IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received March 27, 2021, accepted April 6, 2021, date of publication April 8, 2021, date of current version April 20, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3071874

Resilient Control Systems–Basis, Benchmarking and Benefit

CRAIG RIEGER[1], (Senior Member, IEEE), KEVIN SCHULTZ[2], THOMAS CARROLL[3], AND TIMOTHY MCJUNKIN[1]

[1]Idaho National Laboratory, Idaho Falls, ID 83415, USA
[2]Johns Hopkins University Applied Physics Laboratory, Laurel, MD 20723, USA
[3]Pacific Northwest National Laboratory, Richland, WA 99354, USA

Corresponding author: Craig Rieger (craig.rieger@inl.gov)

This work was supported in part by the Department of Energy through U.S. DOE Idaho Operations Office under Contract DE-AC07-05ID14517, in part by the Resilient Control and Instrumentation Systems (ReCIS) Program of Idaho National Laboratory, in part by the Pacific Northwest National Laboratory's Asymmetric Resilient Cybersecurity (ARC) Laboratory Research & Development Initiative, and in part by the Pacific Northwest National Laboratory is operated by Battelle for U.S. Department of Energy under Contract DE-AC05-76RL01830.

ABSTRACT Because the research area of resilient control systems was pioneered during the last decade, the basis and benchmarking of resilience have continued to mature to achieve what has long been understood as the ultimate benefit of resilience. However, the automation ''ship'' has long since sailed on society's dependence on digital control systems as the basis for all our industries and even the appliances in our homes. While these systems have been in general very reliable and provided for many human and operational efficiencies, the designs were not built on a framework that recognizes and adapts to potential debilitating failures from events such as cyber-attack. In this review, we cover a rapidly maturing framework based upon a disturbance and impact resilience evaluation process that considers both the methodologies for assessing resilience and also how key design principals must be applied within distributed control systems to achieve resilience.

INDEX TERMS Resilience, control, cognitive, cyber, infrastructure, metrics.

I. INTRODUCTION

As the scale and inter-connectedness of critical infrastructure (CI) and cyber-physical systems CPS have increased over the years, their vulnerability (and visibility) to a wide range of failures and disturbances have become more prominent. Events such as component failures, natural phenomena such as storms and earthquakes, and cyber-events (both ''benign'' events caused by unexpected interaction and malicious attacks) can all severely impact performance [1]–[5]. Several notable examples include the 2003 New York City blackout [6], hurricanes Katrina [7] and Sandy [8], the 2013 Okhotsk earthquake [9], and cyber-attacks on the Ukrainian power system [10], [11]. Human error can also lead to cyber-events [12].

The desire to analyze, manage, and mitigate large-scale systems in response to organizational and systemic vulnerabilities, disturbances, and failures has led to a multidisciplinary field of *resilience* engineering [13]–[17].

The following is adapted from [16], which provides a definition of resilience:

> A resilient cyber-physical system is one that maintains state awareness and an accepted level of operational normalcy in response to disturbances, including threats of an unexpected and malicious nature.

The process by which control system resilience is considered in the context of this paper is the disturbance and impact resilience evaluation (DIRE) curve (Figure 1), discussed in the next section. Generally, the time scale effects of cyber-physical disturbances on the system are short, but any damaged components are longer term to accommodate replacement. For a CPS to be considered resilient, therefore, requires it to suitably reconnoiter its environment for threats, resist these threats through adaptation and agility, and recover from degradation. For the purposes of evaluation, normalcy, or the maintenance of an accepted normal,

The associate editor coordinating the review of this manuscript and approving it for publication was Emanuele Crisostomi.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

is defined by the physical-process owner/operator. These resilience attributes are often achieved by a multidisciplinary team and correlated to critical functionality [16].

Due to such high-profile events as [18], there is demand for increased security and resilience valuation to prioritize investments to mitigate the consequences of such events [19]. The recognition of these severe man-made and natural (potentially catastrophic) events has led organizations such as the Department of Energy (DOE) in the United States to advance programs aimed at achieving greater grid resilience through distributed systems, such as microgrids, and advanced cybersecurity [20], [21]. Thorough investment, starting with valuation, allows industry to prioritize investments in technologies and human capital, including training and response activities. This valuation includes tools to understand interdependencies and potential impacts [22]—in addition to recognized critical operations—that are not well characterized.

Often, these concepts are added to fulfill a desire to give the appearance that everything was considered in the design. However, the definition of and technological advances required for resilience are fundamentally different [23]. It is altogether unlikely that these considerations are at the core of developed technologies. With this in mind, one might ask how secure and resilient a new technology would be.

The answer to will depend upon design. For security, it will, in all likelihood, inherit existing, well-recognized security methods that include border defenses, password protections, selective encryption, etc. These are necessary building blocks that require advancement to achieve cyber-resilience [24]. While this might provide some level of comfort when it comes to resilience, it seems likely that the system will have very selective or no resilience benefits if these concerns are not considered as core design concepts. Resilience technologies, to meet the definition, must have state awareness of degradation and must adapt or transform to maintain desired system performance. It appears difficult, at best, to implement these capabilities in a retrofitted capacity without the appropriate design considerations to enable access to both state-estimate and control actuation.

As we consider the resilience of our infrastructure to cyber-attack or other disturbances and threats, we must directly consider the control systems that provide human-in-the-loop mechanisms to monitor and control the power, water, and other utilities upon which we depend. As the desire to automate and achieve efficiencies of labor and operation has grown, so has investment in control systems to allow for integrating different operations, facilities, utilities, and infrastructures. However, the evolutionary integration of control systems has led to complexities of failure, human interaction, and security vulnerability. As control systems evolve toward greater autonomy, reducing and changing the role of the human, the need to consider resilience is underlined. Autonomous systems can react quickly to anomalous conditions—for example, ensuring we have power even if a transformer fails. However, they can also cause a quick escalation to a cascading fault if autonomy has been corrupted by a cyber-attack or an unrecognized failure, or even by a simple flaw in design. Enabling the human in the loop will be necessary throughout, ensuring the ability to adapt to anomalous conditions in ways that the control system cannot.

The next generation of control systems should have a threat-based approach to develop systems that are resilient by nature. In what follows, we first review the DIRE curve, which presents a framework for analysis, design, and implementation of resilience. Next, we summarize a number of metrics used in the analysis of resilience and considerations for incorporating resilience at the design stage of distributed control systems and even larger-scale Systems of Systems (SoS), and how these can be assessed in the context of DIRE. Finally, we demonstrate the flexibility of the DIRE approach by using it as both a tool for analysis at a detailed, physical-model level and as an abstract SoS scenario.

II. INTERPRETING RESILIENCE USING THE DIRE CURVE

The ability to correlate resilience starts with a generalized physical depiction of how it is correlated, which provides a means of evaluating and comparing benefits. This correlation is based upon system performance in meeting desired minimum operation. As a result, an acceptable level of degradation is based on the determination of a resilience threshold. The desired outcome is a mapping of the capabilities and limitations of a CPS to the resilient-control metric that expresses the ''R's'' of resilience in Figure 1, the DIRE curve [25].

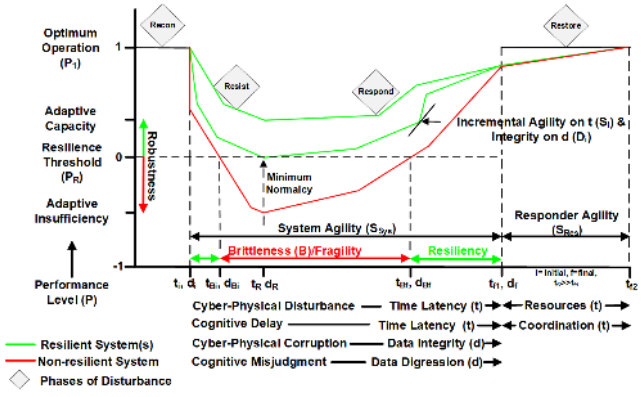

FIGURE 1. Disturbance and Impact Resilience Curve. The resilience of a system can be understood by its ability to initially reduce the impact of a disturbance (resist stage) and recover from it in both the short- (respond) and long term (restore). The red curve indicates the trajectory of a system that is not particularly resilient and falls below some predefined normalcy criterion while the green curves are systems that maintain a minimum level of acceptable operation during this crisis, indicating resilient (as opposed to fragile/brittle) systems (the upper, more so than the lower).

