Federal Register Notice: 89 FR 78915, Federal Register :: Networking and Information Technology Research and Development Request for Information on Cyber-Physical Systems Resilience Research, September 26, 2024.

Request for Information on the National Cyber-Physical Systems Resilience Plan

Meera Sridhar, Chenglong Fu, Miao Wang, Robert Cox, Rick Hudson, Ahmed Helmy, Bojan Cukic

University of North Carolina Charlotte's response to NSF RFI on Networking and Information Technology Research and Development Request for Information on a National Plan for Cyber-Physical Systems Resilience

We are writing in response to the National Science Foundation's Request for Information (RFI) on Networking and Information Technology Research and Development Request for Information on a National Plan for Cyber-Physical Systems Resilience. We appreciate the opportunity to contribute insights and recommendations to inform future research directions and priorities in this critical area. Below, we provide our recommendations for advancing research in this domain.

The University of North Carolina at Charlotte (UNC Charlotte) is a public research university located in Charlotte, North Carolina. Founded in 1946, it has grown into the second-largest institution in the University of North Carolina system, serving over 30,000 students across undergraduate, graduate, and doctoral programs. UNC Charlotte has established itself as a significant contributor to scientific advancement and a driver of regional innovation, preparing students to lead in a rapidly evolving research landscape. UNC Charlotte is known for its emphasis on applied research and innovation, particularly in fields such as cybersecurity, energy systems, advanced manufacturing, and data science. Leveraging its proximity to Charlotte's dynamic financial and technology sectors, the university fosters partnerships with industries to support innovative research and practical applications, The university's connection to Charlotte provides students with unique internship and job opportunities in a vibrant metropolitan area.

The Center for Energy Security And Reliability (CESAR), spearheaded by UNC Charlotte, is a collaborative center (with NC State University and NCA&T University) dedicated to research innovation for the creation of a robust, cybersecure, and carbon-neutral power grid infrastructure through interdisciplinary research, education, and collaborative partnerships. CESAR brings together a multidisciplinary team of researchers, educators, engineers, and industry partners to solve challenging problems and train the future workforce on critical infrastructure cybersecurity and resilience.

Written by: Meera Sridhar, Chenglong Fu, Miao Wang, Robert Cox, Rick Hudson, Ahmed Helmy, Bojan Cukic.

Please note any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the authors' affiliations or the University of North Carolina at Charlotte or the CESAR member institutions.

UNIVERSITY OF NORTH CAROLINA
CHARLOTTE

1. Challenges and Opportunities in Cyber Security and Resiliency for Critical Infrastructure Research

Cyber-Physical Systems (CPS) are naturally complex, with cross-platform, inter-communicating components. CPS such as power grids, water treatment facilities, and transportation control systems, face significant risks from cyber attacks that can lead to catastrophic real-world consequences. The intricate nature of CPS, characterized by numerous interdependent processes and devices, requires extensive manual intervention by highly skilled experts, limiting the feasibility of automated rapid-response mechanisms and potentially providing attackers with opportunities to escalate damage. Furthermore, many CPS devices manage essential physical functions, such as pressure monitoring in nuclear reactors, where improper reconfiguration or shutdown could result in disastrous outcomes. Additionally, while effective response strategies can be tailored to specific CPS configurations, the diversity of CPS applications, structures, and devices across various environments complicates the scalability of these solutions, making it difficult to generalize across deployments.

Some specific challenges are outlined below:

1A. Large, Powerful, High-fidelity Digital Twins: Designing and deploying advanced attacks and natural disasters that have not yet happened in the real-world requires a safe environment for testing. Digital twins can model advanced cyber attacks and other disruptions, and defenses, including AI-based, advanced social engineering etc.

Additionally, a significant challenge in advancing CPS security and resilience lies in the limited background knowledge and resources available for researchers to fully understand threats and conduct meaningful experimentation and evaluation. While cybersecurity researchers may be well-versed in identifying and categorizing various cyber threats, they often lack in-depth knowledge of the operational technology (OT) aspects specific to CPS, such as the control mechanisms, safety protocols, and physical processes that are critical to systems like power grids or industrial automation. This gap hinders their ability to accurately assess the real-world consequences of cyber threats on CPS and evaluate the practical effectiveness of defense mechanisms. Bridging this knowledge gap requires dedicated resources, including access to realistic CPS testbeds and cross-disciplinary collaboration, so that researchers can more effectively simulate, analyze, and respond to threats within these complex systems.

