

Constructing a Science of Cyber-Resilience for Military Systems

IST-153 Workshop on CYBER RESILIENCE (Position Paper Submission)

Alexander Alexeev, Diane S. Henshel, Karl Levitt, Patrick McDaniel,
Brian Rivera, Steven Templeton, and Mike Weisman —July 1, 2017

1. Cyber-Resilience in the Military Context

For over 50 years computer researchers and practitioners have sought to build systems that are secure; systems whose function and data could not be maliciously misused or influenced by an adversary. While substantial gains have been made in addressing security broadly, the complexity, size, malleability and rate of change of software and hardware systems have made achieving perfect security—at least for now—a practical impossibility. Exacerbating the situation are the ever-increasing zero-day attacks. This has led to a rethinking of efforts in countering malicious behavior that asks “*Can systems be designed such that they continue to function while compromised?*” This property, referred to as *cyber-resilience*, provides a means of operating in the face of malicious action.

System resilience is essential to military systems—now and into the future. The operating environment for military systems are increasingly congested and contested. These challenges apply both to the physical environment (including the electromagnetic environment) and cyberspace. Kinetic and digital military operations increasingly rely on computers and networked communications. Further, the uncertain and ambiguous nature of the future operating environments will require the Army to operate in unexpected and nontraditional roles—thus requiring systems to function outside the environments they were designed for.

These challenges lead to several broad requirements for resilience in the military context. First, military networks (which are increasingly complex and interdependent) must be able to operate while under constant attack in both cyberspace and the real world. Second, they must automatically determine when a problem has occurred and reconfigure or adapt to restore services automatically or notify the human if manual intervention is required. Lastly, systems must degrade gracefully to provide minimal services in the event of major failures.

In this position paper, we consider the elements of a *science of cyber-resilience* for military systems and identify connections to other fields of computation and broader sciences.

We ask what can be learned from other domains whose application make systems more resilient, and identify important open questions for the creation of a new Science of Cyber-Resilience. We begin in the next section by identifying other domains in which resilience is important and identify terminology and structures applicable to computer systems.

2. Learning from Other Domains of Resilience

Physical resilience is the ability of a physical material or body to resist mechanical deformation, and recover its original shape and structure after the removal of the stress, acting autonomously or in response to an outside force [Enn2012]. *Resilience involves the ability to resist a stress, minimize the adverse response to that stress, and efficiently and quickly recover, returning to a state as close to the original as possible in structure and/or function.* Resilience exists at multiple scales. At microscale, resilience is a property of a single physical, chemical or biological entity or organism. Physical resilience of a metal or psychological resilience of an individual is an example of microscale resilience [Wag1993]. At mesoscale, resilience is a function of interacting structures or organisms, usually at the same spatial or organizational scale. Team resilience is used to evaluate how well sports and business teams continue to perform effectively in the face of both internal and external stressors [Fle2013]. Group or community resilience helps a social organization preserve its structure and links in response to internal and external stressors, such as in-fighting or politics. At macroscale, resilience is a function of a complex system that includes structures or organisms interacting over multiple scales of time, space and/or organization. Resilience of biological and ecological systems reflect the ability of these systems to maintain functional homeostatic dynamics while internal resources are used and regenerated and external stressors alter biological and ecological demands on those resources. At larger scales, institutional rules and governmental policies may significantly modulate resilience of biological, ecological, and social systems, either preventing catastrophic scenarios or making the systems more vulnerable and less resilient [Fol2006]. Economic and financial systems are extensions of social systems. Economic resilience depends on and interacts with market stability in response to socioeconomic perturbations [Bri2009]. At

the macroscale, climate change resilience is determined by the ability of multidimensional systems (social, economic and environmental) to cope with, respond to, adapt to, and transform following a stressor or system of stressors associated with the climate change phenomenon (extreme climatic events, sea level rise, prolonged weather anomalies [drought, baseline temperature change]) [Bak2017]. Intuitively, resilience as seen in all of these domains has meaningful analogs to those in the cyber systems, and therefore the mechanisms we use to achieve them are fertile ground for identifying and measuring resilience of digital systems.