As evaluated in prior work, the ability to ensure tracking of all states requires that the number of control inputs is at least equal to the number of states controlled [26]. With this assurance, the resiliency limit or stability can also be ensured. Add to this the element of cyber-resilience, and the ability to validate operation is dependent upon the ability to have the same number of responses to the vectors for

compromise. For cognitive resilience, the benign actor is considered, and the responses are not only to ensure correct interpretation and, therefore, human response, but also corrective actions in the event of an undesirable operation. Referring to Figure 1, the impact is summarized in terms of time and data. The phases that characterize this adaptability and agility to threats or disturbances constitute four, with recovery split between two aspects that consider the time scale of the recovery. While similar to other correlations of resilience—e.g., [27]—these phase terms reflect the consideration of the contribution of the control system.

- Recon: Maintaining proactive state awareness of system conditions and degradation
- Resist: Responding to recognized system conditions, both to mitigate and counter
- Respond: Once system degradation has been stopped, returning system performance
- Restore: Restoring longer-term performance, which includes equipment replacement.

While these phases are listed separately, it should be noted that the control system resilience implementations that address these phases can benefit more than one phase. Specifically, recognition of degradation is a necessary need for the recon phase, but is also used to affect resist and respond phases. For example, a mitigation in the cyber environment might include the isolation of a nefarious external connection or in the physical environment to swap over to a redundant control device.

The magnitude and duration of a disturbance that can be withstood depends on the agility of the system [28], [29]. Both the resistance and responsiveness of the system to disturbance, which could include transition of functions to maintain critical operation, consider both the time scale of response and the ultimate operational impact. Within a CPS, these can be very short. However, where equipment has been damaged in the field, longer-lead-time resources are required to recover and restore the system to normal operation. These include both the purchase and installation of equipment. However, the prior stages of a resilient system would still provide context to identify quickly the affected components and, thereby, a more efficient and effective response.

Incremental performance (P_i) is defined as the difference between the performance at any one point in time and the resilience threshold (P_R) relative to optimal performance levels (P_1). The incremental performance of the system is also defined with a positive number ($0 < P_i < 1$, robust), indicating that resilience and an adaptive capacity are maintained in response to the disruption. A negative number ($0 > P_i > -1$, not robust) indicates the system has an adaptive insufficiency. The overall performance (P_S) maintained through a disruption from maximum (P_1) is defined in reference to the resist and respond phases of Figure 1. The incremental agility and integrity are determined at any point, with either t^* and d^* as the calculation point, respectively, with the other held constant. For longer characterizations, additional statistics can be performed to calculate the trends in each phase.

The overall performance, P_S, of a system can be characterized based upon dependent and independent variables of a physics-based model or empirical relationship. However, when this type of relationship does not exist at the system level, the resulting necessary adaptive capacity can be correlated based upon common factors that are representative of system design and by which design requirements and limitations must be defined. Two such factors for deterministic systems typical of aircraft and missile platforms include latency and data integrity, as both represent factors that influence the stability of the overarching communication and control system design. Based upon the platform-engineer specifications, preferably outcomes from physics and/or rigorous testing, acceptable latency and integrity form the basis for scaling adaptive capacity. For example, P_S could be based upon the acceptable latency and data integrity as critical variables associated with aircraft control, like altimeter or speed indicators. As a result, the specifications form the basis for metrics evaluation as follows for any disruption:

- The specification tolerances on communications latency and data integrity that are allowed before an unacceptable event will happen form the resilience threshold (P_R), with the incremental performance (P_i) based upon any value taken in time. When P_i is zero, the subsystem is at the resilience threshold, indicating that the incremental performance evaluation is at the specification limit.
- For overall system performance (P_S), the ability of the system to maintain critical performance is evaluated. Values of the overall degradation of the system are evaluated, and larger numbers are indicative of greater resilience.
- For the resist and respond phases, agility (S_i) and integrity (D_i) are defined at any one point based upon the chosen critical variable (e.g., voltage, current, etc.) and provide an incremental indication of sensitivity to our measures of critical-variable variations. Smaller slopes indicate less sensitivity of the performance measure to disruptions in general. In the resist phase, decreased S_i and D_i are preferred. In the respond phase, the converse is indicated.
- For each of these indicators, an adaptive capacity can be associated to recognize and correlate decreased performance due to a cyber-attack.

Application of this information can be used in design or operation, such as in considering latency issues. In this case, critical functionality can be maintained in a communications system by performing actions such as designing for greater bandwidth, reducing frequency, or using redundant pathways that can be opened upon recognition of the degradation.

III. COMPLEXITY AND RESILIENT CONTROL DESIGN

There are a number of complexity challenges that will need to be addressed in future distributed control-systems designs. First, current automation environments are the result of the organic interconnection of control systems, leading to an

inability to recognize and prevent resulting unanticipated faults. Second, the human element—whether benign human error as the result of data overload and lack of information or the work of a malicious human who exploits current perimeter protections that are insufficient and not designed to adapt rapidly to attacks in order to prevent compromise—have the capability to degrade systems. Finally, current control systems have multiple performance goals, but the lack of necessary identification and prioritization of these goals can lead to an undesirable response from both human operators and the automated design [25].

Nominally, the distributed control elements of control systems are associated with some optimally stabilizable entity. This can be seen by looking at chemical-process plants, where a collection of separate unit operations make up an integral plant [30]. The unit operation, in this case, defines an area of local optimization. Within the operation, many physical variables may exist. Typically, in a plant made up of many unit operations, the process of determining the optimally stabilizable entities normally results in a minimization of interactions between individual operations. That is, only a few physical variables will normally make up the interactions between unit operations. For example, the flow between unit operations must remain within a specified range because the downstream operation is designed to be stable for operation within that range. This decomposition of a complex control system into different functional units with different objectives leads to notional architectures such as the distributed, multi-agent architecture shown in Figure 2.

The process of determining unit operations suggests a complexity-reduction approach for subdividing infrastructure to increase resilience [31], including how the power grid might be subdivided into various forms of micro- or macro-grids [32]. Within these subdivided areas, power stability is maintained against threats or has the ability to regulate effects from destabilizing forces such as intermittent generation. Through minimization of cybersecurity, control, power, and other dependencies and interdependencies, local regions maintain their own stability and prevent cascading affects. Once this distributed architecture is defined, the cybersecurity and resilience of the individual and aggregated system are designed and developed. Where interdependencies and dependencies remain, polymorphic isolation techniques can be concentrated to counter propagation of threats. The result is a foundation or building block for greater efficiencies, where wide-area supervisory strategies can be built without engendering complex failures. The application of fundamental physical characterizations to establish manifold resilience can then be formed with solvable complexity using standard physical formulations [33], [34], as demonstrated in Section IV-B. The parameters associated with this manifold is based upon a power system, defining the real and reactive power over time and presents the impact of disturbances.

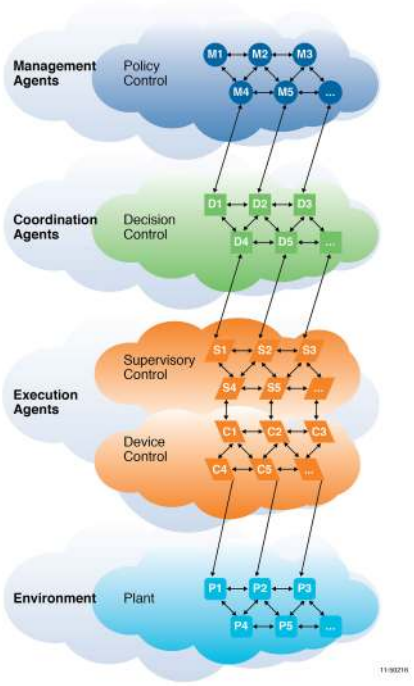

FIGURE 2. Notional breakdown for distributed control systems. A complex system is naturally divided into different levels of control, with different "units" responsible for optimization of limited goals and management of local interfaces between units that combine to contribute to an overall objective.

