Strategic Perspectives

Progress toward Resilient Infrastructures: Are we falling behind the pace of events and changing threats?

David D. Woods[1] and David L. Alderson[2]

[1] Professor Emeritus, Dept of Integrated Systems Engineering, Ohio State University, woods.2@osu.edu

[2] Professor, Operations Research Dept, Naval Postgraduate School, dlalders@nps.edu

[*see Author Capsule Bios below*]

Abstract

The current strategy for achieving resilient infrastructures is making progress too slowly to keep up with the pace of change as evidenced by a continuing stream of "shock" events. How do we better anticipate changing threats and recognize emerging new vulnerabilities in an increasingly interconnected world? We are facing a Strategic Agility Gap that requires us to revise our current perspective and processes if we are to make meaningful progress.

Introduction: The Critical Infrastructure Challenge

For over a decade, societies across the globe have realized the need to make critical infrastructure systems more resilient in the face of threats from natural disasters, risks from technological change, and adversarial actions (e.g., Flynn 2007). In the opening pages of this journal, Krieg (2020, p. 1) restated the challenge of resilient infrastructure: Critical infrastructures "are highly complex, interconnected and sometimes unplanned—and they are evolving at exponential rates. Impairment in one sector can cascade into multiple sector shutdowns leading to serious societal consequences. Each sector encompasses an array of physical assets, organizations and people as well as important cyberspace components. These factors can present unforeseen built-in vulnerabilities, and accidents are likely to be experienced as systems become more complex, opaque and interactive."

In response to this challenge, there has been major investment over more than a decade on the part of universities, national laboratories, and funding agencies in *modeling and simulation* as a means

 doi: 10.18278/jcip.2.2.2

- to describe and predict system behavior within lifeline infrastructures so as to find vulnerabilities (holes);
- for identifying dependencies across infrastructure sectors to project lines of propagation of events to minimize loss of valued services when infrastructures are threatened; and
- to fill holes and/or block propagation to contain consequences.

The basic strategy is invest to build up modeling tools that represent specific infrastructures in specific jurisdictions to capture the interconnections across these lifeline infrastructures. The investment in modeling tools tailored for lifeline infrastructures is intended to provide a base—a modeling "infrastructure" so to speak—that can be used to address new cases or re-examine changes to previous cases. Using the base of modeling capabilities to analyze infrastructures will support identification, assessment, and prioritization of vulnerabilities and associated consequences so stakeholders can the select most important ones for mitigation given limited resources. Continued investment will expand the capability for modeling and, over time, will lead to sufficiently detailed and thorough models of critical lifeline infrastructures and interconnections to support timely, practical decision making. These decisions can be assessed in terms of improved robustness relative to the vulnerabilities identified in the modeling and simulation runs. The modeling and simulation over different types of failure and attack will show which interventions produce measurable gains in either how well systems can withstand threats or how quickly systems will recover normal services.

However, the current strategy has scientific, technical, and practical limits that are revealed when we look at the continuing stream of disruptive events given the growth of complexities (Carlson and Doyle, 2000; Alderson and Doyle, 2010). These limits make the current strategy less responsive than events demand, leaving organizations stranded in the Strategic Agility Gap (Woods, 2020; Figure 1). For example, the current strategy rests on the assumption that modeling can uncover holes and map interdependencies rather completely—whereas work on the fundamentals that give rise to complexity penalties has revealed that, inevitably, past models will miss important aspects as processes of change, improvement, and adaptation continue. One of the new findings is that resilience depends on timely model updating and revision (Woods, 2018).

Falling Behind: The Strategic Agility Gap

Since the critical infrastructure challenge was recognized, the world has not stood still—growth occurred as new technology and opportunities arose, complexities grew, and new threats emerged as other parties hijacked valued capabilities for their own purposes. The pace of change continues to accelerate leading to expand-

ed scales of operation, dramatic new capabilities, extensive and hidden interdependencies, intensified pressures, new vulnerabilities, and puzzling failures with far reaching consequences. Society and organizations face the challenge of how to adapt to the increasing pace of change, scale, capability, risk, and threats, all in more complex, interconnected worlds.