Resilience, then, is a physical property but it is also a process. One can make structural changes that imbue inherent resilience into the materials, components, structures, people, systems, and policies. We consider this to be *structural resilience*. However, the total resilience of any system, meso or macro, must include *active resilience*, the processes by which the material, component, individual organism, system respond, adapt to and recover from the imposed stress. Any consideration of resilience must also recognize that any system, and any component in a system, has both internal and external stresses at play as well as internal (or inherent) and external resources to use to mitigate the stress-imposed disturbance. Homeostasis occurs when a system is in equilibrium. A living or functional system is never static, and always has processes that utilize system resources. Homeostasis is the self-regulating process by which a (biological or ecological) system balances the uses of the available resources in response to internal and external needs and demands (i.e. stresses) without exceeding the limits of any particular resource to the point where the system shuts down or self-destructs in part or in whole. Similarly, cyber systems, whether independent or connected to the Internet of Things or other physical systems, have both structural resilience and active resilience features, which trade off in the attempt to keep the cyber systems in a functional homeostatic equilibrium. When structural resilience (like firewalls, antivirus, policies, and other intrusion detection systems) is strong, it provides a time and safety buffer for the active resilience (autonomous or human-based agility responses), whereas when structural resilience features are minimized, the active resilience must play a larger, more immediate role in both prevention of damage and recovery from any induced system alterations. Within cyber systems, internal stresses include resource constraints, while external stresses include such

factors as attackers, environmental extremes (e.g. temperature, humidity), physical shocks, and electromagnetic radiation, many of which are factors in military operations in the field.

3. Towards a Science of Cyber-Resilience

Understanding how cyber-resilience can improve mission success, in particular when that mission is challenged by an intelligent adversary, requires an understanding of how operational objectives can be degraded, and how systems can be designed to effectively mitigate impact to the mission. In this context resilience focuses on how to design cyber-systems that can operate effectively while under attack and validate measures of the resilience. We argue that validation of resilience measures will be best accomplished through a rigorous, formal study of resilience and security in general, particularly within the context of cyber-operations on the battlefield.

A science of cyber-resilience is a systematic undertaking to obtain and organize knowledge as testable explanations and predictions with respect to the ability of a cyber-system to automatically and reliably recover from events or situations that cause the system to operate outside of defined mission objectives, both security and operational, and in accordance with temporal and other mission constraints.

At its core, science is the art of discovery. It allows us to describe, define, investigate and ultimately try to understand a domain, and will require a defined set of measurements, measures, metrics, and indicators: time-independent, quantified properties that are calculated, not measured [Rag1995]. By increasing our understanding we are able to identify problems, define causes, create solutions and make improvements in quality. A science of cyber-resilience will not in itself create change, however the knowledge gained and the application of science to cyber-resilience can greatly improve mission operations.

Metrics

Many metrics for cyber-systems, in particular in the area of fault-tolerance, have been published; however, additional metrics focused on resilience, rather than on general

operational awareness, will be needed. Characteristics of defensible metrics are comprehensiveness, understandability, feasibility, uniqueness, simplicity, accuracy, timeliness, and repeatability. Metrics can be based on system performance of strategic objectives.

Time is a critical issue on the battlefield and will be a key measure for resilience. How long until degradation has been detected, how quickly a system can recover, and how long a cyber-system operates outside of stated objective are important measures.

External factors must also be considered in evaluating resilience. Examples include scope, complexity, cost and other economic factors. Highly resilient solutions may be possible, but might be infeasible due to cost, time to implement or unrealistic complexity. Such solutions may not be successfully implemented. In [Wof2015] the idea of cyber value-at-risk is discussed as a way of evaluating the utility and need of cyber resilience.