A. DESIGN AND OPERATIONAL CONSIDERATIONS FOR CYBERDEFENSE

Distributed control systems integrate physical components with networked computational elements, and exhibit properties spanning the cyber-, cyber-physical, and physical domains. In this section, we focus on the cyber-domain aspects because this represents an implicit, yet major (if not the dominant) interdependency. Conventional information security is defined around three security objectives: confidentiality, integrity, and availability. These objectives are commonly referred to as the triad, and while we speak of them in isolation, that does not reflect reality because the objectives are not truly orthogonal. An organization pursues a balance of the three that meets its needs. Traditional information technology (IT) security promotes confidentiality and integrity, with availability afforded a lesser role. The safety and reliability design requirements of CPS results in a significantly different balance, one that maps strongly to the integrity and availability objectives [35]. Moreover, current perception is that confidentiality is expensive, coming at the cost of safety and reliability (this can also be thought because an increase in confidentiality is assumed to reduce availability). This mismatch of objectives causes conflict between the IT and operational technology (OT); a product designed and promoted for one is unlikely to meet the requirements of the other. Recent cyber-attacks on power systems, such as

those that occurred in 2015 and 2016 in Ukraine [10], [11], have shown that these complex systems can be compromised remotely, employing cyber domain-only means.

CPS vulnerabilities are commonly caused by (i) violation of isolation assumptions; (ii) increased connectivity, especially to business systems; and (iii) heterogeneity of software and components [36]. Initially, there was a basic assumption that CPS were isolated from the business/IT components of the organization and the outside world. Security was not of great importance, and this led to the development of simplified security controls and a dependency on trust. In practice, the isolation assumption is often violated in the name of business efficacy, which results in business systems linking to CPS to obtain information relevant to business processes and operations. Integration of heterogeneous components, which inherits the vulnerabilities of each component [37], combined with increased connectivity resulting from additional network services and protocols, is a recipe for communications and software vulnerabilities [36].

Security benchmarks may play an essential role in next generation CPS. Security measures, metrics, and benchmarks are tools to facilitate decision making about various aspects of security, including the design, architecture, and efficacy of systems and controls, to acquire and maintain operation and situational understanding necessary to defend against dynamic threats. While there is ongoing research (see, for example, [38]–[41]), robust, measurable, and quantifiable security metrics remain an open challenge. The gap between state-of-the-art metrics and the ultimate goal is significant. One limit is that a nontrivial metric cannot exist that denotes the absolute security state of a system [42].

Metrics can then only measure relative changes or improvement in design or architecture. Another limit is that many security metrics are lagging indicators of threats, which reflect conditions that exist after an attack [43]. More beneficial metrics would be coincident or leading indicators of security conditions. Furthermore, many metrics don't capture the dynamic security state of the system—information that is relevant and necessary for decision making and situation understanding [40]. Unlike general computing, CPS serve a mission in support of well-defined processes. Safety and reliability requirements establish performance requirements that can be measured and baselined. Moreover, we can establish benchmarks of characteristics that are degraded as effect of attacks [25].

Practical performance-based resilience metrics can then be derived and monitored to evaluate the system condition [25], [44]. Operators and components can acquire and maintain CPS state awareness without establishing the fact that an attack has commenced. The methodology comprises four steps: (i) baseline the system, acquiring and repeatedly measuring indicators over time, (ii) identify threats, developing test cases, (iii) identify the benchmarks that are disturbed during testing, and (iv) establish and evaluate normalcy threshold criteria. For example, consider a denial of service (DoS) attack, which may be the attacker's primary motive or a secondary effect of the attacker's methodology. Depending on the DoS specifics, we would expect adverse effect on one or more benchmarks. If the DoS attack is network borne, and the purpose is to overwhelm the ability of a component to respond to requests or the network to function, benchmarks of standard network performance—such as data rates, latency, jitter, and round-trip time—would be advantageous to demonstrate the existence of the attack. We would expect to see other network indicators of the attack, some of which are specific to a component: increasing number of TCP retransmissions, resets, and handshakes; ICMP unreachable notifications; and other artifacts of protocol-specific ''squelch'' mechanisms. If software is instrumented, a general trend of increasing request response time, along with timeouts, would be expected as an indicators of a DoS attack. The benchmarks that are described are sensitive to DoS attack, but are not specific to one. That is, other causes may be precipitating the degradation indicators. For example, a failing network transceiver may result in degrading many of the same indicators as a DoS attack.

The discussion up to this point has primarily focused on availability of the information security concepts confidentiality, integrity, and availability. This is intentional because performance metrics related to availability are readily quantified and measured. The same cannot be said for confidentiality and integrity [45]; they cannot be directly measured. Instead, operationally aligned proxy metrics are derived to obliquely quantify and measure them. Consider a data-injection attack, which affects both integrity and confidentiality of the stored data. If the motive of the attack is to exfiltrate data, benchmarks of disk and network input/output may be illustrious. A change in the amount of data written, the types of commands, and other related input/output indicators can be used to establish the loss of data integrity.

A system can be expressly designed and engineered to express integrity measurements [46]. For example, to confirm the results of a control action, we can model and simulate the system and how the action affects the system and environment. We can then test agreement between actual readings of independent system and environment sensors and modeled readings. A DIRE-compatible metric can then be constructed that measures the difference between actual and modeled readings. While it is feasible to design a system that emits availability and integrity, it is not evident and remains an open challenge how a system may be designed to support direct measurements of confidentiality.

Existing cyber-resiliency metrics may not be compatible with DIRE. Referring to a catalog of nearly 500 representative cyber-resilience metrics [41], we remark that many of the metrics that relate to integrity and confidentiality do not exhibit an explicit system basis. The metrics are akin to process-improvement metrics: they provide a basis to measure improvements in mission, program, and organization resilience postures. While metrics such as ''SI-IC-7: Frequency of hardware/system integrity check,'' ''SI-BV-2: Percentage of mission-critical applications for

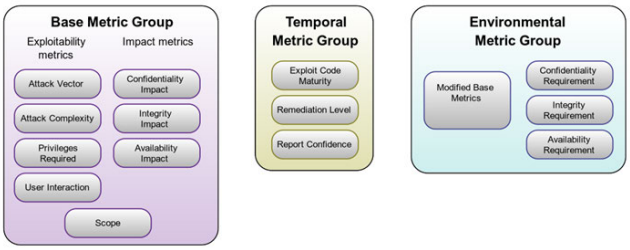

FIGURE 3. CVSS Metric.

which integrity/behavior can be validated," and "SI-BV-6: Frequency of check for faulty processes or services" are indeed helpful in assessing and improving resilience, the character of the metrics means that they are insensitive to a system disturbance. Moreover, they are measured "out-of-band" as the metric manifests itself and is interpreted at the mission, program, and organization levels. Consequently, metrics of this character are unfit for use with the DIRE methodology.

Security score systems, such as common weakness scoring system (CWSS) [47] and common vulnerability scoring system (CVSS) [48], can be adapted for use in DIRE. Consider CVSS. CVSS computes the severity of vulnerabilities in the range 0–10, with 10 being the most severe. It considers three areas of concern (see Figure 3):

- *Base metrics* are metrics intrinsic to the vulnerability and are constant with time and environments,
- *Temporal metrics* are metrics that reflect the characteristic of the vulnerability that change over time, but not across environments, and
- *Environmental metrics* are the elements of the vulnerability that are relevant to a specific environment.

The environmental metrics adjust the score with respect to a system's operational and programmatic context (i.e., importance) and considers the presence of system and organizational controls to mitigate the vulnerability. The CVSS score weights confidentiality, integrity, and availability impacts. The score is then assigned a rating of *low* (0.1–3.9), *medium* (4.0–6.9), *high* (7.0–8.9), and *critical* (9.0–10).

To adapt the CVSS as a DIRE-compatible system performance metric, the following process is performed. First, a system's confidentiality, integrity, and availability requirements are defined. Next, the Common Vulnerabilities and Exposures (CVE) [49] and National Vulnerability Database (NVD) [50], both authoritative sources of vulnerability information, are referenced to identify and catalog system-relevant vulnerabilities. A disclosure describes the issue and references sources for detailed information, where the CVSS base metric score maybe obtained. For each vulnerability, the environmental metric is scored. The maximum score of all the environmental metrics is selected, and the performance metric is then defined as one minus the maximum, obtaining a numeric value that operates in the range 0-10, with 10 corresponding no outstanding vulnerabilities. The performance metric is reevaluated as vulnerabilities applicable to the system are disclosed or revised, when mitigations and counteractions are applied, or when the system's operational, programmatic, or environmental contexts change.