The power grid is an ideal example of a complex CPS that requires detailed testbeds for researchers to understand how to detect and prevent cyberattacks. Today's grid is becoming increasingly reliant on wide-scale, internet-connected distributed energy resources (DERs) such

as photovoltaic (PV) systems, battery energy storage systems, and smart loads such as electric vehicles and HVAC systems. An attacker could gain access to one individual DER, or possibly gain control of many such systems by accessing a back end or cloud-based control system from an entity known as an aggregator that controls many DERs spread out over a wide geographic area. Testbeds are needed that can allow cybersecurity researchers to examine how such attacks can impact the real-time delivery of electricity. For example, a large-scale attack conducted in an appropriate manner on a large number of DERs could cause major disturbances and instability on the bulk power grid. In many cases, even an attack on millions of DERs may cause instability that is immediately detected by protection equipment. Attackers could begin to develop smart approaches, however, that detect the operation of protection systems and subsequently change the nature of the attack so as to cause cascading failures. Such an attack could take down large parts of the bulk power grid for days.

Developing testbeds of the scale described here requires work to develop tightly coupled co-simulation frameworks that can simultaneously simulate multiple interconnected networks. An effective test bed for the power grid, for example, should co-simulate electrical networks, the internet, and protection equipment. Implementing such a testbed at scale is an enormous challenge because of the vastness of the system.

1B. AI-based Attacks and Defenses for Critical Infrastructure: Although advancements in anomaly and intrusion detection have been notable, effective response to these attacks remains challenging with the rise of generative AI. The proliferation of quickly evolving Large Language Models could assist attackers in identifying more covert attack vectors and devising stealthier, more destructive attack strategies, complicating efforts to develop software patches, enforce resource isolation policies, and eliminate attack surfaces effectively. Existing security protocols, intrusion detection systems, and response frameworks are largely unprepared for the sophisticated, adaptive attacking strategies enabled by generative and other advanced AI technologies. This limitation leaves CPS particularly vulnerable, as AI-assisted attacks can dynamically chain together vulnerabilities and execute complex, multi-stage strategies that bypass conventional defenses.

The integration of AI, particularly *deepfakes*, into CPS presents a significant security risk. By manipulating sensor data or control signals, attackers could compromise the integrity and safety of critical infrastructure. This emerging threat necessitates a multi-faceted approach to cybersecurity, including robust authentication mechanisms, advanced threat detection techniques, and resilient system design principles.

1C. Behavioral Science for Critical Infrastructure Resilience: Verizon's 2024 Data Breach Investigative Report indicated that 68% of breaches in the last year involved a non-malicious human element such as a person falling victim to a social engineering attack or making an

error[1]. There are numerous major examples in recent years, including the Colonial Pipeline ransomware hack that took down the largest fuel pipeline in the US, which was caused by a compromised password and password reuse.

Multiple research gaps exist in this area. First, work is needed to understand how to simulate the human element in large-scale digital twins such as those described previously. Power grids with numerous internet-connected distributed energy resources (DERs), for example, constitute large scale CPS. An effective digital twin of the power grid should include models of human users. These models need to be incorporated into time-domain simulations and reflect how humans react to various social engineering attacks at scale, so that system planners can understand the impact that such attacks can have on the power grid.

Work is also needed to mitigate impacts from the human element in cyber-physical systems. Many industrial control systems used today, for example, rely on simple security standards and are installed on old desktop computers found in back rooms with access to a wider network. Examples include Building Automation Systems that control HVAC equipment in hospitals, data centers, and other critical facilities, or industrial control systems in water-treatment plants. Such equipment is particularly vulnerable, and the operators often have minimal training on concepts in cybersecurity. It's not uncommon, for example, to walk into a control room with a password included on a post-it note. Work is needed to train such operators, and similar work is needed as it relates to internet-connected distributed energy resources now finding their way into many homes.

1D. A Science of Cybersecurity: While cybersecurity is being studied extensively, it still lacks a truly scientific foundation. Unlike established sciences like physics or chemistry, cybersecurity often relies on ad-hoc solutions and reactive measures, and often ends up being a "cat-and-mouse" game between attackers and defenders. The rapid evolution of technology, adaptive adversaries, the complexity of systems today, the lack of standardized metrics, and the constant emergence of new threats, makes it challenging to establish a systematic, evidence-based approach.

