ACCEPTED FROM OPEN CALL

# Cyber Meets Control: A Novel Federated Approach for Resilient CPS Leveraging Real Cyber Threat Intelligence

Elias Bou-Harb, Walter Lucia, Nicola Forti, Sean Weerakkody, Nasir Ghani, and Bruno Sinopoli

While almost all works in the literature exclusively tackled the security of one independent aspect of CPS (i.e., cyber or physical), the authors argue that these systems cannot be decoupled. In this context, they present what they believe is a first attempt ever to tackle the problem of CPS security in a coupled and a systematic manner. To this end, this article proposes a novel approach that federates the cyber and physical environments to infer and attribute tangible CPS attacks.

## ABSTRACT

Cyber-physical systems (CPS) embody a tight integration between network-based communications, software, sensors, and physical processes. While the integration of cyber technologies within legacy systems will most certainly introduce opportunities and advancements not yet envisioned, it will undoubtedly also pave the way to misdemeanors that will exploit systems' resources, causing drastic and severe nationwide impacts. While almost all works in the literature exclusively tackled the security of one independent aspect of CPS (i.e., cyber or physical), we argue that these systems cannot be decoupled. In this context, we present what we believe is a first attempt ever to tackle the problem of CPS security in a coupled and a systematic manner. To this end, this article proposes a novel approach that federates the cyber and physical environments to infer and attribute tangible CPS attacks. This is achieved by

- Leveraging real cyber threat intelligence derived from empirical measurements.
- Capturing and investigating CP data flows by devising an innovative CPS threat detector.

An added value of the proposed approach is rendered by physical remediation strategies, which are envisioned to automatically be invoked as a reaction to the inferred attacks to provide CPS resiliency. We conclude this article by discussing a few design considerations and presenting three case studies that demonstrate the feasibility of the proposed approach.

## INTRODUCTION

Critical infrastructure systems are indispensable to the broader health, safety, security, and economic well-being of modern society and governments. In recent years, many of these systems, such as smart grids, nuclear plants, and automated highway systems, have been undergoing large-scale transformations with the infusion of new "smart" cyber-based technologies to improve their efficiency and reliability. These transitions are being driven by continual advances and cost-efficiencies in areas such as integrated networking, information processing, sensing, and actuation. Hence, increasingly, physical infrastructure devices and systems are being tasked to co-exist and seamlessly operate in cyber-based environments. Indeed, tightly coupled systems that exhibit this level of integrated intelligence are often referred to as Cyber-physical systems (CPS) [1].

Nowadays, CPS can be found in significantly diverse industries, including, but not limited to, aerospace, automotive, energy, healthcare, and manufacturing. Undeniably, the development and adoption of such CPS will generate unique opportunities for economic growth and improvement of quality of life. While CPS presents great opportunities, their complexity, which arises from the fusion of computational systems with physical processes, indeed poses substantial security challenges. To this end, novel vulnerabilities will manifest themselves, leading to attack models that are fundamentally new and hard to infer, attribute, and analyze.

Indeed, historical events confirm that industrial control systems have long been the target of disruptive cyber attacks. For instance, in 2010, the prominent Stuxnet malware was employed to target the SCADA control system of a critical uranium enriching facility, which triggered immense plant damage and even endangered human life [2]. Most recently, in March 2016, another malware was inferred to be responsible for the massive power outage that struck Ukraine in December 2015. Given the rapid transformation of industrial control systems toward CPS-based setups, attacks are anticipated to increase in frequency, sophistication, and target diversity. In fact, the latter trend was recently confirmed by the U.S. Department of Homeland Security (DHS), when they published the statistics in Fig. 1, revealing thousands of CPS attacks targeting diverse sectors [3].

