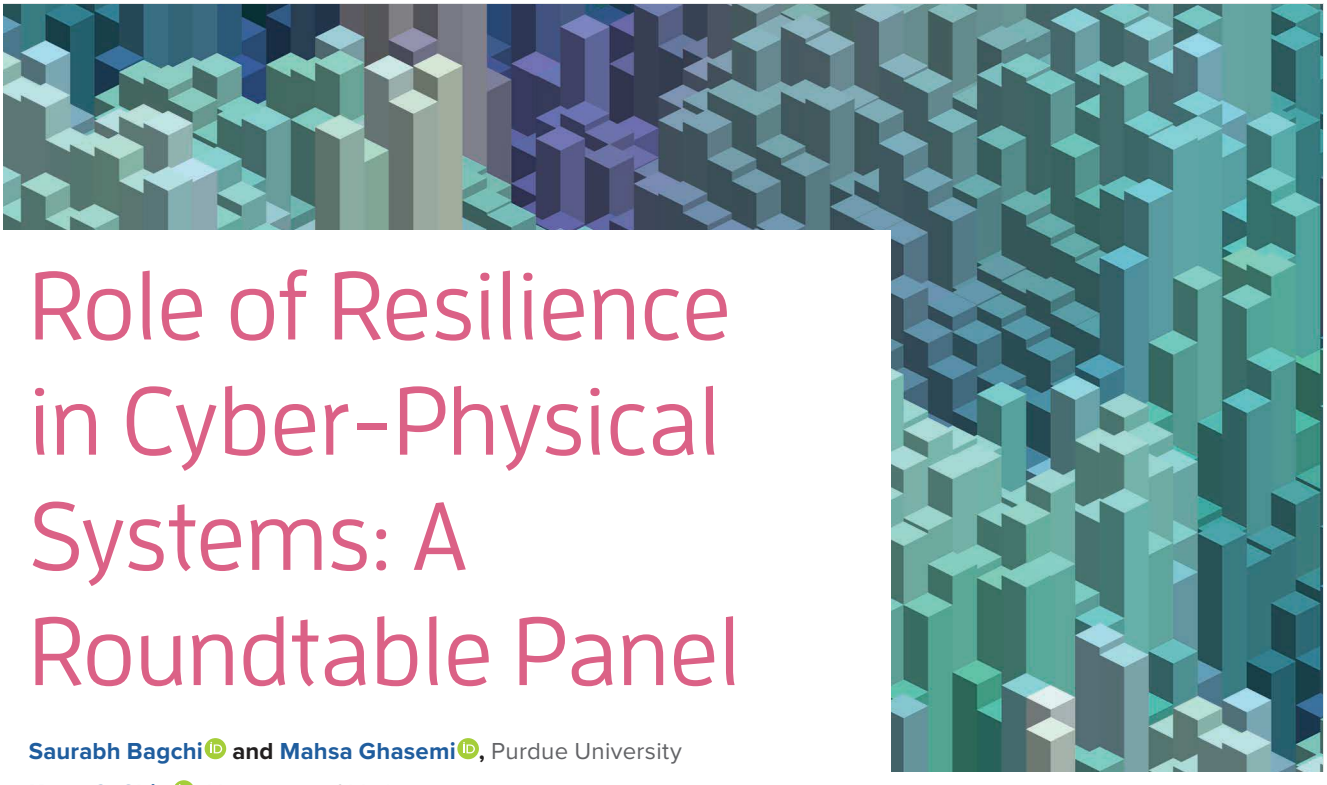

Role of Resilience in Cyber-Physical Systems: A Roundtable Panel

Saurabh Bagchi and **Mahsa Ghasemi**, Purdue University
Kang G. Shin, University of Michigan
Nalini Venkatasubramanian, University of California, Irvine
Dongyan Xu, Purdue University
Saman Zonouz, Georgia Institute of Technology

The panel was held on 14 November 2023 at Purdue University as part of a Grand Challenges in Resilience Workshop sponsored by the U.S. National Science Foundation and organized by our center, the Center for Resilient Infrastructures, Systems, and Processes (CRISP).