Experience across industries, regions, and societies indicates organizations are *slow and stale* to adapt to new threats, as well as to seize new opportunities to build resilience. The result is a *Strategic Agility Gap* evidenced by the regular occurrence of surprising failures at organizational, regional, national scales—breakdowns that trigger or threaten widespread service outages with large financial and human costs (Woods, 2020). The Strategic Agility Gap is the difference between the *rate* at which an organization can *adapt to change* and the rise of new unexpected challenges at a larger industry or society scale. It is a *mismatch in velocities of change and velocities of adaptation* (Figure 1). Can organizations learn how to offset changing risks before failures occur as growth continues? Can organizations build capabilities to be poised to adapt to keep pace with and stay ahead of the trajectory of growing complexity and the penalties that arise as a result?

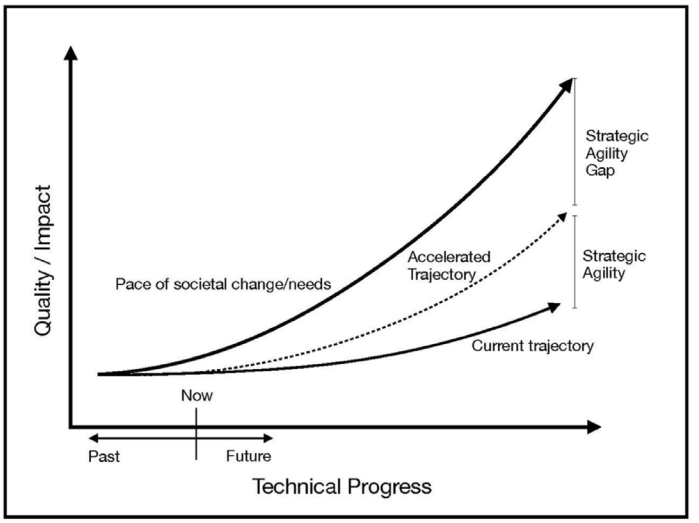

Figure 1. The Strategic Agility Gap. Reproduced from Woods (2020)

Resilience is a verb in the future tense

Starting as early as 2000, the idea emerged that complex systems, as distributed, layered networks with extensive tangles of interdependencies, are fundamentally brittle. But adaptive systems possess capabilities for resilience, as a basic adaptive capacity, to offset the complexity penalties that arise as growth and change go on (Hollnagel et al., 2006; Alderson and Doyle, 2010; Woods, 2019).

The basic signature of complexity penalties is the surprising sudden collapse of function against backdrop of continuous improvement and injection of new capabilities. Carlson and Doyle captured the basic finding as (2000, p. 2529): systems/layered networks "which are robust to perturbations they were *designed to handle*, yet fragile to *unexpected perturbations* and *design flaws*." The core result is that the pursuit of new capabilities under pressure to achieve new levels of productivity, efficiencies, and financial returns inevitably increases the risk of brittle collapse. Worse yet, each case of brittle collapse tends to be cryptic, in the sense that ambiguities about the multiple contributors that produced the breakdown make learning difficult, given the system has a record of past success and improvement. The difficulties are magnified because learning requires revision of past models based on a new understanding of the emergent properties that cut across critical infrastructures and all of the interacting systems, roles, and organizations.

Fortunately, the initial work on resilience also began to derive lessons from biological, human, and human-technology systems that have adapted to flourish in the face of the risk of brittle collapse. The lessons emphasized how these successful systems were *poised to adapt* to keep pace with change. The new work shifts the focus from why a failure occurred. Instead, adaptive capacity was revealed by looking at the challenges that occurred (much more often than stakeholders realized), but were managed by resilient performances (e.g., how does emergency medicine adapt to handle beyond-surge-capacity events successfully despite the threat of overload and bottlenecks).

The critical definition turned out to be "what is adaptive capacity?":

> Definition: *Adaptive capacity* is a system's readiness or potential to change how a system currently works—its models, plans, processes, behaviors, relationships—to continue to fit changing situations, anomalies, and surprises.

All adaptive systems, at all scales, possess the capacity to stretch or extend performance when events challenge their normal competence for handling situations. Without this capability for extensibility, brittle collapse would occur much more often than it is observed (Woods, 2018; Sharkey et al., 2020).

Three of the critical capabilities that support adaptive capacity as extensibility are:

- the ability to *revise* previous models and methods to recognize emerging new vulnerabilities as interconnections change;
- the ability to *synchronize* activities over multiple roles and layers of a network to scale responses to the scope of challenges; and
- the ability to *anticipate* challenges ahead to recognize of emerging new challenges, vulnerabilities and threats before capabilities are overloaded or oversubscribed.