Motivated by the imminent threats targeting CPS in addition to the lack of security approaches that tackle both aspects of such systems in a coupled and a coherent manner [4, 5], we frame the contributions of this article as follows:

•Proposing a new multidisciplinary approach that strives to diminish the gap between cyber security and control systems' science for securing CPS. Contrary to theoretical approaches that only consider the physical aspects of CPS in which some assumption is made regarding an attack

*Elias Bou-Harb is with Florida Atlantic University and the National Cyber Forensic and Training Alliance (NCFTA) of Canada. Walter Lucia is with Concordia University. Nicola Forti is with the University of Florence. Sean Weerakkody and Bruno Sinopoli are with Carnegie Mellon University. Nasir Ghani is with the University of South Florida, and the Florida Center for Cybersecurity.*

Digital Object Identifier: 10.1109/MCOM.2017.1600292CM

and its corresponding countermeasure, the proposed approach uniquely exploits tangible cyber threat intelligence (CTI) to infer real attack scenarios that could realistically affect CPS. Further, the proposed approach also considers the dynamics of CPS by triggering prompt remediation strategies in the physical realm as a reaction to the inferred attacks. To the best of our knowledge, the devised capability is of high impact in the CPS literature and has never been attempted before.

•Generating insightful CTI related to CPS by employing approximately 40 million real malware samples and undergoing dynamic binary analysis to capture CPS insider threats. The latter CTI is further expanded by characterizing and attributing real CPS external attacks by means of designing and deploying a high-interactive CPS honeypot.

•Evaluating and validating the feasibility and effectiveness of the integrated proposed approach by presenting three case studies that depict three different CPS attack scenarios. To this end, we also discuss some insightful design considerations that we expect to be guiding and helpful in the realization of future CPS security approaches.

The organization of this article is as follows. In the next section, we review some CPS security related works. We present the proposed approach by thoroughly discussing each of its components and pinpointing several design considerations. We demonstrate the effectiveness and practicality of the proposed approach by means of three different proof of concept case studies. Finally, closing remarks and future research directions conclude this article.

## Related Work

In this section, we review a number of related works in the context of security challenges in CPS, with special emphasis on control-theoretic, cyber, and hybrid approaches for securing CPS.

Research challenges related to CPS security have been addressed in prior works [4], where unprecedented CPS vulnerabilities and threats were investigated, and new directions for securing control systems were presented. Moreover, the authors of [6] highlighted the need for collaborative approaches that better integrate security into the core design of CPS.

From a control-theoretic perspective, Teixeira *et al.* [7] have introduced and modeled different attack scenarios such as false data injections, replay, and zero-dynamics attacks, where adversarial activities attempt to cause damage to the controlled system while remaining stealthy. Furthermore, Mo *et al.* [5] proposed an active detection method, known as physical watermarking, to authenticate the nominal behavior of a control system. To this end, a known noisy control input is injected to detect replay attacks by analyzing the output of the system. In another closely related work, Weerakkody *et al.* [8] introduced time-varying dynamics, acting as a moving target, to detect integrity attacks. Additionally, Fawzi *et al.* [9] focused on the design, implementation, analysis, and characterization of robust estimation and control in CPS when they are affected by corrupted sensors and actuators.

From a cyber perspective, Caselli *et al.* [10] presented a sequence-aware intrusion detection system that aims at detecting CPS semantic attacks using real empirical measurements from a water treatment facility. In an alternative work, Zonouz *et al.* [11] investigated malware that specifically target programmable logical controllers (PLCs) in CPS environments. The authors proposed a set of big data analytics rooted in symbolic execution that aim at capturing the behavior of malicious code.