I (Bagchi) organized the panel and moderated, and I (Ghasemi) acted as the scribe and in follow-on work, worked with Saurabh on writing the article. The other authors were panelists and modified the written material, in some cases significantly, of their verbal comments at the panel. We set up the panel with the following questions to serve as scaffolding for the discussion. These were shared with the panelists (see Figures 1 and 2) ahead of time:

- What is one or a few examples of resilient cyber-physical systems (CPSs) today?
- What is one or a few examples of a lack of resilience, for example, hack/compromise, security attack, and unpredictable interaction in CPSs, that has/have hurt us?
- What are some principles for creating resilience in CPSs? Likewise, what are some antiprinciples for resilience in CPSs?
- Of the principles in the previous question, which ones are technologically feasible but economically infeasible; for instance, due to a lack of economic or policy incentives?

Digital Object Identifier 10.1109/MC.2024.3394108
Date of current version: 26 June 2024

0018-9162/24©2024IEEE

EDITOR **DIMITRIOS SERPANOS**
CTI DIOPHANTUS and University of Patras; serpanos@computer.org

BAGCHI: The positive examples of resilience in CPS often show up through the syndrome of "the dog that did not bark." For example, when the Los Angeles Aqueduct and Dam survived with little damage from multiple earthquakes in the 20th century or the telecommunication infrastructure in Japan largely stayed intact through the Great East Japan Earthquake of 2011, those are archetypes of resilient CPS, incorporating resilience in the cyber side and in the physical (or infrastructure) side. In the ongoing Ukraine–Russia war, Ukraine's power grid has withstood the destruction surprisingly well, preventing large-scale blackouts. This is particularly noteworthy because back in 2015, malware had severely disrupted the Ukrainian power grid, and thus it is only reasonable to assume that learning lessons from prior failures has, in this case, led to dramatic improvements.

Sometimes, a successful case of resilience is harder to spot due to the still-newsworthy headlines of the devastating impact of failures. A case in point is the ALERTCalifornia wildfire detection system. The program trained AI (artificial intelligence) to detect smoke and other early indications of fire on a feed from a network of more than 1,050 cameras placed in forests across the state. When the system spots something, it alerts the local fire department via text message. In the first two months of its deployment (2019–20), the system had correctly identified 77 fires before any 911 calls came in. Here, resilience comes in the form of redundant cameras, a robust use of AI, and a resilient communication infrastructure—a typical case of the cyber and the physical elements working together to ensure resilience. Nevertheless, stories of devastating wildfires in California leave a much more lasting impression in the minds of the broad public.

We can derive several principles from the positive examples. The design challenge and the engineering challenge lie in the instantiation of these principles for specific application contexts. For example, consider defense in depth; here, for a CPS that does wildfire detection and mitigation, one would need thermal protection for the sensors on the ground and some sensors at a higher level, like on a tree canopy (incidentally, the ALERTCalifornia system does follow these principles). The same principle when applied to a CPS that controls the environment in a manufacturing facility using smart sensors and actuators would need sanity checks for the machine learning algorithm that controls the temperature and the humidity as well as a fallback simple control algorithm.

We can also derive several antiprinciples from the negative examples. A perfectly engineered system will be nonresilient, or synonymously, fragile, if it is very sensitive to the operating conditions. An extreme example is a drone that flies efficiently under "normal" wind conditions, but under "abnormal" but still possible wind conditions, it crashes. Another antipattern is to assume that general human users

FIGURE 1. From left to right: panelists Dongyan Xu, Nalini Venkatasubramanian, Kang Shin, and Saman Zonouz.

FIGURE 2. Panel moderator Saurabh Bagchi addresses attendees.

will be security conscious. They are often in a hurry to use the CPS or are not trained enough to make critical security decisions. A simple example is that a large, distributed denial-of-service (DoS) attack, called the *Mirai Botnet* (2006), happened using IoT (Internet of Things) devices, starting from the fact that their initial default usernames and passwords were left unchanged. A simple design that would have enforced that the initial password be changed upon first use would have avoided this.

The broad point is that the patterns, and antipatterns, for resilience must be in the context of the physical and the cyber operating conditions for the system. And one must "mind the gap" between the design principles and their instantiation in the implementation.

