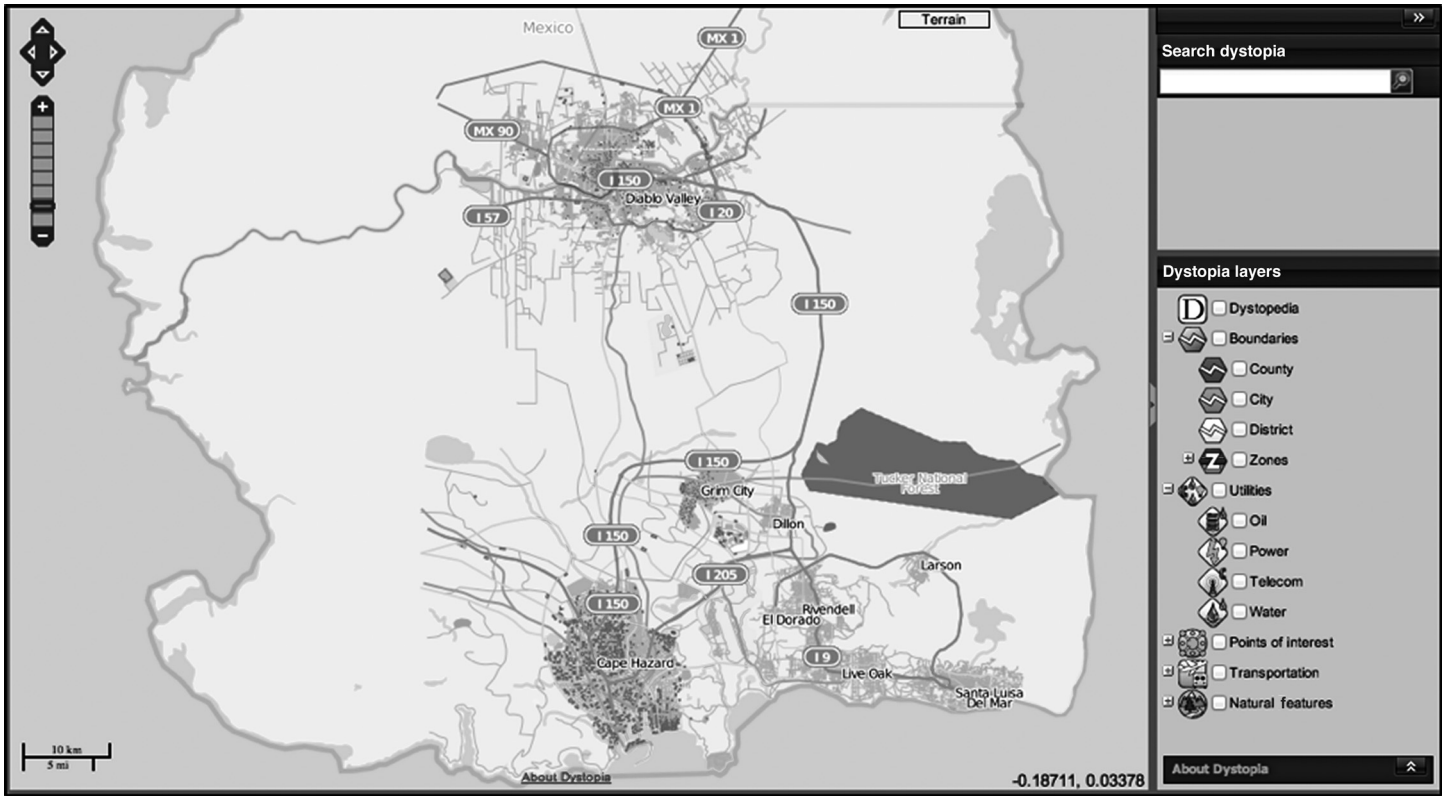

Note: Data about critical infrastructure is presented in layers (along the right) that can be turned on and off as needed. In its current form, the ability to exercise standard models of interdependent infrastructure function for consequence estimation, critical component analysis or mitigation does not yet exist. A platform that facilitates community development and sharing of infrastructure data and models would help to overcome the barriers currently restraining scientific understanding of resilience, among other things.

Source: Center for Homeland Defense and Security, Naval Postgraduate School.

Figure 5.2 Map-based interface to the 'realistic, yet fictitious' world of Dystopia

education, as described above. Recent efforts to develop the Centerville Virtual Community Testbed (Ellingwood et al. 2016) suggest a shared recognition for such an environment.