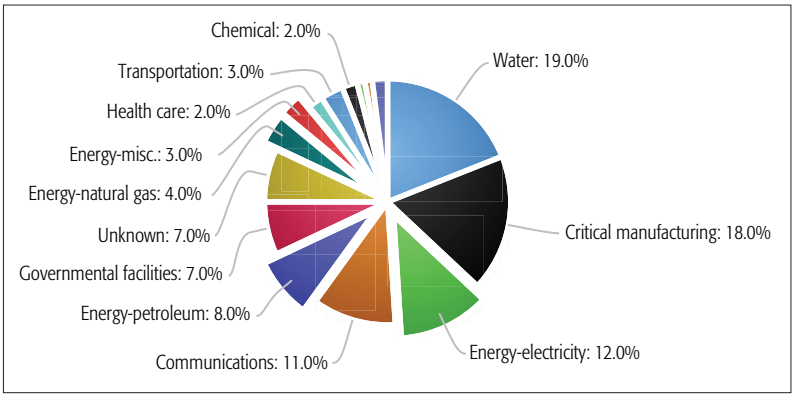


**Figure 1.** Recent attacks on cyber-physical systems as reported by DHS [3].

Complementary hybrid security approaches that attempt to systematically combine cyber and control capabilities are particularly rare. One interesting research specimen was proposed by Zonouz *et al.* [12]. In this work, the authors present an approach that aims at performing detection of corrupted measurements in power grids. Specifically, the proposed approach exploits alert notifications from intrusion detection and firewall systems to generate attack graphs providing an estimate of the compromised set of power grid hosts. Although such an approach leverages information from both the cyber and physical realms, it is neither capable of inferring specific types of attacks nor attributing the attack. Moreover, the proposed method is primarily focused on detection, and thus does not provide any tangible approach on how to provide CPS resiliency during or immediately after an attack.

## Proposed Architecture

In this section, we present and elaborate on the components of our proposed approach as depicted in Fig. 2. The core intuition behind the devised architecture is the unique introduction of the notion of *true maliciousness* to CPS security research. This notion embodies tangible malicious CPS attacks that could realistically affect the stability and security of CPS. The rationale behind this concept is threefold. First, the majority of literature approaches that solely focus on the physical aspects of CPS tend to characterize anomalies by a *deviation* of observed data in comparison to its expected value, generating a significant amount of false positives. Second, it is known that malicious empirical CPS security data from within operational CPS settings is extremely rare [10]. Third, it would be desirable to have an architecture that also provides insights and inferences that would aid in attribution. Undeniably, such attribution evidence will be leveraged to build highly-effective countermeasures to provide CPS resiliency. In the following section, we present and discuss the components of such an architecture that strives to exploit tangible CTI to infer and mitigate real CPS attack scenarios.