Note these capabilities are *verbs*—actions that contribute to resilient performances—and not nouns or states. Even more surprising, adaptive capacity is future oriented. Adaptive capacity consists of readiness to respond and readiness to revise in advance of events that challenge the system's design. These properties are quite evident in studies of resilient performance in emergency medicine when challenged by beyond-surge-capacity events (Chuang et al., 2021).

Hence, the phrase "resilience is a verb in the future tense" provides a compact distillation of the findings. Thinking about the critical infrastructure challenge in terms of these new science results leads to new modeling approaches that complement the current strategy (e.g., Sharkey et al., 2020). Thinking of the capacity for resilient performance in this way represents a shift in how we can meet the goal of resilient critical infrastructures in a complex, changing, limited resource world.

Surprising failures and service outages continue to be regular occurrences

The continuing stream of shock events should remind stakeholders that they do not possess the adaptive capacity to match the trajectory of change and challenge as shown in Figure 1. As a result, the goal of resilient infrastructure remains an aspiration for the future.

New advances on the fundamental science of adaptation and complexity in distributed, tangled layered networks in the last few years have revealed basic patterns, findings, and laws that transcend particular triggering events, lines of interdependencies, and physical parts of infrastructure systems that are involved in any particular failure.

These general patterns are evident in multiple shock events that occurred in 2021. The shock events reveal general constraints that apply to future events as well. We refer to these events as shocks not because of the consequence, but rather because the events reveal gaps and holes in our models of how these systems work, the threats they face and the necessary countermeasures.

In challenge events

a. there is an expanding, but previously hidden, tangle of interdependencies that cross infrastructures and impact valued services;

b. the triggering event produces *effects at a distance* as impacts spread over the tangled lines of dependencies;

c. the effects increase tempos of activity over a wide set of roles across diverse organizations, challenging all to up their pace of activity and coordinate these activities in synchrony with other roles and levels;

d. the scale of effects expands across more organizational and jurisdictional boundaries with regional, societal wide, and even global reverberations;

e. the growing disruptions that flow from the triggering event put greater pressures on each stakeholder and increase the challenges each faces for their own scope of responsibility and for how their activity within their scope supports or hinders other related stakeholder activities; and

f. forces for fragmentation come to the fore and undermine the ability to synchronize actions over layers and roles to scale responses to match spreading disruptions and challenges.

These and many other kinds of emergent system patterns arise regardless of the whether the trigger/driver of the shock comes from extreme weather events (and longer-term climate volatility), the fragilities that accompany growth of capabilities, or new paths for adversarial conflict.

For example, extreme weather triggered the Texas energy crisis in February 2021 with at a minimum 151 deaths (with some estimates running several hundred higher) and at least $200 billion in financial loses. A previous widespread cold snap in 2010 foreshadowed the 2021 crisis, but lessons from the precursor energy system breakdown were weak and led to little or ineffectual remedies (large scale and duration cold snaps did not produce system wide energy system failures prior to 2010). Changes to economic incentive structures in pursuit of efficiencies based on models of deregulation and decentralized markets made the energy system remarkably brittle as spreading disruptions undermined the ability of other resources to come to bear to mitigate the consequences. The economic losses spread beyond Texas to affect ratepayers in distant states. Changes in the energy mix since the last extreme cold weather event had strong implications for how the system would respond to the next event, but modeling did not provide timely reassessments to recognize the potential for new risks and associated costs.

The framework of modeling did not include the general phenomenon of brittle collapse and any notions of resilience were primitive, disconnected from the growing base of general findings about complexity and adaptation. Instead, the distributed, layered network of human and organizational roles struggled to react to the spreading disruptions in local and fragmented ways compounding the consequences. As after the previous cold snap induced energy crisis in 2010, stakeholders continue to struggle to learn and implement systemic changes to build robustness and resilience.