By building infrastructure models into an open-source, sharable and modifiable platform, we can enable broad studies that go beyond the limited, hazard-centered scope currently associated with majority of real-world studies (barrier 2). Previous attempts at releasing modifiable, free online platforms for the general public to participate in citizen science has produced knowledge in a variety of domains. Realistic, yet fictitious models of infrastructure can support the same broad-scale efforts to identify infrastructure vulnerabilities and risks previously unimagined by experts.

Dystopia serves as a provocative example for potential integration of models and data with visualizations to support their analysis, but to some extent it is merely yet another geospatial infrastructure modeling platform. Over the past decade, there has been considerable progress in the development of realistic, domain-specific network models of lifeline CI behavior by the government, at universities and at the national laboratories (for a review, see Yusta et al. 2011; Ouyang 2014). For example, the Environmental Protection Agency (EPA) has standard and openly available models for hydraulic and water quality behavior of water distribution piping systems (EPANET; US EPA 2017a) and storm water management (SWMM; US EPA 2017b). For electric power systems, there is an industry-standard 'direct current (DC) optimal power-flow model' that system operators use to optimize generation to meet demands (for example, Wood and Wollenberg 1996, pp. 108–11), standards for evaluating cascading failure behavior in electric power systems (IEEE Industry Applications Society, Power Systems Engineering Committee 1990), as well as models for analyzing them (for example, Papic et al. 2011; Portante et al. 2011). There are national models for fuel transportation (for example, the national transportation fuels model; Kelic et al. 2015) and models for natural gas transmission (for example, Portante et al. 2007). These are only a subset of the models that are available.

What remains missing is a platform on which these disparate models can be integrated in a reasonable manner. There are several ongoing efforts to develop models of interdependent infrastructure systems – for example, the Interdependent Networked Community Resilience Modeling Environment (IN-CORE) currently under development at the Center for Risk-Based Community Resilience Planning (NIST 2017) – but the complications for integrating these models often result in highly stylized interfaces that make it difficult for anyone beyond the development team to participate.

In order for such a platform to succeed it will be not only need to be sharable, but extensible. It is unlikely that any single research team will have the resources to sustain long-term development of any single platform. This implies the need for content creation and management tools for the community as a whole. Reuse of this content has several key advantages, among them the sharing of validated scenarios and the sharing of validated datasets for comparing results across methods, processes and algorithms. The use of standard data for benchmarking is common in other disciplines, for example, there are standard datasets for satisfiability problems in computer science (Gomes et al. 2008) and standard test networks for power distribution systems (IEEE Power & Energy Society Distribution System Analysis Subcommittee's Distribution Test Feeder Working Group 2017). Even modest investments in the development of standard, or at least commonly available, datasets for interdependent infrastructure systems have the potential for a large impact on the research community as a whole.

RECOMMENDATION 3: RE-THINKING THE MISSION OF CI PROTECTION AND RESILIENCE (AND IMPLICATIONS FOR ANALYSIS, TOOLS AND WORKFORCE DEVELOPMENT)

Modern infrastructure systems have evolved into a complicated and complex synthesis of interdependent technologies, governed by a mix of automation and human operation. These systems operate across potentially vast distances and extremely different timescales, ranging from fractions of a second (for example, telecommunications and electric power) to minutes (for example, highway traffic) to hours and days (for example, water) to months and years (design and construction). Our ability to design, construct and deploy these systems individually has outpaced our ability to understand, predict or control them collectively. This makes it particularly difficult to manage their function in the presence of natural hazards or deliberate threats, and exacerbates the challenge for organizations charged with ensuring national security and welfare.

Concerns on the part of federal and state governments to execute this mission have led to the development of 'watch floors' for situational awareness across CI sectors. For example, DHS operates the National Infrastructure Coordinating Center (NICC) as:

> the dedicated 24/7 coordination and information sharing operations center that maintains situational awareness of the nation's critical infrastructure for the federal government. When an incident or event affecting critical infrastructure occurs and requires coordination between the United States Department of Homeland Security and the owners and operators of our nation's infrastructure, the NICC serves as that information sharing hub to support the security and resilience of these vital assets. (US DHS 2017b)

Similar efforts to link cyber and physical security are exercised at the National Cybersecurity and Communications Integration Center (NCCIC; see US DHS 2017a).

Thus, the US federal government is attempting to maintain situational awareness about infrastructure function for the nation as a whole and preparing to respond in the event of large-scale disaster. Similar efforts are conducted at the state level; for example, the California Office of Emergency Services plays the equivalent role for the State of California, managing planning, preparation and response to disaster events as well as situational awareness for critical infrastructure systems (http://www.caloes.ca.gov/, accessed 6 August 2018). The primary reason for this is described in Bullock et al. (2011, p. 114):

> the private sector looks to the government for assistance when the threat at hand exceeds an enterprise's capability to protect itself beyond a reasonable level of additional investment. In this light, the federal government promises to collaborate with the private sector (and state and local governments) to ensure the protection of nationally critical infrastructures and assets; provide timely warning and ensure the protection of infrastructures and assets that face a specific imminent threat; and promote an environment in which the private sector can better carry out its specific protection responsibilities.