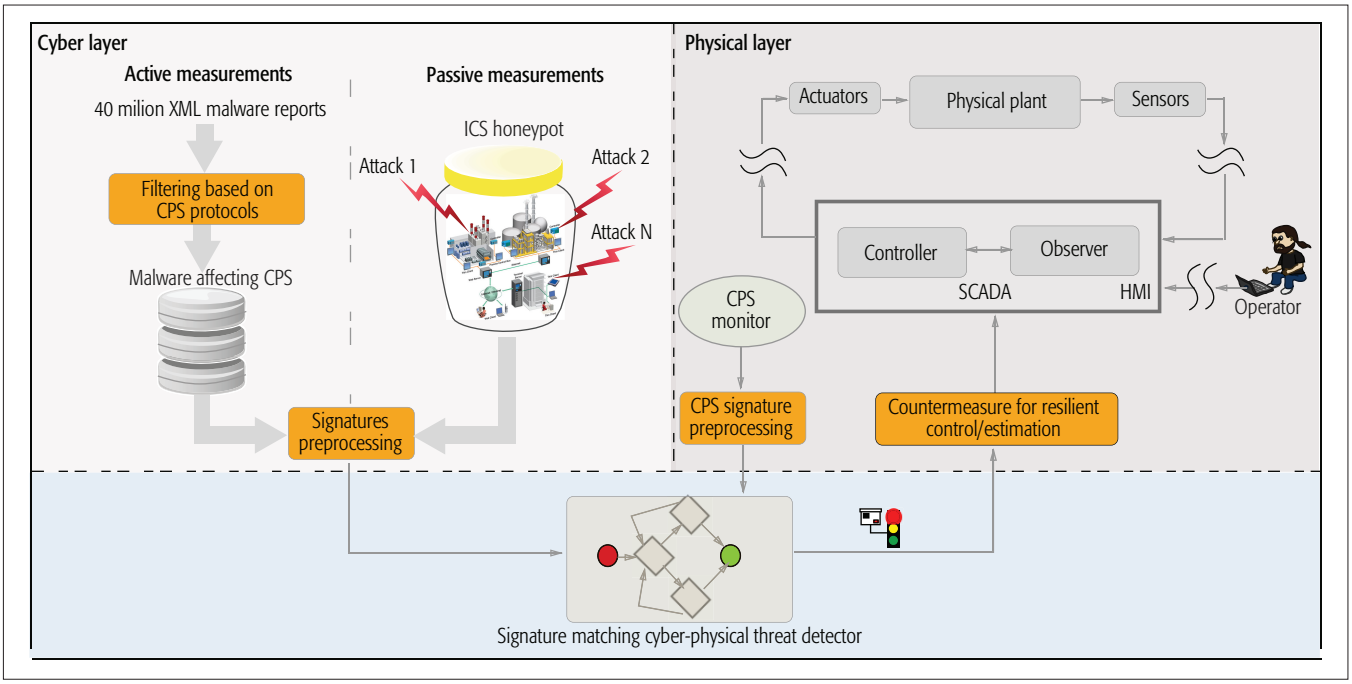


**Figure 2.** A holistic perspective of the proposed architecture.

## Cyber Layer

Threats toward CPS could arise from external as well as internal entities. On one hand, an example of an external threat could be rendered by scanning activities [13], originating from the Internet, in an attempt to probe and quantify CPS vulnerabilities, to be exploited in a subsequent directed attack. On the other hand, the Stuxnet malware would be an accurate example of an insider CPS threat, which leveraged system specific knowledge to execute a stealthy attack from within the boundaries of the CPS. Motivated by such diverse threats, this component aims at capturing real malicious CPS attack signatures for both internal and external threats. In this context, we define attack signatures by a series of consecutive attackers' steps that constitute a well-defined attack scenario. In the following, we describe how the latter CTI is generated.

**Active Measurements:** It is known that malware attacks similar to Stuxnet pose one of the most debilitating internal threats to CPS. To this end, we resort to dynamic malware binary analysis as depicted in Fig. 3 to generate real attack signatures that aim at capturing malware attack scenarios targeting CPS. We are fortunate to have access to malware data provided by Team Cymru Research. Consistent with Fig. 3, each malware binary sample is executed in a controlled client environment. During execution, the client would monitor and record the executed activities by the malware sample at the network and system levels. Consequently, the server processes such received information by producing an XML report, which summarizes the activities of the executed malware. Moreover, to select only those malware that specifically target CPS, we further execute a filtering mechanism on the indexed XML reports. Such filtering mechanisms can be easily modified to match CPS protocols and systems used in particular sectors.

**Passive Measurements:** It is also desirable to possess attack signatures related to external CPS attacks. Undeniably, the optimal approach to achieve this is to employ traffic capturing, measurement and analysis from the external boundaries of an operational CPS. However, due to legal, logistic, and privacy constraints, the latter is not always feasible. For such reasons, we resort to the concept of a honeypot: a set of software modules that can realistically imitate the components and the operations of any CPS. Further, a honeypot could be easily modified and tuned to generically exploit any CPS role within a specialized sector, allowing the capturing of a wide range of tailored external attacks. To this end, we design and implement a honeypot based on the open source industrial control system project dubbed as Conpot. To provide more realism to the CPS honeypot, we have implemented a custom capability that emulates the plant dynamics and configured the honeypot to operate several generic CPS protocols, with a human machine interface (HMI) and Simple Network Management Protocol (SNMP) capabilities. We deployed the honeypot online for a specific duration and enabled high levels of logging. To infer external attacks, we exploited the log files to build the attack strategies by tracking one attacker at a time, throughout its interaction with the honeypot. Please note that there might exist zero-day or unknown attacks that possibly would not be captured by the active and passive measurement modules. Future work will address how to remedy this issue.