Digital infrastructure—that underpins many valued services including all other lifeline infrastructures—represent another important example. Consider just three of the outages of valued digital services that occurred in 2021: the Fastly outage on June 8, 2021; a Facebook outage on October 4, 2021; and the AWS (Am-

azon) cloud services outage on December 7, 2021. These events demonstrate how the growth of capabilities produces hidden interdependencies, expands the scale of disruptions, and increases the difficulties of diagnosis and mitigation. In the AWS outage, computer engineers were hampered by a common problem, namely the tools used to diagnosis problems in the software infrastructure also are software-based and were degraded by the very outage the engineers needed to diagnose via these tools (Woods and Allspaw, 2019). The engineers in this case, and in ones that exhibit the same general pattern, had to develop ways to adapt around the bottleneck while disruptions spread and potential for consequences grew.

These digital infrastructure outages reveal important lessons. In all three of these outages, interdependencies occurred over multiple levels of software services across multiple organizational boundaries. Currently, these interactions are very difficult to recognize until loss of service is threatened. For example, Fastly is one major supplier of a specialized software service (CDN or Content Delivery Network) that operates invisibly in the background to improve responsiveness and load management. This service is valuable to nearly all internet-facing companies. One surprise was how a small unremarkable change by one of their customers revealed a dependency that disrupted their geographically distributed fleet of servers. Rebooting the fleet to restore the functionality meant customers' servers would try to repopulate their data *all at the same time* to regain the valued functionality.

In software engineering this is a "cache stampede" problem which degraded widespread services across multiple organizations. The problem in this case/infrastructure is an instance of a more general pattern in complex systems where common resources become oversubscribed when multiple parties begin to respond to disruptions in ways that overload some resources. Note another general pattern in this event: the story of the outage involved a surprisingly large number of organizations (a) in the genesis of the disruption, (b) those engaged in responding to spreading disruptions, and (c) those who had to cope with widespread secondary service disruptions. Also, note that two of these events began with leading service providers who operate with significant technical, financial, and expert human resources. These events are reminders—no organization is immune from complexity penalties.

One reaction to these incidents is: "but these are not the kind of 'critical' services we refer to when addressing 'critical infrastructures.'" However, the patterns are the same, such as "effects at a distance" which make diagnosis more difficult. In addition, losing valued services is critical to those who depend on them. For example, the Facebook outage affected WhatsApp, but in some parts of the world many critical human and business activities have adapted to take advantage of WhatsApp capabilities and to work around various local constraints. As a result, an outage that undermines this service disrupts a wide range of societal activities

in these regions for a time. This pattern is observed often: new capabilities start as non-essential improvements, but as they provide value, the capabilities migrate to undergird primary activities until they become essential to those activities. Inevitably, as new capabilities provide services valuable to human roles, people adapt to take advantage of that value so that outages become crises or even safety threats.

In addition, all of the patterns evident in the three software outages above appear in recent adversarial threats such as Colonial Pipeline (May 7, 2021), emergence of new tactics in ransomware, Solar Winds (2020), and the Log4j Shell server vulnerability (late 2021). There are interesting overlaps of the Log4j Shell server vulnerability story with non-adversarial software incidents (Allspaw and Cook, 2018).

The trajectory of ransomware threats over time illustrates how our capabilities also need to adapt over time to match the pace of change. The emergence of cryptocurrencies triggered significant growth of ransomware threats which have continued to evolve tactics, frequency, and scale of impact. The continuing story of ransomware is but one illustration of the need to recognize and redress the strategic agility gap.

Finally, a combination of the three classes of triggering events is eminently possible. Adversarial intrusions can hijack or disable services needed when extreme weather events increase criticality and degrade services themselves. Again, the patterns are generic and not simply about the details of any particular trigger event, targeting any particular set of infrastructure systems.

The point of reviewing a sample of visible recent outages isn't about the outages that resulted. From a resilience perspective, the information needed to build and sustain adaptive capacity does not lie in *why* a particular outage happened, but rather, in how other challenges have been occurring yet are dealt with successfully before visible outages resulted (Woods and Allspaw, 2019). Ironically, resilience as a verb in the future tense is most visible in the incidents that do not progress to tangible failures with losses for stakeholders (Hollnagel et al., 2006).

Are We Doomed to be Stuck in the Gap?

The recurrence of increasingly disruptive events in these systems, despite recent investments to avoid or mitigate them, provides evidence of technical and practical limits to the current strategy for building resilient infrastructure.