Beyond the many county- and city-level organizations tasked with local infrastructure planning and emergency response, there exist other types of organizations that have been created or self-organized to coordinate infrastructure-related decision making, typically at the regional level. For example, the Association of Bay Area Governments (ABAG; see http://www.abag.ca.gov/, accessed 6 August 2018) works to cooperate and collaborate

within the San Francisco Bay Area on issues ranging from economic planning to disaster response. Also, the Maritime Transportation Security Act of 2002 (MTSA; Public Law 107–295) requires the establishment of an Area Maritime Security Committee (AMSC) in each of the nation's ports to coordinate the activities of stakeholders – including commercial entities, the US Coast Guard, and other federal, state and local agencies – particularly as they pertain to using limited resources to prepare, prevent and respond to deliberate and non-deliberate security events.

The key recognition is that government agencies and other organizations have adopted protecting critical infrastructure systems as a mission, but what exactly is that mission and how does one measure its success? In military contexts, a mission is defined as 'The task, together with the purpose, that clearly indicates the action to be taken and the reason therefore' (US DOD 2001, p. 349). The key idea here is that a mission requires not only the 'what' but also the 'why', which is essential for measuring success.

Currently within the DOD and the DHS, the mission of protecting critical infrastructure is organized primarily around lists of infrastructure assets, often segregated into different tiers of importance, that are intended to reflect their relative criticality. Each organization has processes for nominating, vetting and selecting the infrastructure assets that are on each of these lists; these lists, in turn, dictate the processes that must be followed for assessing and mitigating the risk associated with these assets. Because these processes can get complicated, however, it is all too easy to focus exclusively on the assets (the 'what') and lose sight of the reasons for their criticality (the 'why').

The DOD has recently reorganized its efforts to protect defense-related critical infrastructure under a broader program of mission assurance (US DOD Directive 3020.40: Mission Assurance). In this context, the recognition is that an asset is important because of its contribution to a military capability (or function) that in turn supports one or more missions. That is, the focus needs to remain on the relationship between the infrastructure asset and the missions it supports.

This leads to an important question for civilian infrastructure: what exactly is the mission? As with the military, there ought to be an explicit connection between an infrastructure asset, the capability (or function) it creates, and the mission it supports. Otherwise there is a tendency to believe we need to watch everything, with the intent to prevent anything bad from happening. Watch floors and emergency response are perhaps necessary, but are unlikely to be sufficient for achieving infrastructure resilience. Instead, it is important that we measure not by the absence of 'bad' events, but by the presence of 'good' events, that is, mission success. This is consistent with recent changes in emphasis within the safety engineering community (see Hollnagel 2014).

Naturally, the definition of mission will likely differ among disparate disciplines and stakeholders (barrier 1). Consider water infrastructure as an example: a mission for ecologists may be watershed restoration and management, for a government practitioner it might be drinking water storage and access, for the military and/or emergency responders it might be flood control, and to another infrastructure system (for example, power grid) it might be reliability. Even though multiple missions potentially mean more confusion over how infrastructure is studied and discussed, recognizing infrastructure mission 'in practice' is an important step to how resilience should be understood and pursued (barrier 4). Safety engineering realizes that 'work as done' is not necessarily 'work as imagined', where even similar infrastructure service providers may interpret and implement the same

national policies differently in practice. Realizing that the imagined definition of mission, and consequentially resilience, by the infrastructure owners and operators may deviate from what is available publicly is important to clarify for researchers.

Infrastructure owners and operators tend to have a very clear notion of mission, typically focused on the safe and reliable delivery of a service (for example, electricity, clean water and communications), and government stakeholders would benefit from a more 'operational view' of critical infrastructure (Alderson et al. 2015). More relevant to this discussion, the scientific research community would benefit from a clearer notion of mission particularly as it pertains to national security and resilience, so that it can inform the types of analyses and tools that are needed by stakeholders to support decision making. Definitions of failure, consequences and threats may shift to more appropriate analyses and shared intent across the CI industry with a greater understanding of mission in infrastructure (barrier 2).