WHAT IS ONE OR A FEW EXAMPLES OF RESILIENT CPSs TODAY?

ZONOUZ: I can mention two examples of power grids and airplanes. Power grid is a large-scale distributed system with remote heterogenous components (not fully securable, under control/contained) and lots of redundancies ("N-1 contingency proof") and is operable without "modern" cyber capabilities [sensors/PMUs (phasor measurement units), SCADA (supervisory control and data acquisition) and so on]. Airplane is a small contained and isolated system with lots of redundancies, has no security within its control network (other than network separation from entertainment/screens and other noncritical functionalities), and is operable without "modern" capabilities (for example, they have sextants on board in case GPS does not work).

The common features of these two systems are 1) redundancy (which has long been proven to be linked to resilience) and 2) operability, even without "modern/cyber" capabilities. The latter is consistent with the fact that most intentional disruptions (attacks) follow cyber vectors.

WHAT IS ONE OR A FEW EXAMPLES OF A LACK OF RESILIENCE IN CPSs THAT HAS HURT US?

ZONOUZ: One example is the 2003 power grid blackout, which was due to the lack of timely situational awareness (did not allow early detection, that is, required for practice recovery). Another one is the 9/11 malicious incidents against avionics, which included physical attack vectors.

Note, from my examples earlier, the most resilient ones now are the ones that hurt us the most in the past due to lack of resilience at the time. There have been lots of efforts/initiatives in both domains by the government and industry since 2001 [for example, DHS (the U.S. Department of Homeland Security), TSA (the Transportation Security Administration), and so on] and 2003 [NERC-CIP (North American Electric Reliability Corporation Critical Infrastructure Protection) and so forth].

WHAT ARE SOME PRINCIPLES OR ANTIPRINCIPLES FOR RESILIENCE IN CPSs THAT CAN BE EXTRAPOLATED FROM TODAY'S SYSTEMS? GROUND THESE PRINCIPLES IN SPECIFIC APPLICATION CONTEXTS.

ZONOUZ: Resilience requires, among other things, 1) diversity, such as N-version programming in software and various technologies in hardware (for instance, various sensing modalities/technologies); 2) "no unnecessary smartness—" interoperability (for example, use of common/many protocols) and remote connectivity [VPN (virtual private network), WI-FI, and so on] are often talked about as a "plus," but they can severely hurt resilience since they provide remote attack vectors if not designed carefully to keep the trusted computing base minimally small; 3) knowledge of system dynamics when it enters the safe–unsafe boundary regions (for instance, drone next to a wall or with a propeller half broken) to allow reactive response when the system is already in unsafe (but still recoverable) state (for example, it has not crashed yet). Reactive response is the last resort if proactive recovery is not accomplished. 4) The resilience solution should not be "gameable" by malicious adversaries—if the solution responds to each incident using a fixed optimal action, it can be gamed by the adversaries, who can take the same action over and over again causing DoS, for instance, if you reboot every time, the same vulnerability is exploited (without patching it), nothing stops the attacker from exploiting it again. 5) Human resilience: fully automated resilience is often not reliable since humans are in the loop (or on the loop). To balance this, one must consider that humans also make mistakes and can be malicious (insider attacks). For example, in airplanes after 9/11 and the Germany incident, they decided to always have two people in the cockpit, that is, having redundancy with humans.

OF THE PRINCIPLES MENTIONED EARLIER, WHICH ONES ARE TECHNOLOGICALLY FEASIBLE BUT ECONOMICALLY INFEASIBLE (THE COMPLEMENT SET IS TECHNOLOGICALLY AND ECONOMICALLY FEASIBLE)?

ZONOUZ: Most of the solutions are technologically feasible, I believe, except when it comes to resource-limited settings with hard, real-time constraints. From an economical perspective, the application and extent of the aforementioned solutions in practice would depend on many factors: the safety criticality of the operation, resilience against what (malicious or accidental failures), and the financial or other damages in case of failure.

WHAT IS ONE OR A FEW EXAMPLES OF RESILIENT CPSs TODAY?