## Physical Layer

This module deals with the physical (i.e., control) aspects of CPS. These include hardware and software modules capable of:

- Characterizing the physical process
- State estimation and reconstruction
- Stabilizing the physical plant
- Monitoring and managing the system

It is worthy to note that this module is generic to any CPS environment, including those operating in diverse domains such as power, wireless sensor networks, or manufacturing. To support CPS

attack detection and mitigation, we extend this layer by introducing an additional component that we refer to as the CPS monitor, as depicted in Fig. 2. The monitor exploits CPS communication channels and protocols to tap, gather, and amalgamate cyber-physical (CP) data flows that are circulating through a CPS. Moreover, the monitor coordinates with the CP threat detector to generate attack-resilient estimation and control remediation strategies, which are invoked as a reaction to an inferred attack.

### Cyber-Physical Threat Detector

This component lies at the intersection between the cyber and physical layers. Indeed, the core aim of the CP threat detector module is to investigate whether the CP data flows extracted from the CPS communication channels are susceptible to any tangible external or internal threats. An imperative auxiliary aim of this module is to characterize the severity of any inferred attack.

To achieve the intended tasks, the CP threat detector initially models both the malicious attack signatures generated from the cyber layer and the CP data flows extracted from the physical layer, into a common structure that we refer to as semantic behavioral graphs. Such directed graphs capture the activity performed as well as the semantics of such actions in the context of the observed protocol. To clarify the notion of semantic behavioral graphs, consider Fig. 4, which depicts a miniature specimen of such a graph capturing the dynamics of a benign CP data flow of the Siemens proprietary communication s7comm.

The CP threat detector proceeds in an attempt to infer any similarities between semantic graphs, as an indicator of an ongoing malicious activity on the CPS. However, given the fact that such graphs could possibly be of large scale due to the excessive captured/modeled network activity, and of high dimensionality due to the appended feature vectors, computing graph similarities in practice would indeed be challenging. In an attempt to overcome these issues, we devise a twofold approach.

First, the CP threat detector applies the notion of graph kernels borrowed from [14] on the formed behavioral graphs (discarding any semantics at this stage). The rationale here is to transform the similarity computation procedure of complex graphs into a linear space. To achieve this, sub-graphs from the behavioral graphs are initially extracted based on a certain criterion. Subsequently, compact representations of the created sub-graphs are generated based on a specific fingerprinting approach. Lastly, a mapping technique is employed to transform the latter representations into a linear space defined by a kernel matrix [14, 15], where the similarity computation is executed in linear time.

Second, we employ a threshold mechanism to deem when there is a significant similarity between those two types of graphs in order to flag that an ongoing malicious activity might be occurring on the physical plant. In this work, we set 60 percent as a conservative similarity threshold, to indicate a possible attack. To further reduce any other false positive cases and to confirm the attack, the CP threat detector proceeds by investigating the semantics of those graphs that exceed the similarity threshold. To this end, the CP threat detector performs binary comparisons between each corresponding pair of features of the semantic vectors. Currently, the header layer is only used in the comparisons for efficiency purposes; however, this can be easily extended to include the protocol layer and data layer. Any similarity between the semantics of the graphs of the CP data flows and the semantics of the graphs of the malicious signatures will likely confirm the existence of an ongoing attack. In this context, it is important to note that:

- We employ the similarity measure as a severity score.
- We are capable of providing tangible evidence attributing the attack to a specific malware specimen, in case the matching semantic behavioral graph was captured from the active measurements of the cyber layer.

Please note that the components of the proposed architecture that were discussed are fully automated, from an implementation perspective.