In its first decade, the DHS has sought tools that allow existing homeland security professionals, who typically have little or no expertise in infrastructure, to perform assessments of CI systems. This requirement puts extreme limitations on the types of tools that can be deployed and analyses that can be supported. Moreover, it is inconsistent with the norms in other industries. In medicine, we do not (yet) deploy tools that allow individuals without medical training to diagnose and treat human illness. Rather, there is a process by which we educate and then train doctors and nurses, who are subject to industry evaluation and must apprentice to recognized experts and be certified before they are allowed to practice. Similarly, within the financial sector, analysts who make buy and sell recommendations on company stocks typically specialize by industry sector, having apprenticed under the mentorship of a senior colleague before they are given any real responsibility. They must also pass industry certification before they are allowed to serve as investment advisors. Although automated medical and financial advice tools are undergoing rapid development, common wisdom suggests that it is not yet prudent to fully entrust our medical treatment of care or financial welfare to automated technologies. Interconnected infrastructure systems are starting to become as complex as living organisms, and despite ongoing progress in artificial intelligence (AI) it might be naive to think that the use of an automated tool is sufficient to qualify someone to analyze critical infrastructure.

Thus, for CI systems there is a fundamental connection between the types of analyses to be performed, the supporting tools, and the professional development of the workforce who will execute this. A starting point for unraveling all of this is a clearer definition of the ultimate mission to be performed.

CONCLUSION

During the past 20 years, the US has embraced the mission of critical infrastructure protection as a means of securing national welfare. However, critical infrastructure systems are complicated, often complex entities with behavior that can be hard to understand and predict, much less control. The interdisciplinary nature of these systems holds promise for greater understanding (Seager et al. 2017), but also creates practical barriers to collaboration across disciplines. Recent progress in our scientific understanding has been

hampered because research efforts have been overly focused on threat-based analysis, ultimately revealing vulnerabilities that must be protected and cannot be disseminated in a meaningful way. Moreover, we are still struggling to develop a theory of resilience, and have yet to discover best practices for building resilience in an integrated manner.

Nonetheless, there is a need to educate and train a new generation of professionals about critical infrastructures and resilience. Since the domain-specific details of these systems matter, the use of case studies and immersive role play, possibly coupled with simulation, are promising directions for ongoing development and investment. These case studies, either real or fictional, will require data that can be shared for model validation and benchmarking.

However, in order for these efforts to be successful, we require a clear sense of mission. Presidential Policy Directive 21 calls for proactive and coordinated efforts 'to strengthen and maintain secure, functioning, and resilient critical infrastructure'. It is perhaps time to revisit what exactly is the shared mission of the various stakeholders, including the research community.

This chapter considers only a subset of the open issues. Most notably absent in this discussion is the growing role and vulnerability associated with the cyber-physical nature of all critical infrastructure systems. Comparable efforts to understand resilience in the cyber domain will be will be of equal, or perhaps even greater, importance.

ACKNOWLEDGMENTS

The author is indebted to many colleagues for their ongoing conversations about resilience in critical infrastructure systems, most notably Jerry Brown, Matt Carlyle, Rudy Darken, Daniel Eisenberg, Steve Flynn, Tom Seager and David Woods. This work was supported by funding from the Office of Naval Research and the Defense Threat Reduction Agency.

REFERENCES

Adachi, T. and B.R. Ellingwood (2008), 'Serviceability of earthquake-damaged water systems: effects of electrical power availability and power backup systems on system vulnerability', *Reliability Engineering & System Safety*, **93** (1), 78–88.

Alderson, D.L. and R.P. Darken (2018), 'Dystopia: a virtual environment for education, training, and exercises', NPS Technical Report, Naval Postgraduate School, Monterey, CA.

Alderson, D.L. and J.C. Doyle (2010), 'Contrasting views of complexity and their implications for network-centric infrastructures', *IEEE Transactions on Systems, Man, and Cybernetics-Part A*, **40** (4), 839–52.

Alderson, D.L., G. Brown and W.M. Carlyle (2013), 'Sometimes there is no "most vital" arc', *Military Operations Research*, **18** (1), 21–37.

Alderson, D.L., G.G. Brown, W.M. Carlyle and R.K. Wood (2011), 'Solving defender-attacker-defender models for infrastructure defense', in R.K. Wood and R.F. Dell (eds), *Operations Research, Computing, and Homeland Defense*, Hanover, MD: Institute for Operations Research and the Management Sciences, pp. 28–49.

Alderson, D.L., G.G. Brown and W.M. Carlyle (2014), 'Assessing and improving operational resilience of critical infrastructures and other systems', in A. Newman and J. Leung (eds), *Tutorials in Operations Research: Bridging Data and Decision*, Hanover, MD: Institute for Operations Research and Management Science, pp. 180–215.