KANG SHIN: The following are examples of resilient CPS today: medical

devices; power and communication grids (Starlink satellite communication); factories and hospitals; national infrastructure; military weapons, platforms, and C4I (command, control, communications, computers, and intelligence); transportation platforms; and infrastructure.

WHAT ARE SOME EXAMPLES OF WHERE THE LACK OF RESILIENCE IN CPSs HAS HURT US?

KANG SHIN: The following are some examples: disaster recovery systems (failure of communication and coordination), transportation platforms and infrastructure thereof (need to support both autonomy and coordination of individual platforms), and IoT devices and their coordination (they are useful for monitoring due to the possibility of pervasive deployment; however, they are easy to hack and difficult to coordinate at scale).

Here is a partial list of complex CPS failures: Denver baggage handling system [US$300 million (1990s)], power blackout in New York (2003), Ariane 5 [US$370 million (1996)], Mars Pathfinder (1997), Mars Climate Orbiter [US$125 million (1999)], the Patriot Missile (1991), USS *Yorktown* (1998), Therac-25 (1985–1988), London Ambulance System [£9 million (1992)], pacemakers (500,000 recalls during 1990–2000), numerous computer-related incidents with commercial aircraft (http://www.rvs.uni-bielefeld.de/publications/compendium/incidents_and_accidents/index.html).

WHAT ARE SOME PRINCIPLES FOR RESILIENCE IN CPSs THAT CAN BE EXTRAPOLATED/IMPROVED FROM TODAY'S SYSTEMS?

KANG SHIN: Creating resilient CPS needs the following design principles: 1) advancing cybersecurity, which has as the two primary elements, detection of errors followed, by damage assessment and recovery. The recovery can involve reconfigurability of systems. It is important that there is high coverage for each step. There needs to be latency reduction in computation systems and communication networks. 3) One can possibly learn from biological systems, like immune systems.

WHAT IS ONE OR A FEW EXAMPLES OF RESILIENT CPSs TODAY?

VENKATASUBRAMANIAN: At the outset, one must lay down that resilience is a cross-layer concept: physical infrastructure resilience, component/device resilience, network resilience, software resilience, and societal resilience. The aforementioned question, in my view, is therefore a bit underspecified. The relevant question is one of "resilient under what circumstances and at what scale,"—one can refer to this as contextual resilience. Under fairly localized events and failures, many of our current systems, including intelligent transportation systems (air traffic, transit systems, road networks) and energy systems, are resilient. They are able to tolerate limited levels of (specific) failures and degrade gracefully by providing a reduced but adequate level of service. For example, interairline arrangements can reroute travelers with alternate airline partners. Public health and health-care systems have built-in resilience to small events (the annual flu) and limited surges. It is when large disruptions or catastrophic events occur (COVID-19, flash floods, earthquakes) and cause unexpected damages to critical lifelines that these same systems and services are found to be not resilient.

WHAT IS ONE OR A FEW EXAMPLES OF LACK OF RESILIENCE IN CPSs THAT HAS HURT US?

VENKATASUBRAMANIAN: Rather than specific examples, let me abstract out some recurring patterns. In several CPS platforms, preparedness is limited to a very localized scope of events (fire in a building, failure of traffic lights). Preparedness across agencies that must address different pieces of the puzzle under larger failures is missing, leading to a lack of resilience.

The cost of failure and the time to avert the failure are important factors to consider. For example, the aviation industry uses formal verification methods to explore possibly anomalous situations, including for providing timeliness guarantees. Other systems must do the same. Autonomous vehicles that gracefully exit from traffic and stop when faced with uncertain situations, rather than taking more complicated decisions that may lead to failures, are desirable.

WHAT ARE SOME PRINCIPLES OR ANTIPRINCIPLES FOR RESILIENCE IN CPSs THAT CAN BE EXTRAPOLATED FROM TODAY'S SYSTEMS? GROUND THESE PRINCIPLES IN SPECIFIC APPLICATION CONTEXTS.

VENKATASUBRAMANIAN: I distill these principles into two categories, depending on when they are needed.