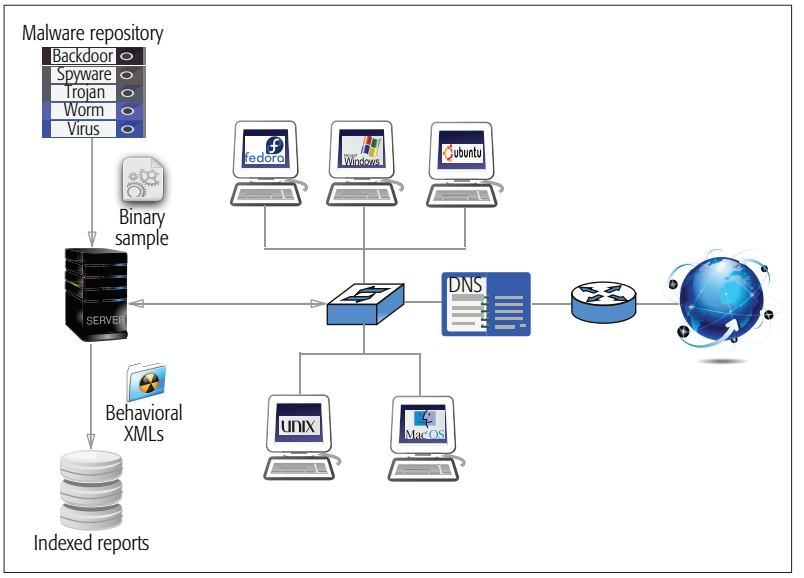


**Figure 3.** Dynamic malware binary analysis to capture real internal CPS threats.

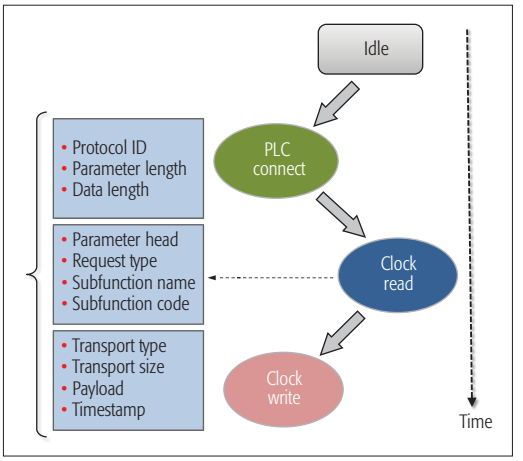


**Figure 4.** A high-level depiction of a semantic behavioral graph of a Siemens communication protocol.

While CP data flows in the physical layer are characteristically in the order of milliseconds or seconds, attacks in the cyber realm require time intervals of larger magnitude. Thus, when designing the CP threat detector, as detailed earlier, this issue ought to be taken into account.

### Design Considerations

The proposed architecture of Fig. 2 introduces a new paradigm that is rendered by the tight coupling and systematic fusion of the cyber and physical layers in the context of CPS security. Given that such an approach has never been attempted before, we thought that it would be of added value in this article to provide some design blueprints that elaborate on a few considerations related to the proposed scheme.

**On Time-Scale Discrepancy:** Architectures that aim at inferring, characterizing, and mitigating tangible attacks on CPS should pay special attention to time-scale discrepancies. While CP data flows in the physical layer are characteristically in the order of milliseconds or seconds, attacks in the cyber realm require time intervals of larger magnitude. Thus, when designing the CP threat detector, as detailed earlier, this issue ought to be taken into account. A feasible solution could be rendered by an approach that operates in a sliding time-window fashion; sampling periods from the physical realm could be employed in conjunction with attack signatures derived within specific time-window cycles. In this context, the time-window duration would be treated as a system parameter that would be tuned in order to achieve a balance between attack inference rates and computational load.