Alderson, D.L., G.G. Brown and W.M. Carlyle (2015), 'Operational models of infrastructure resilience', *Risk Analysis*, **35** (4), 562–86.

Alexander, D.E. (2013), 'Resilience and disaster risk reduction: an etymological journey', *Natural Hazards and Earth System Sciences*, **13** (11), 2707–16.
American Society of Civil Engineers (ASCE) (2017), '2017 Infrastructure Report Card', accessed 6 August 2018 at https://www.infrastructurereportcard.org/.
Bowman, R.E. (2016), 'Toward a more complete theory of disaster resilience as it relates to homeland security, in general, and critical infrastructure, in particular: better understanding the challenges we face', PhD dissertation, School of Public Policy and Urban Affairs Northeastern University, Boston, MA.
Brown, K. (2006), *Critical Path: A Brief History of Critical Infrastructure Protection in the United States*, Fairfax, VA: Spectrum.
Bullock, J., G. Haddow and D.P. Coppola (2011), *Introduction to Homeland Security: Principles of All-Hazards Risk Management*, Waltham, MA: Butterworth-Heinemann.
Center for Homeland Defense and Security (CHDS) (2011), 'Dystopia: where bad things happen', Naval Postgraduate School, accessed 2 July 2017 at https://www.chds.us/ed/items/291.
Cook, R.I. (2006), 'Being bumpable: consequences of resource saturation and near-saturation for cognitive demands on ICU practitioners', in D.D. Woods and E. Hollnagel (eds), *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*, Boca Raton, FL: Taylor & Francis and CRC Press, pp. 23–35.
Critical Infrastructure Resilience Institute (CIRI) (2017), accessed 21 May 2017 at http://ciri.illinois.edu/.
Cutter, S.L. (2016), 'The landscape of disaster resilience indicators in the USA', *Natural Hazards*, **80** (2), 741–58.
Dekker, S. (2014), *Safety Differently: Human Factors for a New Era*, Boca Raton, FL: CRC Press.
Dueñas-Osorio, L., J.I. Craig and B.J. Goodno (2007), 'Seismic response of critical interdependent networks', *Earthquake Engineering & Structural Dynamics*, **36** (2), 285–306.
Eisenberg, D.A. (2018), 'How to think about resilient infrastructure systems', PhD dissertation, Arizona State University, Tempe, AZ.
Ellingwood, B.R., H. Cutler, P. Gardoni, W.G. Peacock, J.W. van de Lindt and N. Wang (2016), 'The Centerville virtual community: a fully integrated decision model of interacting physical and social infrastructure systems', *Sustainable and Resilient Infrastructure*, **1** (3–4), 95–107.
Executive Order 13010, US Federal Register, Vol. 61, No. 138, July 17, 1996.
Flynn, S. (2007), *The Edge of Disaster: Rebuilding a Resilient Nation*, New York: Random House.
Flynn, S.E. (2015), 'Bolstering critical infrastructure resilience after Superstorm Sandy: lessons for New York and the nation', Global Resilience Institute, Northeastern University, Boston, MA, accessed 6 August 2018 at http://hdl.handle.net/2047/D20241717.
Folke, C. (2006), 'Resilience: the emergence of a perspective for social–ecological systems analyses', *Global Environmental Change*, **16** (3), 253–67.
George Mason University, Center for Infrastructure Protection & Homeland Security (GMU) (2017), accessed 2 July 2017 at https://cip.gmu.edu/.
Gomes, C.P., H. Kautz, A. Sabharwal and B. Selman (2008), 'Satisfiability solvers', *Foundations of Artificial Intelligence*, vol. 3, Amsterdam: Elsevier, pp. 89–134.
Gómez-Baggethun, E., V. Reyes-García, P. Olsson and C. Montes (2012), 'Traditional ecological knowledge and community resilience to environmental extremes: a case study in Doñana, SW Spain', *Global Environmental Change*, **22** (3), 640–50.
Haimes, Y.Y. (2009), 'On the definition of resilience in systems', *Risk Analysis*, **29** (4), 498–501.
Hale, A. and T. Heijer (2006), 'Defining resilience', in E. Hollnagel, D. Woods and N. Leveson (eds), *Resilience Engineering: Concepts and Precepts*, Aldershot: Ashgate, pp. 95–123.
Hart, S.D., J.L. Klosky, S. Katalenich, B. Spittka and E. Wright (2014), 'Infrastructure and the operational art: a handbook for understanding, visualizing, and describing infrastructure systems', US Army Corps of Engineers, Engineer Research and Development Center Report TR-14-14, September.
Harvard Business School (HBS) (2017), 'The HBS case method', accessed 2 July 2017 at http://www.hbs.edu/mba/academic-experience/Pages/the-hbs-case-method.aspx.
Holling, C. (1973), 'Resilience and stability of ecological systems', *Annual Review of Ecology and Systematics*, **4** (November), 1–23.
Hollnagel, E. (2014), *Safety-I and Safety-II: The Past and Future of Safety Management*, Farnham: Ashgate.
Hollnagel, E., D. Woods and N. Leveson (eds) (2006), *Resilience Engineering: Concepts and Precepts*, Aldershot: Ashgate.
Hollnagel, E., C.P. Nemeth and S.W.A. Dekker (eds) (2008), *Resilience Engineering Perspectives, Volume 1: Remaining Sensitive to the Possibility of Failure*, Aldershot: Ashgate.
Hollnagel, E., J. Pariès, D.D. Woods and J. Wreathall (eds) (2011), *Resilience Engineering Perspectives, Volume 3: Resilience Engineering in Practice*, Farnham: Ashgate.
Institute of Electrical and Electronics Engineers (IEEE) Industry Applications Society. Power Systems Engineering Committee (1990), 'IEEE Recommended Practice for Industrial and Commercial Power Systems Analysis: Approved May 31, 1990, IEEE Standards Board: Approved October 26, 1990', American National Standards Institute.