1E. Under-Resourced Areas: The United States has a vast number of small organizations responsible for maintaining critical infrastructure. For example, 56% of the American landscape is served by electric cooperatives, many of which have a very small staff[2]. There are also numerous small & medium sized manufacturing facilities found in rural areas in the United States that rely heavily on automation but have very limited staff capabilities to prepare for cyberattacks.

---

[1]Verizon, https://www.verizon.com/business/resources/reports/dbir/
[2] https://www.electric.coop/electric-cooperative-fact-sheet

1F. Closing the Loop and Aging Infrastructure: CPS rely on the seamless integration of digital and physical components, forming a *closed-loop* system. These systems continuously monitor physical components, send data to digital twins for analysis, and use the results to adjust physical actuators.

However, the increasing sophistication of cyberattacks, particularly those leveraging AI-powered deepfakes, poses a significant threat to the security and reliability of CPS. Real-time communication and control are essential for effective response, but network latency and reliability issues can hinder timely actions. The increased connectivity between physical and digital components introduces new attack vectors—-malicious actors can exploit vulnerabilities in communication protocols, network infrastructure, or software to disrupt operations or compromise sensitive data. Data manipulation or unauthorized access can lead to incorrect decisions and compromised security. Physical access to critical infrastructure components can be exploited by attackers or be affected by natural disasters. Additionally, the supply chain for critical infrastructure components can be vulnerable to cyberattacks. Malicious actors may introduce compromised hardware or software, leading to security breaches.

Additionally, a big challenge in achieving “closing the loop” and rapid response is the typically aging infrastructure in most CPS.

2. Proposed Research Directions

To strengthen the robustness of essential infrastructure systems, research initiatives should bridge the fields of the critical infrastructure domain science science with cybersecurity and resilience innovations, targeting automatic and scalable solutions on both the discovery and reduction of system vulnerabilities. Based on current trends and emerging technologies, we propose the following research directions that NSF should consider prioritizing:

2A. Large, Powerful, High-fidelity Digital Twins that can model all aspects (e.g., the IT and OT sides, the humans involved, and the diverse and evolving nature) of the complex critical infrastructure ecosystem. These types of digital twins will allow modeling and testing of advanced natural disasters and cyber attacks and defenses, including AI-powered scenarios, advanced social engineering attacks etc.

Prominent research avenues should include developing advanced tools for simulating and modeling real-time interactions within critical infrastructures that account for the increasingly decentralized nature of critical infrastructure systems. New research programs should foster the collaboration between electrical engineers, scientists that study natural disasters (such as geologists, hydrologists etc.), cybersecurity specialists, high-performance computing experts, and data analysts. Also, programs should encourage the development of digital twins that

enable researchers from diverse backgrounds to easily deploy and test their solutions, even if they have limited knowledge of other involved fields. Such a platform would lower the barriers to interdisciplinary collaboration, allowing researchers to validate and refine their methods in a broader context, thereby fostering knowledge sharing and technological integration across different domains.

Developing realistic representations of real CPS can be extremely complicated because of the vast nature of the interconnected networks. The power grid, for example, consists of many independently modeled electrical circuits running in conjunction with the internet. Constructing realistic models can be an extremely complex challenge requiring proprietary and sensitive data from utilities. In many cases, subject matter experts will need to get deeply involved to develop appropriate models. Advances in techniques such as federated learning could perhaps expedite the development of appropriate models.

2B. Advances in AI Techniques and Cyber Physical Resilience: Scalable and reliable vulnerability discovering and patching that aims to automate the existing security solutions with AI-empowered intelligent agents are an important area of research. These automation solutions will help overcome the scalability and cost issues of traditional security solutions that require specific expert knowledge or intervention. This helps the future autonomous systems such as vehicles and power grid to achieve self-sustainability against cyberattacks with acceptable costs.

Future research should also focus on integrating traditional security solutions and tools with the advanced reasoning capabilities of generative AI. Potential research tasks could range from developing security-focused AI agents to enhancing the reliability and robustness of these integrated systems. Achieving this goal requires approaches that address several key facets: (1) development of advanced AI models equipped with security-focused reasoning capabilities and up-to-date knowledge bases, enabling them to perform effectively in security tasks; (2) a comprehensive suite of security tools optimized for seamless integration with AI models, ensuring efficient information exchange and coordination; (3) improved organizational structures that incorporate AI-enabled agents into the existing workflows of critical infrastructures. New research programs should promote collaboration among experts in AI, cybersecurity, and operational sciences, encouraging applicants from each field to demonstrate how their proposed tasks will contribute to enhancing AI-enabled security solutions as a cohesive whole.