The reliability of a metric must also be considered. If the supporting measure is unreliable, perhaps due to unreliable sensing of a system state or uncertain detection of an attack, we have an increased risk of taking inappropriate recovery actions. Risk will be an important factor for resilience, for example determination of the risk of an attack to the mission. Reliability of resilience measures, but also in determining the importance and prioritization of developing and initiating resilience measures.

Overall a science of cyber-resilience will provide useful information to characterize the security and operational state of the cyber-system, but ultimately, "*Resilience Indicators*" are needed. These will provide the ability to compare the resilience of different cyber-systems or system components. For example, life-cycle status of supporting operating systems can be used as an indicator of the risk to a system or to suggest actions to reduce risk. Resilience indicators can be used to provide warnings or help signal that the resilience requirements of a system are approaching or are below required objectives. Not focusing directly on measures or metrics as our goals, helps avoid what is now known as Goodhart's Law. Economist Charles Goodhart observed in 1975 that "... when a measure becomes a target, it ceases to be a good measure" [Goo1981]. This law essentially captures the reality that optimization of a system to a metric that approximates a phenomenon can rob the metric of its assessment value.

The effectiveness and correctness of resilience metrics, indicators, derived principles and techniques must be validated. This can be done both experimentally and empirically. We need to verify that, for a given cyber-objective, the metrics provide useful, correct, timely, and reliable information to support taking recovery actions. For example, recovery may be insufficient, response too slow or not effecting restoration of resources needed to meet mission objectives. Recovery might also be disproportionate, causing problems for other systems or generating excessive cost. Of particular concern is over-response to transient events, where the resilience mechanisms are taking action to recover from problems that self-resolve faster than the system can handle them. Resilience techniques may also be incorrectly deployed, not only failing to mitigate the problem, but causing degradation of other parts of the system. Application of a resilience measure may also inadvertently decrease security. In the parlance of security, the operations in support of resilience may themselves be subject to attack, thus potentially increasing the overall attack surface [Man2004] of the system. In order to build and manage resilient cyber-systems, all elements of resilience, including the resilience of the cyber defense teams, must be addressed in an overall assessment of resilience.

Principles

Our work is motivated in part by the ongoing effort to develop a *science of cyber-security*. Schneider [Sch2011] uses succinctly stated laws to enable a system to be designed and reasoned about as a layering of abstractions, but does not address adaptive defense or resilience, and focuses primarily on methods of protection, defense and classification of attacks, primarily those impacting confidentiality. While knowing the type of attack may be useful, the impact of the attack to the mission is, in our view, equally important for resilience, particularly when dealing with unknown threats. The goal of resilience is not to just stop the attack, but to support recovery to an acceptable level of operation while under attack; how the attack affects the cyber-system is of greater value. This suggests that, while overlapping with the ongoing work on a science of cyber-security, a science of cyber-resilience will focus on different principles. A related area, a science of cyber decision-making [Mcd2016], also supports our sought after science of cyber-resilience.

A science of cyber-resilience will require developing a set of principles such as redundancy, diversity, modularity, visibility, managing connectivity, managing change, feedback, reducing risk, and innovation. Dependencies between cyber-system, components, and mission objectives must also be considered. The science will also need a set of principles of how attacks/compromises impact operational and security objectives. This requires that we are able to understand and characterize, clearly and effectively, all mission objectives, including functional, time, strategic, safety, economic, security, risk and recovery objectives.

A system's resilience can be characterized both at the component level and as a holistic system. The individual components support the resilience of the system, and the system supports the resilience of the components. The reliability mechanisms of system components may improve the resilience of each other, while in other cases, such as when the resilience mechanisms compete for a common resource, increasing resilience in one component can increase the risk that another component may not retain resilience objectives.