Institute of Electrical and Electronics Engineers (IEEE) Power & Energy Society Distribution System Analysis Subcommittee's Distribution Test Feeder Working Group (2017), 'Distribution test feeders', accessed 2 July 2017 at https://ewh.ieee.org/soc/pes/dsacom/testfeeders/.

Kaszynski, W. (2000), *The American Highway: The History and Culture of Roads in the United States*, Jefferson, NC: McFarland.

Kelic, A., D.H. Hart, T.F. Corbet and M.L. Wilson (2015), 'Refinery control systems in national transportation fuels modeling' (No. SAND2015-6951C), Sandia National Laboratories (SNL-NM), Albuquerque, NM.

Krishnan, A. (2009), 'What are academic disciplines? Some observations on the disciplinarity vs. interdisciplinarity debate', Economic and Social Research Council (ESRC) National Centre for Research Methods Working Paper No. 03/09, University of Southampton, January.

Longstaff, P.H., N.J. Armstrong, K. Perrin, W.M. Parker and M.A. Hidek (2010), 'Building resilient communities: a preliminary framework for assessment', *Homeland Security Affairs*, **6** (3), 1–23.

Longstaff, P.H., T.G. Koslowski and W. Geoghegan (2013), 'Translating resilience: a framework to enhance communication and implementation', paper presented at the Fifth Resilience Engineering International Symposium, Soesterberg, The Netherlands, 25–27 June, pp. 12–23.

Madni, A. and S. Jackson (2009), 'Towards a conceptual framework for resilience engineering', *IEEE Systems Journal*, **3** (2), 181–91.

Martin, K.M. (2014), 'A geographic and functional network flow analysis tool', M.Sc thesis, Naval Postgraduate School, Monterey, CA.

Martin-Breen, P. and J.M. Anderies (2011), 'Resilience: a literature review', Rockefeller Foundation, New York.

Meerow, S. and J.P. Newell (2015), 'Resilience and complexity: a bibliometric review and prospects for industrial ecology', *Journal of Industrial Ecology*, **19** (2), 236–51.

Meerow, S., J.P. Newell and M. Stults (2016), 'Defining urban resilience: a review', *Landscape and Urban Planning*, **147**, 38–49.

Milly, P.C.D., J. Betancourt, M. Falkenmark, R.M. Hirsch, Z.W. Kundzewicz, D.P. Lettenmaier, et al. (2008), 'Stationarity is dead: whither water management?' *Science*, **319** (1), 573–4.

Morlok, E.K. and D.J. Chang (2004), 'Measuring capacity flexibility of a transportation system', *Transportation Res. Part A*, **38** (6) 405–20.

Moteff, J.D. (2015), 'Critical infrastructures: background, policy, and implementation', Report No, RL30153, Congressional Research Service, Washington, DC.

National Institute for Standards and Technology (NIST) (2017), A NIST-funded Center of Excellence: Center for Risk-Based Community Resilience Planning: A, accessed 21 May 2017 at http://resilience.colostate.edu/.

National Research Council (NRC) (2006), *Facing Hazards and Disasters: Understanding Human Dimensions*, Washington, DC: National Academies Press.

National Research Council (NRC) (2012), *Disaster Resilience: A National Imperative*, Washington, DC: National Academies Press, accessed 6 August 2018 at https://doi.org/10.17226/13457.