B. DESIGN AND OPERATIONAL CONSIDERATIONS FOR PREVENTION OF CASCADING FAILURES

In interconnected CI, initial points of component failure can propagate to other components in the system, resulting in so-called cascading failures or exacerbate existing independent failures, resulting in so-called escalating failures [51]. These failure modes are perhaps most notoriously observed in power systems, but computer networks, transportation systems, and even abstract networks such as those between markets and economies in financial networks are susceptible. Cascading failures can even cross from one network to another when the systems are interdependent—for example, power failures resulting in widespread roadway congestion, or the recent computer network failures that have severely impacted airline operations. An essential step in the analysis of cascading failure modes is to determine the dependencies (one-way interactions) and interdependencies (bidirectional interactions) in a system. In a SoS context [52], the interactions between the different systems must also be assessed. Additionally, the input and output connections of each system to be analyzed, sometimes called the upstream and downstream dependencies [51], must also be considered because these connections may not be connections within the SoS itself. Dependencies can take many forms, with a common taxonomy that classifies the interactions into one or more classes of physical, cyber, geographic, and logical dependencies such as in Figure 4.

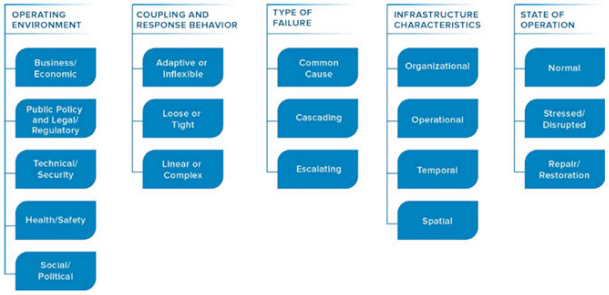

FIGURE 4. Dimensions of dependencies.

Overall, this SoS approach to the analysis of critical infrastructure seeks to identify and categorize all manner of interactions as a necessary first step to enable the prediction and prevention of cascading and escalating failures. The notion of assessing the components in a SoS and assessing their interconnections as dependencies immediately suggests graph-theoretical models for CI and SoS where the nodes or vertices of the graphs are the individual SoS components, and the links or edges in the graph are the (inter)dependencies between them. Indeed, in the nonlinear physics and complex systems communities, cascading failures in interconnected

systems have been studied using graph theoretical models in the context of network and percolation theory. The most abstract form of these types of models (e.g., [53]) treat the spread of failures in a manner analogous to the spread of disease in epidemiological models [54]. Other approaches build on the graph abstraction and seek to add elements of physical realism, such as the capacity of the networks [55] or power-flow and load-frequency control [56].

Network-theoretic modeling approaches are flexible enough to address a wide variety of scenarios and contexts, are amenable to mathematical analysis, and are generally computationally tractable and can simulate networks of many thousands of nodes. Much of the power of these approaches comes from their abstraction: however, recent work suggests that abstract models of cascading failure can produce starkly different results from more-detailed models of the systems in question, especially when the impact of control systems are considered [56]. Increasing the physical realism of the network models in question potentially limits the size and scope of the analysis that can be performed, but it is still possible to analyze large SoS at varying levels of physical fidelity. Even as the physical realism in these models increases, the intuitive representation of the SoS as a set of interconnected graphs is prevalent, and characterization of networks is performed using the abstract graph model.

The literature on network and graph theory offers a wide range of metrics to analyze SoS for the purposes of risk assessment or investment. These range from metrics assessing individual vertices or edges that can be used to assess the relative importance of individual components or links in the SoS, to other metrics that seek to quantify aggregate characteristics beyond a single component, and all the way to global properties of the graph model [57]. Single-component metrics, such as the various flavors of centrality (e.g., degree, between-ness, closeness, etc.) are useful in identifying which components are especially likely to result in cascading failures due to their ''influence'' on other components in the network. A recent contribution to this class of metric is percolation centrality, which is designed to assess the criticality of nodes while a network is undergoing percolation [58]—i.e., while failures are currently cascading in this context. Thus, this metric and its generalizations may be useful for identifying upcoming failures in a cascade for the purposes of online prevention.

Alternatively, one can consider global metrics of the abstract graph model. One class of global graph metrics includes averages, variances, or other functions such as entropy over lower-level metrics, with the intuition that some form of ''balance'' in the metric over the graph will result in a more resilient SoS. Other examples of global graph metrics quantify stability or reliability of various graph properties, such as how the removal of edges or vertices affects graph diameter or connectivity. It is easy to see an operational context behind this latter class of metrics because it, in some sense, indicates the ability of the network model to maintain functionality under duress. Furthermore, many of the metrics in this class can be generated from the algebraic and spectral properties of the graph adjacency and Laplacian matrices, which provides a connection to the field of distributed control where these matrices can be used to derive control-theoretic properties.

Another approach to global metrics is to broadly categorize SoS graphs and subgraphs into general classes, such as small-world or scale-free graphs, and apply these concepts to make general statements about the SoS. In between the two extreme classes of graph metrics above, there are those that use more ''regional'' properties of the underlying model. For example, there are metrics based on the number of cycles that a given vertex or edge belongs to, or those based on the largest fully connected subgraph (called a clique) to which a vertex or edge belongs. Another example of a metric in this class is the number of unique paths between two given vertices, often used as a measure of redundancy in the network.

While the network-theoretic approaches to the modeling and analysis of cascading failures described above have dominated the research literature, other fields have contributed to this area. For example, an important consideration regards the nature of the generation mechanism of the cascading failures. The failure mechanisms can range from purely random to adversarial, which naturally leads to applications of game theory [59] to the analysis and defense against cascading failures in SoS. Contributions to this topic can be found in other fields. Other approaches involve more-traditional control-theoretic stability analysis of interconnected systems—e.g., [60]—or other techniques from the stochastic process literature, such as [61]. In addition to approaches in the research literature, a number of tools are commercially available, such as those discussed in [62]. These simulation tools, as well as any of the more-abstract modeling and simulation approaches above, can be used to assess the likelihood of cascading failures and to infer critical components and interconnections in the SoS.

The DIRE curve can be used for an individual or systemic view of an operation, such as an individual turbine or a full generation station. However, when the competing complexities of large interdependent systems must be considered, the SoS provides a rationale for considering what these impacts might be in planning. While methods to link disparate infrastructure and weather models have been developed to enable the confirmation of some impacts [63], [64], these are not necessarily definitive in all aspects (e.g., political) of SoS, and the methods have limitations of scale in considering the global.

IV. INSTANTIATING A TANGIBLE MEASURE OF RESILIENCE

A. GENERIC FORM OF A RESILIENCE MANIFOLD

Achievement of control-system resilience within any domain starts with an extensible approach across all domains, including complexity, benign and malicious humans, and conflicting goals. Fulfilling this objective starts with definition of those physical capabilities, over time, that form the envelope

of an ability to adapt and margin to maneuver to maintain normalcy before, during, and after a disturbance. Given a physical basis as the top-level indicator of performance in Figure 1, any cognitive, cyber-, or physical disturbance will impact that performance indicator over time. Knowing this, we can suggest that, given a physical relationship, whether empirical or first principles, any impacts of a disturbance can be directly evaluated based on variations in the data or timing.

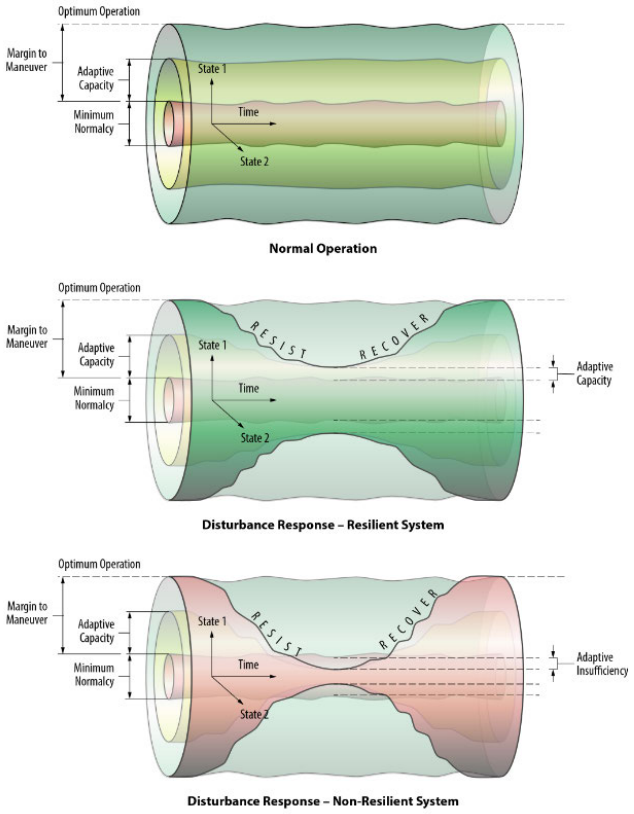

FIGURE 5. Resilience Manifold.