**On Detection Practicality:** Depending on the quantity and type of the derived CTI, a general design hurdle would be how to build effective and efficient models of such CTI that can be employed for detection. Commonly, attack detection and remediation in CPS realms require some near real-time requirements; any significant detection delay or poorly-timed reactions might result in cascading failures or damage of system components. Although in this work, we devised a CP threat detector based on efficient graph models, other approaches that have yet to be investigated could be more practical, easier to manage, and provide stronger analytical capabilities. Further, given that the extracted CTI could be extensive, as we expect and have observed throughout this work, one should initially model it offline and subsequently incrementally enrich it as new CTI becomes available.

**On Cyber-Physical Countermeasures:** Any approach that aims at achieving CPS resiliency should endeavor to provide cyberphysical remediation strategies to combat the effects of the inferred malicious activity. In this context, we advocate and stress the importance of designing countermeasures that cooperatively leverage information and capabilities from both the cyber and physical realms. Indeed, if only cyber mitigation strategies are executed, then the safety of the CPS under an attack cannot be guaranteed. Conversely, if only physical countermeasures are adopted, an attack-free CPS environment would be impossible to achieve. In this work, we attempted to capture the latter desired features by introducing the notion of semantic behavioral graphs as a modeling and a detection approach in the context of CP data flows and malicious attack signatures. Another desirable design goal to be considered would be the ability to characterize attacks through severity metrics, which would be highly beneficial from two perspectives, namely, situational awareness and prioritized mitigation. Additionally, one should postulate CTI approaches that not only can infer attacks, but more imperatively, can generate attack signatures that aim at disclosing attackers' strategies, aims, and intentions. Indeed, the active measurements approach described earlier, which exploits real malware strategies, attempts to achieve exactly the latter. In this case, the concrete knowledge of the disruptive resources available to the attacker allows the design of more effective cyber-physical countermeasures. Specifically, the observer and control modules could be re-designed by discarding the information originating from the corrupted CPS channels. Moreover, if an estimate of the attack vector is available, a resilient control solution can be achieved by means of an adaptive compensator, which simply adds a compensating term to the nominal control signal, avoiding redesign of the controller. Simultaneously, in the cyber realm, various information technology (IT) security countermeasures can be rapidly deployed to minimize attackers resources, and thereby limit their damage to targeted CPS assets. Such strategies could include dropping network traffic originating from attackers' IP addresses or dynamically adapting firewall rules, among numerous other available techniques.

## Proof of Concept: Case Studies

In this section, we verify the effectiveness of the proposed architecture by demonstrating its capability in generating tangible cyber threat intelligence related to CPS environments. Furthermore, we consider three case studies that capture two real external CPS attacks and an internal CPS attack launched by the eminent Stuxnet malware.

By operating the customized CPS honeypot as described previously for almost one month, we were able to infer around 500 unique attackers generating thousands of diverse malicious activities. By executing IP geolocation, Fig. 5a illustrates the countries where these attacks originated from, while Fig. 5b shows the Internet service providers (ISPs) responsible for some of those attacks in a one-day specimen. Note that we are aware that the University of Michigan performs regular benign scanning attempts toward various Internet services, including CPS services, and thus its appearance on this list verifies and validates the setup of the CPS honeypot and its capability in inferring malicious attempts. It is also worthy to mention that the statistics behind Fig. 5 could be relatively skewed by the abundant use of spoofing attacks. Indeed, given that IP spoofing is still present on almost 20 percent of the Internet autonomous systems, the reported statistics should not be taken as an absolute representation but rather as a figurative and a relative view of the status of CPS security in the context of real cyber threat intelligence.