National Science Foundation (NSF) (2017), 'Program Solicitation NSF 16-618: Critical Resilient Interdependent Infrastructure Systems and Processes FY17 (CRISP)', accessed 21 May 2017 at https://www.nsf.gov/pubs/2016/nsf16618/nsf16618.htm.

Naval Postgraduate School (NPS) (2017), Center for Homeland Defense and Security, accessed 2 July 2017 http://www.chds.us/.

Nemeth, C.P. and E. Hollnagel (2014), *Resilience Engineering in Practise, Volume II: Becoming Resilient*, Farnham: Ashgate.

Nemeth, C.P., E. Hollnagel and S.W.A. Dekker (eds) (2009), *Resilience Engineering Perspectives, Volume 2: Preparation and Restoration*, Farnham: Ashgate.

Northeastern University (NEU) (2017), 'Master of Science in Security and Resilience Studies program', accessed 2 July 2017 at https://www.northeastern.edu/graduate/program/master-of-science-in-security-and-resilience-studies-14294/.

Ouyang, M. (2014), 'Review on modeling and simulation of interdependent critical infrastructure systems', *Reliability Engineering & System Safety*, **121** (January), 43–60.

Ouyang, M. and L. Dueñas-Osorio (2014), 'Multi-dimensional hurricane resilience assessment of electric power systems', *Structural Safety*, **48** (May), 15–24.

Papic, M., K. Bell, Y. Chen, I. Dobson, L. Fonte, E. Haq, et al. (2011), 'Survey of tools for risk assessment of cascading outages', in Proceedings of the 2011 IEEE Power and Energy Society General Meeting, Detroit, MI, 24–28 July, pp. 1–9.

Park, J., T. Seager, P. Rao, M. Convertino and I. Linkov (2013), 'Integrating risk and resilience approaches to catastrophe management in engineering systems', *Risk Analysis*, **33** (2), 356–67.

Poljanšek, K., F. Bono and E. Gutiérrez (2012), 'Seismic risk assessment of interdependent critical infrastructure systems: the case of European gas and electricity networks', *Earthquake Engineering & Structural Dynamics*, **41** (1), 61–79.

Portante, E.C., B.A. Craig and S.M. Folga (2007), December. 'NGfast: a simulation model for rapid assessment