Figure 5 provides the resulting resilience manifold perspective on considerations for measuring resilience. Note that, in considering the comparisons among different performance indices, state variable relationships can be defined based upon the particular domain. Examples of process domain-state variables include pressure and flow for process industries or voltage magnitude and phase for steady-state electricity transmission and distribution systems. The tangible inputs influencing these examples are the asset capabilities that can respond to the disturbance in the state of the system.

Within the manifold of Figure 5, the margin to maneuver is the maximum performance of the system during optimum normal operation, which considers the operating set point dictated by the system owner. When evaluating resilience for a system, it is particularly important to recognize native pinch points where the margin to maneuver is reduced. For a power grid, the native pinch points generally reflect a narrowing of the band between available generation and loads that occur during peak demand. Within the pinch points, adaptive capacity is the amount of flexibility the system has to absorb disturbances such as a tornado taking out power routes or a cyber-attack on a substation. Where the margin to maneuver is reduced, the available adaptive capacity is also generally less. This implies that two considerations provide the bounding of the dynamic attributes of resilience: the margin to maneuver, which is dependent upon the constraints of the equipment and system-operator settings, and the adaptive capacity, which is limited by the margin to maneuver and provides the maximum capacity of the system to address disturbances.

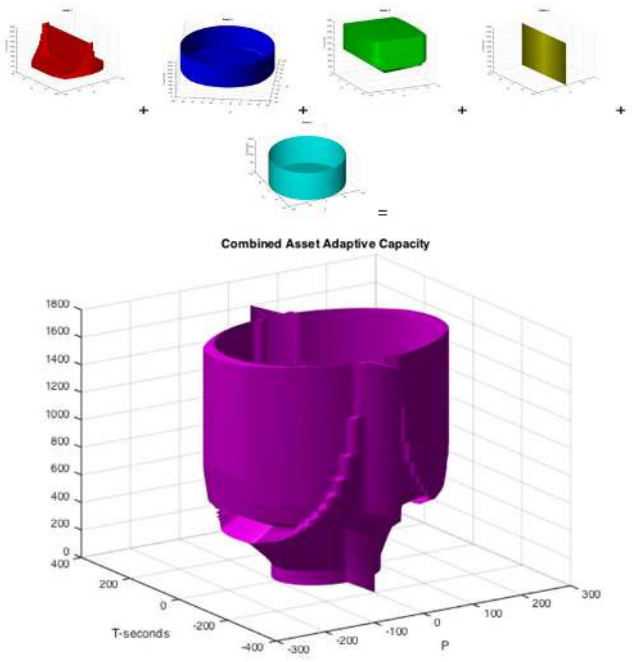

FIGURE 6. Power Asset Aggregation Manifold [5].

B. APPLICATION OF MANIFOLD TO BENCHMARK PERFORMANCE IN THE ELECTRIC GRID

For power, the capabilities that can be brought to bear against a disturbance are the real (P) and reactive (Q) contributions for multiple distribution power assets. These can be combined mathematically to form the manifold, as seen in Figure 6 [65]. For this figure, these assets include: a) a battery, b) an alternate-tie line source, c) an asymmetric P/Q-conjectured source, d) a DSTATCOM, and e) a low latency, four-quadrant source with no energy limit. While these provide individual considerations that are of localized interest, mass and energy provide the basis for overall resilience of an infrastructure system. In developing this understanding, a systematic approach is needed. This approach evaluates individual systems based upon the physics, but also integrates with communications and control dynamics to correlate impacts of natural and man-made attacks in terms of the originating-disturbance type.

The manifold of Figure 6, describing the capabilities assets (a, b, c, d, e), are found by first defining the key simplified parameters needed to describe the assets:

- $P_{ck}m$ and $P_{ck}M$ - limits of real power level of device m for minimum and M for maximum
- $Q_{ck}m$ and $Q_{ck}M$ – limits of reactive power
- $E_{ck}m$ and $E_{ck}M$ - energy limits in the device
- t_l - device latency
- t_P and t_Q - agility or time to ramp real or reactive power.

Note this is a simplification of asset descriptions, assuming symmetric manifolds for illustration. The full description of the assets in an operational context are presented in [66]. The control system in the Recon phase should be designed to drive the bias in a direction that is optimized for the anticipated situation. From a steady-state perspective, each component attached to drops of a distribution network has an apparent power range over the complex S plane, $S = P + jQ$. The relevant control goal for a distribution network is to drive $S = \sum_i S_i = 0$. Some devices, like batteries or generators with a limited amount of fuel, have energy constraints. This energy limit restricts the duration for which the asset can be applied. Other temporal characteristics of the device are captured in the device latency and the agility of the device to ramp either real or reactive power.

Based upon these parameters, a full description of the S-plane manifold is provided that specifies the maximum adaptive capacity, over time, of the collection of assets applicable to a disturbance, as shown in Figure 6. The complex shape is bounded by the roll-up of the adaptive capacity, with rate constrained by the agility of the components and any pure latency. The surface grows based on agility up to the maximum magnitude available and continues in that direction until energy is depleted and the P goes to zero, reducing the limitation on any trade-off between P and Q.

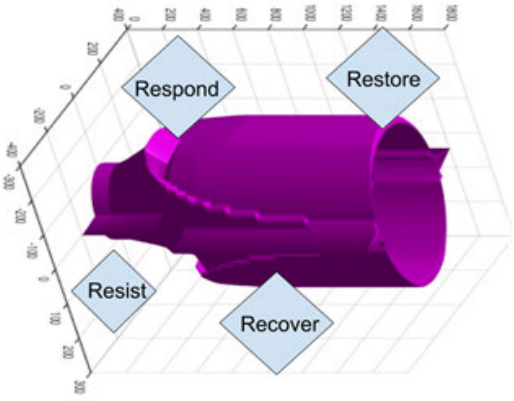

FIGURE 7. Mapping Manifold to DIRE.

The mapping to the DIRE curve is illustrated in Figure 7. The Resist phase provides the intrinsic support that may slow the effects of the disturbance (e.g., synchronous machine-supplied inertia or synthetic inertia provided by power electronics). The Respond phase consists of the capacity of short-to-medium time-period use of the very fast and low-level control loops designed to react without supervisory decision. Finally, the Restore phase is the portion that is required to bring the depleted assets back to bias points where they would again be available to respond to another disturbance. The Recovery phase considers the Respond and Restore phases. Total P, Q, and E that can be applied over a given time period defines the magnitude and duration of disturbance that can occur without dropping below minimum normalcy, as defined by the stakeholders.

For the distribution-system power example shown in Figure 6, the disturbance at time zero is associated with a steep loss in remaining adaptive capacity, followed by a gradual recovery to full adaptive capacity. However, one or more disturbances can be considered at any time and are of particular importance at the pinch points. The resulting time scale of the aggregated response to one or more disturbances then provides the resilience of the individual system that, when decomposed to a distributed, dynamic hierarchy as discussed in Section III, characterizes the overall system of systems resilience by characterizing the limits of magnitude and duration of disturbances for which minimum normal operation can be maintained.

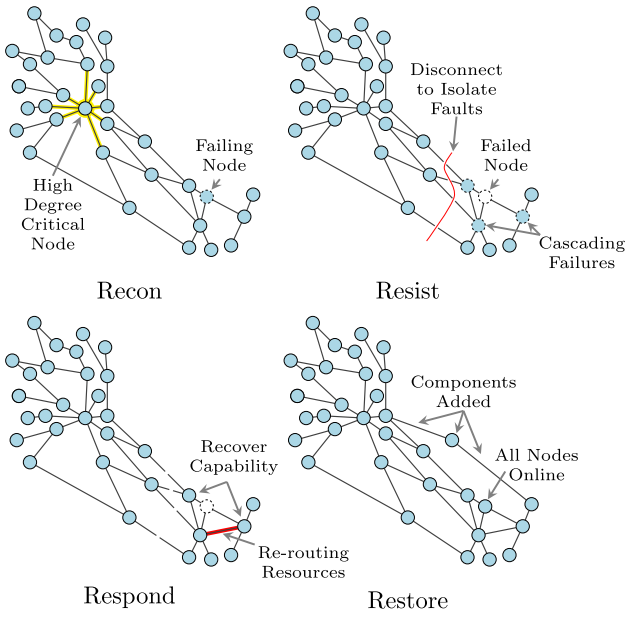

FIGURE 8. Sample abstract graph interactions of a SoS in the DIRE framework.