For proactive resilience as well as preparedness, 1) decentralized solutions are critical to prevent a single point of failure, but purely decentralized solutions lack globalized awareness and are therefore not always desirable; 2) hierarchy is critical to address scale and provides levels at which decision support can be infused; 3) redundancy of critical resources and associated data are important, while keeping in mind that not everything needs to be equally redundant; 4) monitoring is critical, but the approaches must be cognizant of security/privacy implications, especially with mission-critical and sensitive data; 5) AI techniques that use a priori data (during nondisasters) to learn patterns of CPS use and behavior must be resilient to bias that most of the data are from normal operation.

For reactive resilience as well as after failure happens, 1) it is important to have timely and accurate situational awareness—as that can help with decision making—to recover from the failure.

A challenge here is that developing situational awareness after the failure has to deal with accuracy issues since information may be lost, unavailable, or erroneous; 2) distributed decision making will be needed, under various constraints, including time and information flow constraints; 3) the distributed decision making should lead to a well-prioritized set of response strategies; 4) the resource allocation decisions need to be revisited postfailure to aid in the recovery.

It is also important to design fault-tolerant solutions that are dual use; that is, they work seamlessly under nonfailure situations.

Regarding antiprinciples for resilience, one size fits all does not work. This means that not all failures are equivalent in terms of impact on lives and property, so solutions must not be either. Besides, the human should not be considered merely as an impacted user or as a passive observer. In several CPSs, humans play roles that intrinsically affect CPS functioning as the designer or the operator/maintainer of a service. Therefore, possible disruptions due to interactions between the human (in all these roles) and the CPS must be anticipated.

FURTHER COMMENTS

VENKATASUBRAMANIAN: The lack of fine-grained sensing instruments in many forms of infrastructure complicates the detection (where and how) of disruptions. In above and underground infrastructures, we need to determine where to instrument—further research is required to understand how this decision should be made to handle operational and extreme situations. While this problem bears similarity to the traditional sensor-placement problem, we should note that the importance of sensing in different locations is not equal since the consequences of failures in different locations are not equal. Therefore, we need to develop CPS and infrastructure design solutions that incorporate consequence/impact level, often requiring iterative approaches, active learning, and solving inverse problems.

FIGURE 3. A photo of the StarCraft robotic food delivery vehicle on the campus of Purdue University as it moves deftly to avoid even unpredictable obstacles.

WHAT IS ONE OR A FEW EXAMPLES OF RESILIENT CPSs TODAY?

XU: Starship robots (see Figure 3) on the Purdue campus are resilient under different missions, payloads, and transportation variability while being courteous. (Starship robots are seen all around Purdue's campus at all times of day and night doing food delivery to campus buildings and dorms.) They are designed as level-4 autonomy, and the human-on-the-loop aspect only triggers in critical conditions.

The biggest e-commerce company is resilient in terms of being weather-proof and pandemic-proof, having adaptation in the supply chain, and operating sustainably in its cloud services.

WHAT IS ONE OR A FEW EXAMPLES OF LACK OF RESILIENCE IN CPSs THAT HAS HURT US?

The meltdown of one of the U.S. major airlines during the 2022 holiday season due to winter storm. The root fundamental cause has been an outdated IT system not meeting the needs of employees and physical operations under extreme weather. Moreover, their point-to-point network architecture, as opposed to the hub-and-spokes architecture, is more robust in normal conditions but turns out to be vulnerable in the face of wide-area winter weather with geographically distributed employees.

WHAT ARE SOME PRINCIPLES OR ANTIPRINCIPLES FOR RESILIENCE IN CPSs THAT CAN BE EXTRAPOLATED FROM TODAY'S SYSTEMS? GROUND THESE PRINCIPLES IN SPECIFIC APPLICATION CONTEXTS.

A four-step methodology for CPS resiliency that has been applied to a variety of applications, such as autonomous vehicles [for example, research effort under ONR (U.S. Office of Naval Research) RHIMES] and manufacturing [for instance, research effort under DOE (U.S. Department of Energy) CyManII]:

1) modeling of and mapping between cyber and physical components of a subject CPS, 2) exploration and vetting of the overall CPS based on the model and mapping in "1)," 3) demonstration and assessment of weaknesses/vulnerabilities discovered in "2)" to quantify the risk and consequences (for decision making), 4) mitigation and system hardening, based on assessment result from "3)," with the most cost-effective mitigation decision/action to the proportion of the risks and consequences identified.