of impacts of natural gas pipeline breaks and flow reductions at US state borders and import points', in *Proceedings of the 39th Conference on Winter Simulation*, Washington, DC: IEEE Press, pp. 1118–26.
Portante, E.C., B.A. Craig, L. Talaber Malone, J. Kavicky, S.F. Folga and S. Cedres (2011), 'EPfast: a model for simulating uncontrolled islanding in large power systems', in S. Jain, R. Creasey, J. Himmelspach, K.P. White, and M.C. Fu (eds), *Proceedings of the Winter Simulation Conference (WSC'11)*, Cantonsville, MD: Institute for Operations Research and the Management Sciences, pp. 1763–74.
President's Commission on Critical Infrastructure Protection (PCCIP) (1997), 'Critical foundations', technical report, The White House, Washington, DC.
Rockefeller Foundation (2017), '100 resilient cities', accessed 6 August 2018 at http://www.100resilientcities.org/.
Rose, A. (2004), 'Defining and measuring economic resilience to disasters', *Disaster Prevention and Management*, **13** (4), 307–14.
Rose, A. (2007), 'Economic resilience to natural and man-made disasters: multidisciplinary origins and contextual dimensions', *Environmental Hazards*, **7** (4), 383–98.
Rose, A. (2009), 'Economic resilience to disasters', Technical Report CARRI Research Report 8, Community and Regional Resilience Institute, Oak Ridge, TN.
Rose, A. (2017), *Defining and Measuring Economic Resilience from a Societal, Environmental and Security Perspective*, Singapore: Springer.
Rouse, W.B. (1982), 'On models and modelers: N cultures', IEEE Transactions on Systems, Man, and Cybernetics, **12** (5), 605–10.
Ruether, J.P.H. (2015), 'A virtual environment for resilient infrastructure modeling and design,' MSc thesis, Naval Postgraduate School, Monterey, CA.
Seager, T.P., S. Spierre Clark, D.A. Eisenberg, J.E. Thomas, M.M. Hinrichs, R. Kofron, et al. (2017), 'Redesigning resilient infrastructure research', in I. Linkov and J. Palma Olivera (eds), *Resilience and Risk: Methods and Application in Environment, Cyber and Social Domains*, Dordrecht: Springer, pp. 81–119.
The White House (2013), 'Presidential policy directive: critical infrastructure security and resilience', The White House, Washington, DC, accessed 20 August 2018 at https://www.dhs.gov/sites/default/files/publications/PPD-21-Critical-Infrastructure-and-Resilience-508.pdf.
Thoma, K., B. Scharte, D. Hiller and T. Leismann (2016), 'Resilience engineering as part of security research: definitions, concepts and science approaches', *European Journal for Security Research*, **1** (1), 3–19.
United States Department of Defense Multidisciplinary University Research Initiative (US DOD MURI) (2017), Fiscal Year (FY) 2018 Department of Defense Multidisciplinary Research Program of the University Research Initiative, Funding Opportunity Number N00014-17-S-F006, accessed 2 July 2017 at https://www.grants.gov/view-opportunity.html?oppId=292607.
United States Department of Homeland Security (US DHS), 'DOD Directive 3020.40: Mission Assurance', 29 November 2016, accessed 20 August 2018 at http://www.dtic.mil/whs/directives.
United States Department of Defense (US DOD), 'Joint Publication 1-02: Department of Defense Dictionary of Military and Associated Terms, 12 April 2001, (As Amended Through 31 October 2009)', accessed 6 August 2018 at http://jitc.fhu.disa.mil/jitc_dri/pdfs/jp1_02.pdf.
United States Department of Homeland Security (US DHS) (2017a), National Cybersecurity and Communications Integration Center, accessed 2 July 2017 at https://www.dhs.gov/national-cybersecurity-and-communications-integration-center.
United States Department of Homeland Security (US DHS) (2017b), National Infrastructure Coordinating Center, accessed 2 July 2017 at https://www.dhs.gov/national-infrastructure-coordinating-center.
United States Department of Homeland Security (US DHS) (2017c), 'Regional Resilience Assessment Program', accessed 2 July 2017 at https://www.dhs.gov/regional-resiliency-assessment-program.
United States Environmental Protection Agency (US EPA) (2017a), 'EPANET: software that models the hydraulic and water quality behavior of water distribution piping systems', accessed 4 July 2017 at https://www.epa.gov/water-research/epanet.
United States Environmental Protection Agency (US EPA) (2017b), 'Storm water management model (SWMM)', accessed 4 July 2017 at https://www.epa.gov/water-research/storm-water-management-model-swmm.
United States Government Accountability Office (US GAO) (2006), 'Critical infrastructure protection: Progress Coordinating Government and Private Sector Efforts Varies by Sectors' Characteristics', Report GAO-07-39, 16 October, Government Accountability Office, Washington, DC.
Westrum, R. (2006), 'A typology of resilience situations', in E. Hollnagel, D. Woods, N. Leveson (eds), *Resilience Engineering: Concepts and Precepts*, Aldershot: Ashgate Press, pp. 49–60.
Wood, A.J. and B.F. Wollenberg (1996), *Power Generation, Operation and Control*, 2nd edn, New York: Wiley.
Woods, D. (2006), 'Essential characteristics of resilience', in E. Hollnagel, D. Woods and N. Leveson (eds), *Resilience Engineering: Concepts and Precepts*, Aldershot: Ashgate, pp. 49–60.
Woods, D.D. (2005), 'Creating foresight: lessons for resilience from Columbia', in W.H. Starbuck and M. Farjoun (eds), *Organization at the Limit: NASA and the Columbia Disaster*, Malden, MA: Blackwell, pp. 289–308.
Woods, D.D. (2009), 'Escaping failures of foresight', *Safety Science*, **47** (4), 498–501.

Woods, D.D. (2015), 'Four concepts for resilience and the implications for the future of resilience engineering', *Reliability Engineering and System Safety*, **141** (April), 5–9.
Woods, D.D. and M. Branlat (2011), 'Basic patterns in how adaptive systems fail', in E. Hollnagel, J. Pariès, D.D. Woods and J. Wreathall (eds), *Resilience Engineering Perspectives Volume 3: Resilience Engineering in Practice*, Farnham: Ashgate, pp. 127–44.
Woods, D.D. and R.I. Cook (2006), 'Incidents: are they markers of resilience or brittleness?', in E. Hollnagel, D.D. Woods and N. Leveson (eds), *Resilience Engineering: Concepts and Precepts*, Aldershot: Ashgate, pp. 69–76.
Yusta, J.M., G.J. Correa and R. Lacal-Arántegui (2011), 'Methodologies and applications for critical infrastructure protection: state-of-the-art', *Energy Policy*, **39** (10), 6100–119.
Zeihan, P. (2014), *The Accidental Superpower: The Next Generation of American Preeminence and the Coming Global Disorder*, New York: Twelve.
Zolli, A. and A. Healy (2012), *Resilience: Why Things Bounce Back*, New York: Free Press.