As indicated above, resilience is both structural and active. Structural resilience is a function of the architecture and interaction of the components of a system. The broad peer based communication in an ad hoc mesh network provide structural resilience, not directly through redundancy, but by providing a mechanism by which the individual nodes can find communication paths. Other structural mechanisms can support resilience by detecting problems faster, slowing the spread of an attack, or allowing the system to respond faster. Active resilience can be reactive or adaptive and may involve dynamic-control. Reactive resilience directly mitigates the impact of an attack (or other problem), while adaptive resilience intelligently alters operations to mitigate the problem and minimize the impact of future problems. Adaptive resilience can enhance the structural resilience of a system.

Defining resilience for cyber-security objectives can be more abstract than other objectives. With respect to the traditional CIA triad (confidentiality, integrity, availability), availability is the easiest to understand. Increasing availability or decreasing availability requirements need to be balanced against the resource needs and risk associated with the loss of CIA across all mission needs. Resilience for integrity will include methods to repair damage to suspect information, which will require both the ability to determine that the information may

have been altered and a means to correct the damaged information or a way to alter the mission to use alternative sources. Resilience for confidentiality may be the hardest to manage. Like integrity, we must know that the information has been disclosed. In response, the utility of that information (to the attacker) must be reduced, if not eliminated. In general, this would trigger reconfiguring relevant aspects of the cyber-system or mission such that the information is no longer useful (e.g. changing passwords or rerouting). While these agility responses may not always be feasible, characterizing the different types of information and how their disclosure can affect a mission will support the design of systems with increased resilience to loss of confidentiality. An alternative model for framing the components of security (other than CIA) may be useful in studying cyber-security resilience.

Improving cyber-resilience requires creation and use of metrics and resilience indicators effective for assessment the design, development, and implementation of resilience methods, both technical and procedural. A science of resilience needs a framework to organize the underlying principles of cyber-resilience. This requires a resilience analysis process, a variety of techniques for measuring resilience, a taxonomy to systematize the elements of cyber-resilience, and a framework for developing and evaluating resilience metrics. Efforts to study resilience such as that done by Sandia and Argonne National Labs [Wat2014, Car2012], can provide inspiration.

4. Conclusions

In this paper, we have highlighted some areas of investigation for a new science of cyber-resilience. We argue that we need to look to the broader scientific community and look to biological, environmental, and physical systems to see how they achieve resilience. In this way, we can develop a principled approach to measuring how systems can fight through malicious action, and therein provide safer, more reliable military systems. We can draw inspiration from the bubbling but promising effort towards a science of security.

Bibliography

- [Bak2017] Bakkensen, L. A., Fox-Lent, C., Read, L. K. and Linkov, I. (2017), "Validating Resilience and Vulnerability Indices in the Context of Natural Disasters". *Risk Analysis*, 37: 982–1004. doi:10.1111/risa.12677.
- [Bri2009] Briguglio, L., Cordina, G., Farrugia, N., & Vella, S. (2009). "Economic vulnerability and resilience: concepts and measurements". *Oxford development studies*, 37(3), 229-247. doi:10.1080/13600810903089893.
- [Car2012] Carlson, L and Bassett, G and Buehring, W and Collins, M and Folga, S and Haffenden, B and Petit, F and Phillips, J and Verner, D and Whitfield, R. "Resilience: theory and applications". Argonne National Laboratory. Argonne, Illinois, USA, 2012. <http://www.ipd.anl.gov/anlpubs/2012/02/72218.pdf>
- [Enn2012] Ennos, Roland. "Solid biomechanics". Princeton University Press, 2012, 9-10.
- [Fle2013] Fletcher, D., and Sarcar, M. "Psychological Resilience. A Review and Critique of Definitions, Concepts, and Theory". *European Psychologist* (2013): 18(1),12-23. doi:10.1027/1016-9040/a000124.
- [Fol2006] Folke, C. (2006). "Resilience: The emergence of a perspective for social–ecological systems analyses". *Global environmental change*, 16(3), 253-267. doi:10.1016/j.gloenvcha.2006.04.002
- [Goo1981] Goodhart, Charles AE. "Problems of monetary management: the UK experience." *Monetary Theory and Practice*. Macmillan Education UK, 1984. 91-121.
- [Man2004] Manadhata, Pratyusa, and Jeannette Marie Wing. "Measuring a system's attack surface." (2004). <https://www.cs.cmu.edu/~wing/publications/tr04-102.pdf>
- [McD2016] Patrick McDaniel and Ananthram Swami, The Cyber Security Collaborative Research Alliance: Unifying Detection, Agility, and Risk in Mission-Oriented Cyber Decision Making. *CSIAC Journal, Army Research Laboratory (ARL) Cyber Science and Technology*, 5(1), December, 2016.
- [Rag1995] Ragland, Bryce. "Measure, Metric, or Indicator: What's the Difference?." *Crosstalk* 8.3 (1995): 29-30.
- [Sch2011] Schneider, Fred B. (2011) "Blueprint for a Science of Cybersecurity" <https://www.cs.cornell.edu/fbs/publications/SoS.blueprint.pdf>