C. SoS APPLICATION TO DIRE PHASES

To emphasize the flexibility of the DIRE approach, we present another example that considers a larger system at a higher level of abstraction. Within Figure 8, a notional graph-theoretic representation of the SoS approach is illustrated by DIRE phase. In this example, the nodes of the graph represent different subsystems in the SoS, and the edges represent the interdependencies between them as described

in Section III-B. As a concrete example, these nodes could represent generators and transformers in an electrical system, with the edges representing the lines between them. Alternatively, the nodes could represent a mixture of electrical and hydraulic components, with the edges between electrical components representing transmission lines, the edges between hydraulic component pipes, and the edges between nodes of differing type more-abstract dependencies, such as cooling requirements of the hydraulic components on the electrical components or power requirements for the hydraulic components. We emphasize that it is possible—indeed, likely—for all four stages of the DIRE framework to be occurring in different components of SoS simultaneously as individual subsystems in the SoS should have their own DIRE analyses that implement the parts of the four stages that apply to the collective DIRE analysis of the composite SoS. In particular, the Recon stages should be continually active across all four stages to monitor any active faults and the potential for new faults.

The Recon phase of the DIRE framework is represented in the upper left of Figure 8, where a critical node (highlighted in yellow) has been identified through analysis of the SoS interdependencies in the abstract graph model—here, by means of its high degree (primarily chosen for illustrative purposes). Other metrics from Section III-B are, of course, also applicable as defined by the context of a specific SoS. We point out that, again, depending on the specifics of the SoS of interest, the analysis of critical nodes may be primarily static if the overall network is static. However, criticality may vary over the course of time if dependencies are dynamic and/or the edges in the model include any actual usage or capacity that would cause the relative importance of a node to vary over time. In either case, the Recon phase of the DIRE framework is responsible for both understanding the real-time importance of different components in the SoS and also monitoring the SoS for potential issues such as component failure or cyber-attack.

In the context of a disturbance detected during the Recon stage and resulting in, e.g., the failure of a node such as a generator, the Resist stage of DIRE is responsible for attempting to locally stabilize the disturbance in the overall SoS, as depicted in the upper right panel of Figure 8. Overall, the purpose of the Resist stage is to slow degradation in SoS capability. In the case of a power system with a disturbance exclusively affecting reactive power, the Resist phase would buffer the disturbance from reaching any critical nodes by first absorbing the disturbance at an induction generator or other reactive source or sink. These initial attempts at resisting degradation in capability may not be sufficient to mitigate the disturbance, and the failures may begin to spread to nearby nodes through the dependency links indicated in the graph model. As noted in Section III-B, cascading failures are a phenomenon in complex SoSs that are capable of crossing system boundaries in a SoS and can be particularly devastating. As a corrective action to a cascading failure in a power system, the cascading faults could be isolated through a control-system policy or procedure that removes pathways from the fault through the operating environment. Another example of such a resistance measure is the isolation of compromised components in an IT system to prevent further spread of a cyber-attack. In a graph-theoretical context, this isolation amounts to a graph cut that isolates the cascading fault from the critical node. This sort of drastic action serves two purposes: 1) it limits the overall spread of the cascading failures, but also 2) it increases the ability to stabilize the divides in SoS independently by reducing the complexity and dependencies between them.

While the Resist stage of DIRE is responsible for slowing the rate of capability loss, the Respond and Restore stages (collectively referred to as the Recover stage) of the DIRE framework (depicted in the bottom left and bottom right panels of Figure 8, respectively) are responsible for actively recovering any lost capability. These stages are really a continuum based on their respective time scales; the Respond stage is responsible for short-term recovery of capability while the Restore stage is concerned with longer-term repair and recovery issues. During the Respond stages, individual components in the SoS can be recovered to restore some of the overall SoS capability induced by a given disturbance. Additionally, it is possible that a portion of the SoS capability could be momentarily lost, as in the case of unrecoverable faults or targeted isolations. In cases such as these, the Recovery phase would apply a contingency—for example, rerouting power through a separate feeder to reestablish the power within the system or switching to alternative computational resources (i.e., servers) in an IT component. More abstractly, this could manifest itself as the use of an alternative capability that is on hand to meet a similar need (e.g., cellular hot spots to replace traditional IT connectivity). As the combined Recover phases continue, the Restore phase implements any longer-term, required repairs that would push the overall SoS capability beyond some minimum required baseline and potentially increase future overall capability and resilience. Such actions include the repair of maligned communications or wind-damaged assets within the failing components, bringing them fully back into operation.

In summary, the benefit in applying the DIRE approach in the abstract SoS context is the ability to visualize dependencies, not only of, e.g., a power system, but also the operating considerations and pathways to cascading failures in complex SoS. We have primarily discussed the DIRE framework in terms of what a resilient SoS *should* do in each of these stages. However, we reiterate that in order to incorporate resilience at the design stage or improve resilience of an existing system, it is necessary to have baseline capability metrics for each component in the SoS. Additionally, a catalog of potential mitigations and actions to stabilize the subsystem and restore capabilities is necessary to actually implement the desired resilience. This is further expanded in scope when considering a SoS, as now it is necessary to understand all dependencies in the SoS that can induce failures in other components and to understand capabilities in the abstract so

that the full array of Resist and Recovery actions can be realized.

V. CONCLUSION

This article covers a proposed basis for the measurement, and ultimately benefit, of planning for resilience within resilient control designs. The DIRE framework relates resilience to a baseline level of operation, denoted normalcy. A manifold is given to correlate adaptive capacity and margin to maneuver, based upon the physical constraints of operation, which are readily understood. It also provides the measurable impact to reflect both time and data effects from cyber-attacks to the control-system dynamics and, ultimately, the physical infrastructure facility that is monitored and controlled. To achieve resilience, therefore, the system should be decomposed into the appropriate distributed dynamics and controlled within those dynamics to prevent the propagation of destabilizing effects. Both design and operational measures can be used to develop and maintain an inherently resilient system.

Given the potential for drastic consequences, characterizing the full situational analysis requires an understanding of root causes and the potential for propagation of failures throughout a CI system. The root cause of a failure, whether from cyber- or physical attack or damaging storm, requires separate analysis, but can be decomposed into impacts throughout the DIRE framework. This framework was presented in terms of relevant benchmark metrics for both system and SoS contexts. In the context of SoS, the DIRE methodology can be understood as a framework to correlate overlaying factors that influence the operation and associated dependencies. Visibility of the conditions that can lead to a cascading failure, for instance, also allows design of potential mitigations. From a cybernetic root-cause context, correlating the complexity and factors influencing the success of an attack and the overall vulnerability are defined in the CVSS. The integration of these contributing metrics approaches can enable a more-comprehensive understanding of resilience in both design and operation.

REFERENCES

[1] M. Shinozuka, "Resilience of integrated power and water system," *Seismic Eval. Retrofit Life Time Syst.*, pp. 65–86, 2004.

[2] A. Boin and A. McConnell, "Preparing for critical infrastructure breakdowns: The limits of crisis management and the need for resilience," *J. Contingencies Crisis Manage.*, vol. 15, no. 1, pp. 50–59, Mar. 2007.

[3] L. Dueñas-Osorio and S. M. Vemuru, "Cascading failures in complex infrastructure systems," *Struct. Saf.*, vol. 31, no. 2, pp. 157–167, Mar. 2009.

[4] M. Rudner, "Cyber-threats to critical national infrastructure: An intelligence challenge," *Int. J. Intell. CounterIntell.*, vol. 26, no. 3, pp. 453–481, Sep. 2013.

[5] Y. Wang, C. Chen, J. Wang, and R. Baldick, "Research on resilience of power systems under natural disasters—A review," *IEEE Trans. Power Syst.*, vol. 31, no. 2, pp. 1604–1613, Mar. 2016.

[6] M. E. Beatty, S. Phelps, C. Rohner, and I. Weisfuse, "Blackout of 2003: Public health effects and emergency response," *Public Health Rep.*, vol. 121, no. 1, pp. 36–44, Jan. 2006.