The most obvious problem is one of scale. The current strategy of building and deploying tools to model specific systems in specific jurisdictions is never going to address the entire problem. There are too many systems and too many interdependencies that continue to evolve. *We are never going to have a complete view of critical infrastructure such that we can close all the holes; such a vantage point does not exist.* Moreover, the current piecemeal approach prevents sharing

details and/or insights about our infrastructure systems (often for valid security reasons) that are slowing our collective learning (Alderson, 2019).

Modeling individual critical infrastructures remains necessary and important because it forces a closer look at how systems actually work, which is often misunderstood. Operational models of specific infrastructures are incredibly helpful as decision support tools (Alderson et al. 2015). Many modern critical infrastructure systems, such as the electric grid, are so complicated that it is near impossible to operate them without the support of modeling and simulation tools. But it is important to recognize "the model" of any particular infrastructure system is never complete because the system is not static. Rather, the system continues to evolve as new opportunities are explored and new vulnerabilities exploited.

Confident that the complex world will continue to throw novel challenges at us, we need to expand our current perspective and processes if we are really going to build resilient infrastructure.

Pivoting to Increase Strategic Agility

New fundamental findings provide opportunities to pivot from the current approach, enhance our tempo of progress, and reduce the strategic agility gap. These new insights have revealed general patterns and laws about complexity and adaptation (e.g., Chiang et al., 2007; Woods, 2018; Nakahira et al., 2021). These lawful patterns play out in particular settings as new capabilities are deployed, as processes of adaptation transform these improvements in unexpected ways, and as competitors/adversaries hijack new capabilities to pursue their own ends. In this paper we have illustrated just a small portion of these patterns using recent infrastructure outages arising from external events (extreme weather), fragilities that accompany growth (complexity penalties) and adversarial intrusions to degrade valued services.

One of the fundamentals about adaptive capacity is the ability to revise—building our models and tools in ways that make updating, revising, reframing and reconceptualizing straightforward. This requires an ability to recognize the significance of early, often "weak" signals that change is afoot in ways that might matter. Rather than discount such signals, resilient performance depends on seeing discrepancies as unexpected anomalies that trigger re-examination of previous beliefs. For example, the most common statement in after-action reviews of incidents that threaten loss of service in critical digital infrastructure is "I didn't know it worked that way" (e.g., Woods, 2017).

When organizations adopt processes to discover hidden interdependencies—rather than assume all are well-mapped and well-guarded—learning, updating, and revising models of how a system works is facilitated. The organization introduces probes, generally in non-risky periods or situations, in order to chal-

lenge and update its understanding of the challenges it faces and the processes it uses to handle challenges. This approach has become widespread in critical digital services as part of continuous development and deployment processes (Rosenthal and Jones, 2020).

In one common technique (which software reliability engineers call chaos engineering), system operators conduct frequent experiments by injecting small disruptions into their system as a means to uncover/discover how it actually responds to various disruptions. As complexity increases, this technique quicky shows how the system is more vulnerable than previously understood. The insights guide interventions that increase system robustness to the specific previously 'hidden' vulnerabilities. The probes also provide insights about what factors contribute to resilient performances in many different kinds of anomalous situations could occur that were not part of the test event (Allspaw and Cook, 2018; Woods and Allspaw, 2019). Techniques like this enhance the ability of the entire cyber-physical-human system to adapt to novel challenges—a boost in adaptive capacity and resilient performance.

Underlying these efforts is a recognition that the abstract capabilities of the organization are more important than the physical assets (Finkel, 2011). In addition, emergency response plans, no matter the efforts at pre-planning, cannot be sufficiently robust in themselves when events call them into action (this has turned out to be a fundamental constraint; e.g., Alderson and Doyle, 2010; Woods, 2018; Rosenthal & Jones, 2020). Thus, it is important to study the way organizations can mobilize and generate additional adaptive capacity when non-routine events threaten to overload (saturate) initially deployable responses as breakdowns spiral into accidents or disasters (Mendonça at al., 2007; Chuang et al., 2021). Recent studies investigating the resilience of organizations in response to 2011 Superstorm Sandy (Zhang and Mendonça, 2021) and the 2015 refugee influx to Sweden (Degerman, 2021) are promising steps in the right direction. Moreover, there is a need to continue investment in training and exercises for developing adaptive capacity at upper echelons of organizations and how these layers synchronize with other organizational layers (e.g., Bergström et al., 2011; Alderson et al., 2022).