2C. Behavioral Science for Critical Infrastructure Resiliency: Developing effective digital twins such as the ones described above requires realistic models of human behaviors. Models that capture a range of possible responses to a social-engineering attack, for example, could be deployed as part of a multi-domain power-grid simulation. Work is needed to understand the statistical dependencies between certain attack responses and demographics, educational

background, etc. Stochastic models could be deployed as part of a larger power-grid testbed. This is an area requiring fundamental work in behavioral science and close collaboration between researchers in different fields, including behavioral science, data science, and computation.

Work is also needed to understand how to better train OT infrastructure operators how to mitigate the impact of cyber breaches. Such operators can be operating critical systems, but do not commonly use even the simplest of best practices, including the use of effective passwords. Work is needed to understand how to best train and prepare these operators. Similar work is needed to address the homeowners that now have many DERs.

2D. A Science of Cybersecurity: Formal methods, memory-safe programming languages and other mathematical approaches to security are more critical than ever for CPS, where violation of policies can be life-threatening. By providing a rigorous approach to model, analyze, and verify systems, formal methods help identify and analyze threats, design secure systems, detect and respond to attacks, and develop *cybersecurity and resilience* metrics for a CPS[3].

While formal methods offer significant advantages, in order to effectively apply them, future research should tackle challenges such as complexity and scalability to handle the complexity of CPSs. Additionally, history has shown that formal methods can be heavy-weight, and difficult to adapt for non-experts. Future research should focus on lightweight formal methods, and ease of policy specification and enforcement for non-experts.

An important research area is developing effective techniques for combining formal methods with AI and behavior science of cybersecurity for critical infrastructure.

2E. Closing the Loop: Addressing these cybersecurity and resiliency challenges to “close the loop” in CPS requires a multifaceted approach, including robust cybersecurity measures, resilient network infrastructure, advanced threat detection and response capabilities, and a strong focus on human factors and supply chain security. Robust physical security measures, including access controls, surveillance, and intrusion detection systems, are essential to prevent unauthorized access and sabotage. Operators and engineers must be trained to respond effectively to cyber threats and avoid mistakes that could exacerbate security incidents.

Future CPS must incorporate resilient physical components that can actively respond to attacks in real-time, based on feedback from both the physical and digital realms. By intertwining these components with large, powerful, high-fidelity digital twins, we can create a more robust and secure CPS ecosystem.

---

[3] https://www.whitehouse.gov/wp-content/uploads/2024/02/Final-ONCD-Technical-Report.pdf

3. Education and Workforce Development

To ensure the sustainability and future growth of research in this field, NSF should also consider enhancing workforce development. We recommend:

- Supporting educational initiatives that train the next generation of scientists and engineers in the *intersection of AI and cybersecurity.*
- Support for long-term training of a cadre of personnel to maintain all elements of a CPS. This extends to those who will maintain elements on a daily basis (at the technician level and above), under normal circumstances and who must respond to unforeseen situations involving physical and cyber disruptions—-either through maintenance issues, natural events or as a result of malicious activities.
- Testbeds can serve a role in training and developing operational scenarios for normal operation as well as response to out of the ordinary conditions.
- Red Team exercises for CPS.
- Expanding funding for programs that promote diversity and inclusion in STEM, and specifically in cyber-physical systems resilience.
- Providing funding to develop unique cross-disciplinary educational programs designed to provide appropriate training to students in multiple fields such as power systems, behavioral science and computer science.
- Provide resources for academic institutions and industry partnerships for training talent to serve under-resourced smaller entities (such as rural electric cooperatives) who are critical to America's national security.

4. Collaboration and Infrastructure

- DOD and civil systems can benefit from the cross-pollination of approaches. For example, as AI becomes an element of the Tactical Microgrid concept and that construct infuses the domain of fixed facilities (as US Army DEVCOM has identified as a desired objective) there will be a growing need for capabilities to assess, validate, and train.
- Creating mechanisms to facilitate multi-institutional collaborations, including shared research infrastructure.
- Collaboration between researchers across institutions and industries is crucial for advancing this field.
- Increasing funding for the construction (hardware, software, expertise) of large-scale experimental testbeds and computational infrastructure to support advanced research in CPS.