FURTHER COMMENTS

Human users and operators of CPS should be considered explicitly to detect vulnerabilities. Humans can be the weakest links; hence, more attention to their modeling and interaction is required through multidisciplinary solutions.

ADDITIONAL DISCUSSIONS

BAGCHI: There is an inherent latency in the detection of disruptions. Does CPS make this latency a less significant issue since the physical component has a long latency (for motion and so on), buying us some more time for detection and recovery?

ZONOUZ: The answer depends on the dynamics of the system, whether it is fast or slow. Sometimes, in the case of a physical system, a very fast response in the order of a few milliseconds is needed, for example, a synchronous power generator. However, compared to cyber components, the physical component is usually slower. In the case of cyber component, sometimes, the companies or designers are hesitant to apply security layers since it defeats real-time, fast responses.

SOMALI CHATERJI: The companies or designers may prefer not to accessorize/overengineer systems with too many components due to cost-related concerns, for instance, the number of sensors on (semi-) autonomous cars. Nevertheless, sometimes those additional components are needed to build trust in the system. We should investigate the tradeoff and interplay of economics with resiliency when dealing with private commercial sectors.

XU: Overengineering may offer rich functionality but increase a CPS's attack surface; a minimal design may lead to poor functionality but achieve better security. We advocate a more balanced CPS engineering methodology that balances functionality, security, and resiliency. Another important factor in many CPS operations is the human factor, which deserves more attention during the full lifecycle of CPS development and operation.

This article summarized the discussion from a panel in November 2023 and captured the distinct perspectives of researchers in various aspects of CPSs. This should hopefully trigger deep thoughts and follow-on action on how to make CPSs more resilient.

SAURABH BAGCHI is a professor in the Elmore Family School of Electrical and Computer Engineering and the Department of Computer Science, Purdue University, West Lafayette, IN 47907 USA, and CTO of the cloud computing startup KeyByte. His research interests include reliable and secure software systems. Bagchi received a Ph.D. in computer science from the University of Illinois at Urbana-Champaign. Contact him at sbagchi@purdue.edu.

MAHSA GHASEMI is an assistant professor in the Elmore Family School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA. Her research interests include theoretical and foundational advancements for trustworthy sequential decision-making. Ghasemi received a Ph.D. in electrical and computer engineering from the University of Texas at Austin. She is a Member of IEEE. Contact her at mahsa@purdue.edu.

KANG G. SHIN is the Kevin and Nancy O'Connor Professor of Computer Science, University of Michigan, Ann Arbor, MI 48109 USA. His research interests include various QoS issues of cyber-physical systems including timeliness, fault-tolerance, security, privacy, and usability. Shin received a Ph.D. in electrical engineering from Cornell University. He is a Life Fellow of IEEE. Contact him at kgshin@umich.edu.

NALINI VENKATASUBRAMANIAN is a professor of computer science in the Donald Bren School of Information and Computer Sciences, University of California, Irvine, CA 92697 USA. Her research interests include distributed systems, adaptive middleware, pervasive and mobile computing, cyber-physical systems, resilient IoT, and formal methods. Venkatasubramanian received a Ph.D. in computer science from the University of Illinois at Urbana-Champaign. She is a Senior Member of IEEE. Contact her at nalini@ics.uci.edu.

DONGYAN XU is the Samuel Conte Professor in the Department of Computer Science, Purdue University, West Lafayette, IN 47907 USA. His research interests include computer systems security and malware defense and forensics. Xu received a Ph.D. in computer science from the University of Illinois at Urbana-Champaign. Contact him at dxu@purdue.edu.

SAMAN ZONOUZ is an associate professor in the School of Cybersecurity and Privacy and the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30308 USA. His research interests are in cyber-physical systems security. Zonouz received a Ph.D. in computer science from the University of Illinois at Urbana-Champaign. Contact him at saman.zonouz@gatech.edu.