[7] C. E. Colten, R. W. Kates, and S. B. Laska, "Community resilience: Lessons from new orleans and hurricane katrina," *CARRI Rep.*, vol. 3, pp. 2–4, Sep. 2008.

[8] T. Comes and B. Van de Walle, "Measuring disaster resilience: The impact of hurricane sandy on critical infrastructure systems," *ISCRAM*, vol. 11, pp. 195–204, May 2014.

[9] Y. Chen, L. Wen, and C. Ji, "A cascading failure during the 24 may 2013 great Okhotsk deep earthquake," *J. Geophys. Res., Solid Earth*, vol. 119, no. 4, pp. 3035–3049, Apr. 2014.

[10] R. Lee, M. J. Assante, and T. Conway, "Tlp: White-analysis of the cyber attack on the Ukrainian power grid-defense use case," in *Proc. Electr. Inf. Sharing Anal. Center (E-ISAC)*, 2016, pp. 1–29.

[11] R. M. Lee, M. Assante, and T. Conway, "Crashoverride: Analysis of the threat to electric grid operations," Dragos, Hanover, MD, USA, Tech. Rep., 2017.

[12] D. Bisson, "7 data breaches caused by human error: Did encryption play a role?" Venafi, Salt Lake City, UT, USA, Tech. Rep., 2020.

[13] E. Hollnagel, D. D. Woods, and N. Leveson, *Resilience Engineering: Concepts and Precepts*. Farnham, U.K.: Ashgate, 2006.

[14] A. M. Madni and S. Jackson, "Towards a conceptual framework for resilience engineering," *IEEE Syst. J.*, vol. 3, no. 2, pp. 181–191, Jun. 2009.

[15] R. Patriarca, J. Bergström, G. D. Gravio, and F. Costantino, "Resilience engineering: Current status of the research and future challenges," *Saf. Sci.*, vol. 102, pp. 79–100, Feb. 2018.

[16] C. G. Rieger, D. I. Gertman, and M. A. McQueen, "Resilient control systems: Next generation design research," in *Proc. 2nd Conf. Hum. Syst. Interact.*, May 2009, pp. 632–636.

[17] L. Fraccascia, I. Giannoccaro, and V. Albino, "Resilience of complex systems: State of the art and directions for future research," *Complexity*, vol. 2018, pp. 1–44, Aug. 2018.

[18] R. M. Lee, M. J. Assante, and T. Conway, "Analysis of the cyber attack on the Ukrainian power grid defense use case," SANS, Bethesda, MD, USA, Tech. Rep., 2016.

[19] B. J. Pierre, B. Arguello, A. Staid, and R. T. Guttromson, "Investment optimization to improve power system resilience," in *Proc. IEEE Int. Conf. Probabilistic Methods Appl. Power Syst. (PMAPS)*, Jun. 2018, pp. 1–6.

[20] *DOE Announces Investment for Resilience, Reliability of Nation's Energy Infrastructure*, Transmiss. Distrib. World, Overland Park, KS, USA, 2018.

[21] *DOE Invests $28 Million to Advance Cybersecurity of U.S. Energy Infrastructure*, Transmiss. Distrib. World, Overland Park, KS, USA, 2018.

[22] *INL Tool Helps Emergency Managers Understand Hidden Impacts of Disaster*, EDMS, Emergency Manage. Issues Special Interest Group, 2020.

[23] C. G. Rieger, "Notional examples and benchmark aspects of a resilient control system," in *Proc. 3rd Int. Symp. Resilient Control Syst.*, Aug. 2010, pp. 64–71.

[24] C. Rieger, C. Kolias, J. Ulrich, and T. R. McJunkin, "A cyber resilient design for control systems," in *Proc. Resilience Week (RWS)*, Oct. 2020, pp. 18–25.

[25] C. G. Rieger, "Resilient control systems practical metrics basis for defining mission impact," in *Proc. 7th Int. Symp. Resilient Control Syst. (ISRCS)*, Aug. 2014, pp. 1–10.

[26] C. Rieger and D. S. Naidu, "New techniques for implementing linear quadratic methods with aerospace and other industrial control applications," in *Proc. IASTED Int. Conf. Intell. Syst. Control*. Calgary, AB, Canada: ACTA Press, 2004, pp. 388–393.

[27] M. Panteli and P. Mancarella, "The grid: Stronger, bigger, smarter?: Presenting a conceptual framework of power system resilience," *IEEE Power Energy Mag.*, vol. 13, no. 3, pp. 58–66, May 2015.

[28] K. Young, R. Muehlhaeusser, R. Piggin, and P. Rachitrangsan, "Agile control systems," *Proc. Inst. Mech. Eng., D, J. Automobile Eng.*, vol. 215, no. 2, pp. 189–195, 2001.

[29] K. Schultz, "Towards agile control of ship auxiliary systems," in *Proc. 4th Int. Symp. Resilient Control Syst.*, Aug. 2011, pp. 154–157.

[30] C. G. Rieger, K. L. Moore, and T. L. Baldwin, "Resilient control systems: A multi-agent dynamic systems perspective," in *Proc. IEEE Int. Conf. Electro-Inf. Technol. (EIT)*, May 2013, pp. 1–16.

[31] *Resilience in Distributed Systems*, Infosys, Bangalore, India, 2019.

[32] T. V. Vu, B. L. H. Nguyen, Z. Cheng, M.-Y. Chow, and B. Zhang, "Cyber-physical microgrids: Toward future resilient communities," *IEEE Ind. Electron. Mag.*, vol. 14, no. 3, pp. 4–17, Sep. 2020.

[33] J. J. Grainger and W. D. Stevenson, *Power System Analysis*. New York, NY, USA: McGraw-Hill, 1994.

[34] C. J. Geankoplis, *Transport Processes and Separation Process Principles: (Includes Unit Operations)*. Upper Saddle River, NJ, USA: Prentice-Hall Professional Technical Reference, 2003.