Conclusion

In this paper we pose a simple question: despite investments and progress in the development of critical infrastructure systems, *are societies, especially our own, falling behind the pace of events and changing threats?* The current strategy to develop resilient infrastructure, while a necessary part of the portfolio for progress, has been running into scientific, technical, and practical limits leaving progress stalled in the strategic agility gap. Increasing the tempo of progress requires a strategic course adjustment synchronizing efforts across stakeholders, disciplines, and agencies in new ways. Ironically, the guide for a shift in course—as researchers, managers, op-

erators, funders—is to apply the results on adaptive capacity and resilient performance to ourselves to keep pace with changes, growth, interdependencies.

The advances in the underlying science create a timely opportunity now. The next step is for all the stakeholders to come together to pivot from how they cooperate today. This is an invitation to the operational, research, and management communities to chart an expanded course of action and take advantage of the new advances to build strategic agility.

Acknowledgments

David Woods was supported by funding from the SNAFU Catchers Consortium. David Alderson was supported by the Office of Secretary of Defense Strategic Environmental Research and Development Program (SERDP) Project #RC21-1233.

Author Capsule Bios

David Woods, Professor Emeritus in Department of Integrated Systems Engineering at the Ohio State University and Principal at Adaptive Capacity Labs (PhD, Purdue University) has worked to improve systems safety in high-risk complex settings for 40+ years. He began developing Resilience Engineering on the dangers of brittle systems and the need to invest in sustaining sources of resilience beginning in 2000–2003 as part of the response to several NASA accidents (see the 2006 book Resilience Engineering and many subsequent works available at ResearchGate). He developed the first comprehensive theory on how systems can build the potential for resilient performance despite complexity. The results of this work on how complex human-machine systems succeed and sometimes fail has been cited over 39K times (H-index > 92).

He is Past-President of the Resilience Engineering Association. He has received many awards; for example, the Laurels Award from Aviation Week and Space Technology (1995). He has carried out accident investigations in aviation, nuclear power, critical care medicine, crisis response, military operations, and space operations (advisor to the Columbia Accident Investigation Board). He is frequently asked for advice by many government agencies, and companies, both domestically and abroad (e.g., DoD, NASA, FAA; ir France, TNO, IBM; UK MOD, NHS, Haute Authorité de Santé).

David Alderson is a Professor in the Operations Research Department and serves as Founding Director for the Center for Infrastructure Defense at the Naval Postgraduate School in Monterey, CA. His research focuses on the function and operation of critical infrastructures, with particular emphasis on how to invest limited resources to ensure efficient and resilient performance in the face of accidents,

failures, natural disasters, or deliberate attacks. His research explores tradeoffs between efficiency, complexity, and fragility in a wide variety of public and private cyber-physical systems. Dr. Alderson has been the Principal Investigator of sponsored research projects for the Navy, Army, Air Force, Marine Corps, and Coast Guard.

He received his doctorate from Stanford University and his undergraduate degree from Princeton University. He has held research positions at the California Institute of Technology (Caltech), the University of California Los Angeles, the Xerox Palo Alto Research Center (PARC), and the Santa Fe Institute. He has extensive industry experience and has worked for several venture-back startup companies. His early career was spent developing technology at Goldman Sachs & Co. in New York City.

References

Alderson, D.L., (2019). Overcoming Barriers to Greater Scientific Understanding of Critical Infrastructure Resilience. In M. Ruth and S. G. Reisemann (Eds)., *Handbook on Resilience of Socio-Technical Systems*. Edward Elgar Publishing, pp. 66-88.

Alderson, D.L., Brown, G., and Carlyle, W.M. (2015). "Operational Models of Infrastructure Resilience," *Risk Analysis 35*(4): 562-586.

Alderson, D.L., Darken, D.P., Eisenberg, D.A., and Seager, T.P. (2022). Surprise is inevitable: How do we train and prepare to make our critical infrastructure more resilient? *International Journal of Disaster Risk Reduction*, forthcoming.

Alderson, D.L., and Doyle, J.C. (2010). Contrasting Views of Complexity and Their Implications for Network-Centric Infrastructures. *IEEE Transactions on Systems, Man, and Cybernetics-Part A*, 40(4): 839-852.