[Tri2011] Trimintzios, P. "Measurement Frameworks and Metrics for Resilient Networks and Services." Discussion Draft. European Network and Information Security Agency (2011). https://www.enisa.europa.eu/publications/metrics-tech-report/at_download/fullReport

[Wag1993] Wagnild, Gail M.; Young, Heather M., "Development and psychometric evaluation of the Resilience Scale", *Journal of Nursing Measurement*, Vol 1(2), 1993, 165-178.

[Wat2014] Watson, Jean-Paul, et al. "Conceptual framework for developing resilience metrics for the electricity, oil, and gas sectors in the United States." Sandia Nat. Lab., Albuquerque, NM, USA, Tech. Rep. SAND2014-18019 (2014).
<https://cfwebprod.sandia.gov/cfdocs/CompResearch/docs/EnergyResilienceReportSAND2014-18019o.pdf>

[Wof2015] World Economic Forum, "Partnering for Cyber Resilience: Towards the Quantification of Cyber Threats", January 2015.

Author Contact Information

Alexander Alexeev, PhD

Lecturer in Data Analysis and Modeling, Indiana University
School of Public and Environmental Affairs
1315 E 10th Street, Bloomington, IN 47405
E-mail: aalexeev@indiana.edu
Phone: 812-856-3294

Diane S. Henshel, PhD

Associate Professor, Indiana University
School of Public and Environmental Affairs
1315 E 10th Street, Bloomington, IN 47405
Email: dhenshel@indiana.edu
Phone: 812-855-4556

Karl Levitt, Ph.D.

Professor of Computer Science
College of Engineering, University of California, Davis
1 Shields Ave. Davis, CA 05616
Email: levitt@cs.ucdavis.edu
Phone: 650-799-6474

Patrick McDaniel, Distinguished Professor,
School of Electrical Engineering and Computer Science,
The Pennsylvania State University,
360A Information Sciences and Technology Building,
University Park, PA 16802-6823
Email: mcdaniel@cse.psu.edu
Phone: 814-863-3599
URL: <http://www.patrickmcdaniel.org>

Brian Rivera, Ph.D., CISSP, CEH, Chief (A),
Network Science Division
Chief, Tactical Network Assurance Branch,
U.S. Army Research Laboratory
2800 Powder Mill Road, Adelphi, MD 20783
Email: brian.m.rivera.civ@mail.mil
Phone: 301-394-2298

Steven Templeton, Ph.D. (9/17), CISSP
Computer Security Research Lab
Department of Computer Science,
University of California, Davis
One Shields Avenue, Davis CA 95616
Email: sjtempleton@ucdavis.edu
Phone: 530-554-2779

Mike Weisman, Ph.D.,
Network Science Division, Network Security Branch,
U.S. Army Research Laboratory, ICF International
2800 Powder Mill Road, Adelphi, MD 20783
Email: michael.j.weisman2.ctr@mail.mil
Phone: 301-394-1237