[35] W. A. Conklin, "Security in cyber-physical systems," in *Proc. Workshop Future Directions Cyber-Phys. Syst. Secur.* Newark, NJ, USA: Gateway Center, 2009.
[36] A. Humayed, J. Lin, F. Li, and B. Luo, "Cyber-physical systems security—A survey," *IEEE Internet Things J.*, vol. 4, no. 6, pp. 1802–1831, Dec. 2017.
[37] S. Amin, G. A. Schwartz, and A. Hussain, "In quest of benchmarking security risks to cyber-physical systems," *IEEE Netw.*, vol. 27, no. 1, pp. 19–24, Jan. 2013.
[38] J. A. Wang, H. Wang, M. Guo, and M. Xia, "Security metrics for software systems," in *Proc. 47th Annu. Southeast Regional Conf.*, 2009, pp. 1–6.
[39] Y. Cheng, J. Deng, J. Li, S. A. DeLoach, A. Singhal, and X. Ou, "Metrics of security," in *Cyber Defense and Situational Awareness*. Düsseldorf, Germany: Springer, 2014, pp. 263–295.
[40] M. Pendleton, R. Garcia-Lebron, J.-H. Cho, and S. Xu, "A survey on systems security metrics," *ACM Comput. Surv.*, vol. 49, no. 4, pp. 1–35, 2016.
[41] D. J. Bodeau, R. D. Graubart, R. M. McQuaid, and J. Woodill, "Cyber resiliency metrics catalog," MITRE, McLean, VA, USA, Tech. Rep. MTR180450, 2018.
[42] M. Torgerson, "Security metrics for communication systems," in *Proc. 12th Int. Command Control Res. Technol. Symp.*, Newport, RI, USA, 2007, pp. 1–15.
[43] W. Jansen, "Research directions in security metrics," *J. Inf. Syst. Secur.*, vol. 7, no. 1, pp. 3–22, 2011.
[44] E. D. Vugrin, D. E. Warren, and M. A. Ehlen, "A resilience assessment framework for infrastructure and economic systems: Quantitative and qualitative resilience analysis of petrochemical supply chains to a hurricane," *Process Saf. Prog.*, vol. 30, no. 3, pp. 280–290, Sep. 2011.
[45] G. Cybenko, "Quantifying and measuring cyber resiliency," *Proc. SPIE*, vol. 9825, May 2016, Art. no. 98250R.
[46] A. M. Azab, P. Ning, E. C. Sezer, and X. Zhang, "HIMA: A hypervisor-based integrity measurement agent," in *Proc. Annu. Comput. Secur. Appl. Conf.*, Dec. 2009, pp. 461–470.
[47] MITRE. (2014). *Common Weakness Scoring System (CWSS)*. Accessed: Nov. 30, 2020. [Online]. Available: https://cwe.mitre.org/cwss/cwss_v1.0.1.html
[48] Forum of Incident Response and Security Teams. (2019). *CVSS User Guide*. Accessed: Nov. 9, 2020. [Online]. Available: https://www.first.org/cvss/user-guide
[49] National Cybersecurity FFRDC. *CVE—Common Vulnerabilities and Exposures (CVE)*. Accessed: Nov. 30, 2020. [Online]. Available: https://cve.mitre.org/
[50] NIST. *National Vulnerability Database*. Accessed: Nov. 30, 2020. [Online]. Available: https://nvd.nist.gov/
[51] F. Petit, D. Verner, D. Brannegan, W. Buehring, D. Dickinson, K. Guziel, R. Haffenden, J. Phillips, and J. Peerenboom, "Analysis of critical infrastructure dependencies and interdependencies," Argonne Nat. Lab.(ANL), Argonne, IL, USA, Tech. Rep. ANL/GSS-15/4, 2015.
[52] I. Eusgeld, C. Nan, and S. Dietz, "'System-of-systems' approach for interdependent critical infrastructures," *Rel. Eng. Syst. Saf.*, vol. 96, no. 6, pp. 679–686, 2011.
[53] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, "Catastrophic cascade of failures in interdependent networks," *Nature*, vol. 464, no. 7291, pp. 1025–1028, Apr. 2010.
[54] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani, "Epidemic processes in complex networks," *Rev. Mod. Phys.*, vol. 87, no. 3, p. 925, 2015.
[55] P. Crucitti, V. Latora, and M. Marchiori, "Model for cascading failures in complex networks," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 69, no. 4, Apr. 2004, Art. no. 045104.
[56] M. Korkali, J. G. Veneman, B. F. Tivnan, J. P. Bagrow, and P. D. H. Hines, "Reducing cascading failure risk by increasing infrastructure network interdependence," *Sci. Rep.*, vol. 7, no. 1, p. 44499, Apr. 2017.
[57] M. J. Alenazi and J. P. Sterbenz, "Comprehensive comparison and accuracy of graph metrics in predicting network resilience," in *Proc. 11th Int. Conf. Design Reliable Commun. Netw. (DRCN)*, Mar. 2015, pp. 157–164.
[58] M. Piraveenan, M. Prokopenko, and L. Hossain, "Percolation centrality: Quantifying graph-theoretic impact of nodes during percolation in networks," *PLoS ONE*, vol. 8, no. 1, Jan. 2013, Art. no. e53095.
[59] Q. Zhu and T. Basar, "Game-theoretic methods for robustness, security, and resilience of cyberphysical control systems: Games-in-games principle for optimal cross-layer resilient control systems," *IEEE Control Syst.*, vol. 35, no. 1, pp. 46–65, Feb. 2015.
[60] A. Alessandri and R. Filippini, "Evaluation of resilience of interconnected systems based on stability analysis," in *Critical Information Infrastructures Security*. Berlin, Germany: Springer, 2013, pp. 180–190.
[61] J. Kim and I. Dobson, "Approximating a loading-dependent cascading failure model with a branching process," *IEEE Trans. Rel.*, vol. 59, no. 4, pp. 691–699, Dec. 2010.
[62] M. Papic, K. Bell, Y. Chen, I. Dobson, L. Fonte, E. Haq, P. Hines, D. Kirschen, X. Luo, S. S. Miller, N. Samaan, M. Vaiman, M. Varghese, and P. Zhang, "Survey of tools for risk assessment of cascading outages," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, Jul. 2011, pp. 1–9.
[63] N. Ahmad, M. Chester, E. Bondank, M. Arabi, N. Johnson, and B. L. Ruddell, "A synthetic water distribution network model for urban resilience," in *Sustainable and Resilient Infrastructure*. Abingdon, U.K.: Taylor & Francis, Jul. 2020, pp. 1–15.
[64] *Inl Tool Helps Emergency Managers Understand Hidden Impacts of Disasters*, EMI SIG, 2020.
[65] T. R. McJunkin and C. G. Rieger, "Electricity distribution system resilient control system metrics," in *Proc. Resilience Week (RWS)*, Sep. 2017, pp. 103–112.
[66] T. Phillips, T. McJunkin, C. Rieger, J. Gardner, and H. Mehrpouyan, "An operational resilience metric for modern power distribution systems," in *Proc. IEEE 20th Int. Conf. Softw. Qual., Rel. Secur. Companion (QRS-C)*, Dec. 2020, pp. 334–342.

CRAIG RIEGER (Senior Member, IEEE) received the B.S. and M.S. degrees in chemical engineering from Montana State University, in 1983 and 1985, respectively, and the Ph.D. degree in engineering and applied science from Idaho State University, in 2008. He is currently the Chief Control Systems Research Engineer and a Directorate Fellow with the Idaho National Laboratory (INL), pioneering interdisciplinary research in next generation resilient control systems. The grand challenge provided an integrated research strategy to address the cognitive, cyber-physical challenges of complex control systems into self-aware, trust-confirming, and threat-resilient architectures. In addition, he has organized and chaired thirteen co-sponsored symposia and one National Science Foundation workshop in this new research area and authored more than 45 peer-reviewed publications. His Ph.D., coursework and dissertation focused on measurements and control, with specific application to intelligent, supervisory ventilation controls for critical infrastructure. He has 20 years of software and hardware design experience for process control system upgrades and new installations. He has also been a supervisor and a technical lead for control systems engineering groups having design, configuration management, and security responsibilities for several INL nuclear facilities and various control system architectures.

KEVIN SCHULTZ received the B.S. degree in mathematics and in electrical computer engineering and the M.S. and Ph.D. degrees in electrical computer engineering from The Ohio State University, in 2005, 2007, and 2010, respectively. He is currently a Project Manager and a Senior Staff Scientist with the Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory. His dissertation focused on modeling and analyzing collective behaviors in swarms and applying them as motivation in distributed control problems. His research interests include the control, characterization, and modeling of quantum systems, as well as distributed control and signal processing applications outside of quantum information. His research interests also include characterization and control problems in quantum information, neuromimetic approaches to distributed control problems, and applications of graph signal processing to emergence in distributed systems. As a principal investigator or a project manager, he has executed research funded by various sponsors, as well as numerous internal efforts.

THOMAS CARROLL received the B.Sc. degree in chemistry and in computer science and the M.Sc. and Ph.D. degrees from Wayne State University, Detroit, MI, USA, in 2001, 2007, and 2009, respectively. He is currently a Senior Cybersecurity Researcher with the Computational Science Division, Pacific Northwest National Laboratory (PNNL). As a Ph.D. Student, he received a NSF Graduate Research Program Fellowship award to study the incentive-centered design of distributed computer scheduling. He is responsible for more than 30 peer-reviewed publications, a book chapter, and two patents on topics of game theory, situation awareness, cyber security, smart grid, and vehicle systems. His research interests include the development and application of resiliency and zero-trust strategies to enterprise, cyber-physical, and cyber-vehicle systems. He is also the principle investigator of a project examining the cyber security and resiliency of high-powered electric vehicle charging infrastructure and another project designing, developing, and demonstrating zero-trust network access mechanisms to better defend cyber-physical systems in the electric power domain.

TIMOTHY MCJUNKIN received the M.Sc. degree in electrical and computer engineering from Utah State University. He is currently pursuing the Ph.D. degree with the Electrical Engineering Department, University of Idaho. He has been a Distinguished Researcher with the Idaho National Laboratory (INL), since 1999, with current research and development in resilient control of critical infrastructure, smart grid for renewable energy integration, cybersecurity, robotics and automation, intelligent systems, and acoustic based non-destructive examination. His has published more than 20 peer review journal articles, twp book chapter, and been awarded 13 patents on topics of computer systems, analytical chemistry instrument systems, industrial automation, smart grid, and nondestructive examination. Prior to joining INL, he was with Compaq Computer Corporation's Industry Standard Server Group, from 1994 to 1999, leading board level motherboard design of multiple server products. At Utah State, he was awarded a Rocky Mountain NASA Space Grant Consortium fellowship for his work on autonomous planetary vehicles. He has served an Adjunct Professor for the Electrical Engineering Department, Idaho State University.