Allspaw, J., Cook, R. I. (2018). SRE cognitive work. In *Seeking SRE: Conversations About Running Production Systems at Scale*, ed. D. Blank-Edelman, 441-465. O'Reilly Media.

Bergström, J., Dahlström, N., Dekker, S., Petersen, K., (2011). Training organisational resilience in escalating situations, in: E. Hollnagel, Paries, J., Woods, D.D., and Wreathall, J., Eds., *Resilience engineering in practice*. Ashgate, pp. 45–57.

Carlson, J. M., & Doyle, J. (2000). Highly optimized tolerance: Robustness and design in complex systems. *Physical Review Letters*, *84*(11), 2529.

Chiang, M., Low, S. H., Calderbank, A. R., & Doyle, J. C. (2007). Layering as optimization decomposition: A mathematical theory of network architectures. *Proceedings of the IEEE, 95*(1), 255-312.

Chuang, S.-W., Woods, D. D., Reynolds, M., Ting, H-W., Balkin, E. A. and Hsu, C-H. (2021). Rethinking preparedness planning in disaster emergency care: Lessons from a beyond-surge-capacity event. *World Journal of Emergency Surgery*. https://doi.org/10.1186/s13017-021-00403-x

Degerman, H. (2021). Barriers towards Resilient Performance among Public Critical Infrastructure Organizations: The Refugee Influx Case of 2015 in Sweden. *Infrastructures, 6*(8), 106.

Finkel, M. (2011). On flexibility: recovery from technological and doctrinal surprise on the battlefield. Stanford University Press.

Flynn, S.E. (2007). *The Edge of Disaster: Rebuilding a Resilient Nation*.

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

Krieg, R.M. (2020). Editor's Letter, *Journal of Critical Infrastructure Policy*, 1, 1 Spring/Summer 2020.

Mendonça, D., Jefferson, T., and Harrald, J. (2007). Collaborative adhocracies and mix-and-match technologies in emergency management. *Communications of the ACM, 50*(3), 44-49.

Nakahira, Y., Liu, Q., Sejnowski, T. J., & Doyle, J. C. (2021). Diversity-enabled sweet spots in layered architectures and speed–accuracy trade-offs in sensorimotor control. *Proceedings of the National Academy of Sciences, 118*(22):e1916367118. doi: 10.1073/pnas.1916367118.

Rosenthal, C., & Jones, N. (2020). Chaos engineering: System resiliency in practice. O'Reilly Media.

Sharkey T.C., Nurre Pinkley S.G., Eisenberg D.A., Alderson D.L. (2020). In search of network resilience: An optimization-based view, *Networks*, 1-30. https://doi.org/10.1002/net.21996

Woods, D. D. & Branlat, M. (2011). How Adaptive Systems Fail. In Hollnagel, E., Paries, J., Woods, D.D., & Wreathall, J. (Eds.), Resilience Engineering in Prac3ce (pp. 127-143). Aldershot, UK: Ashgate.

Woods, D. D., ed, (2017). SNAFUcatchers Workshop on Coping with Complexity. *Brooklyn NY, March 14-16, 2017 Download from stella.report*

Woods, D. D. (2018). The Theory of Graceful Extensibility: Basic rules that govern adaptive systems. *Environment Systems and Decisions*, *38*(4), 433-457.

Woods, D.D. (2019). Essentials of Resilience, Revisited. In M. Ruth and S. G. Reisemann (Eds)., *Handbook on Resilience of Socio-Technical Systems*. Edward Elgar Publishing, pp. 52-65.

Woods, D.D. (2020). The Strategic Agility Gap: How Organizations are Slow and Stale to Adapt in a Turbulent World. In Journé, B., Laroche, H., Bieder, C. and Gilbert, C. (Eds.), *Human and Organizational Factors: Practices and Strategies for a Changing World*. Springer Open & the Foundation for Industrial Safety Culture, Springer Briefs in Safety Management, Toulouse France, pp. 95-104 https://doi.org/10.1007/978-3-030-25639-5

Woods, D. D. and Allspaw, J., Eds., (2019). Revealing the Critical Role of Human Performance in Software. Special Issue, *ACM Queue*, 17(6), November-December, 2019.

Zhang, X. and Mendonça, D., (2021), Co-evolution of work structure and process in organizations: improvisation in post-disaster debris removal operations. *Cognition, Technology & Work* 23: 343–352.