We also had a brief chance to observe and investigate the network traffic packets arriving at the CPS honeypot. Our analysis revealed a staggering 10,000 generic TCP and UDP scanning attempts and close to 2,000 TCP flooding denial of service attacks on various CPS communication protocols, including those targeting the open source DNP3 and Modbus CPS protocols. We also generate supplementary material related to such misdemeanors including geo-location information per source, organization, city, and region. Although we refrain from publishing this informa-

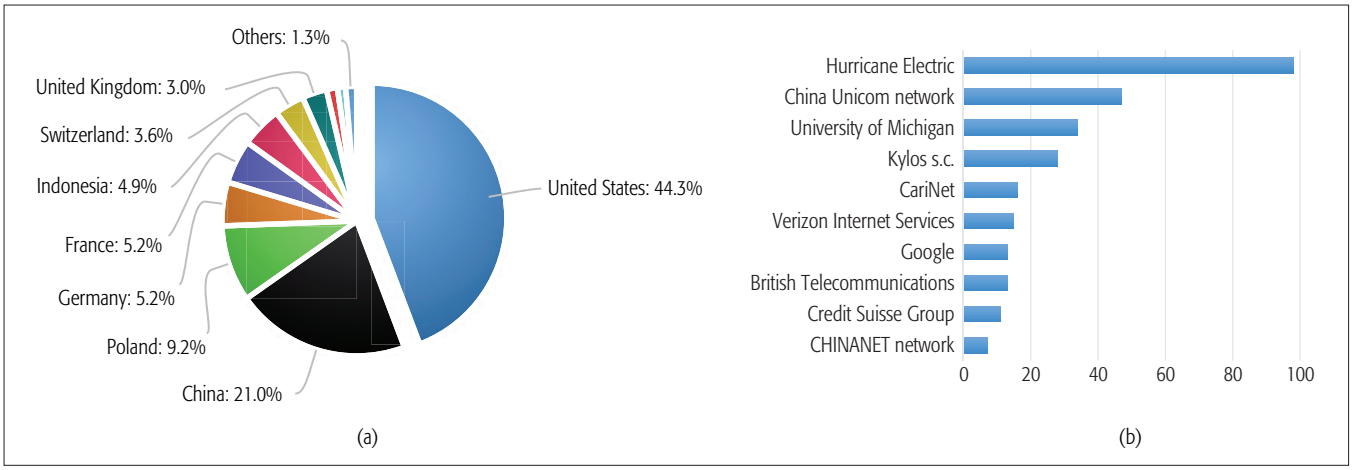


**Figure 5.** Real external CPS threats as inferred by the CPS Honeypot: a) distribution of all inferred threats per originating country; b) threats from top 10 Internet service providers during a 24-hour period.

tion due to sensitivity/legal issues, we can note that these attempts originated from 177 diverse operational providers, 124 distinct ISPs, and 71 different cities. Given that our analyzed period is relatively short, we concur that all the inferred attacks in terms of source diversity, target protocol diversity, frequency, and machinery are quite interesting and alarming.

We proceed by selecting and reporting on three real case studies that the proposed architecture was able to infer and remediate. The first case study captured an external CPS attacker that was inferred by employing the CPS honeypot. This scenario represented an attacker who attempted to gain elevation of privilege, and hence complete control over the plant, by striving to exploit the session manager of the CPS HMI. Part of this attack's inferred request is illustrated in Fig. 6a. Given that such an attack did not threaten the physical/control aspect of the CPS, a simplistic yet effective cyber strategy is applied to mitigate the attack by blocking any subsequent traffic from the attacker's IP address.

The second case study pinpointed another CPS external attack that was also perceived by leveraging the CPS honeypot. In this scenario, the attacker initially scanned the CPS plant using the SNMP to retrieve the map of all operational services. Once this has been achieved, the attacker read the Modbus variables and subsequently crafted some invalid input that is beyond a specified safe range. The latter attack appears to be an attempt to cause some sort of damage to the CPS equipment. Part of this attack's captured request is illustrated in Fig. 6b. In this attack scenario, a physical control remediation is obtained by simply discarding the injected malicious data, while a cyber countermeasure is enforced by filtering any further incoming packets from the attacker's IP address.

The third case study captured an instance of an internal CPS threat. In this attack scenario, detailed experimentation was conducted using a sample of the Stuxnet malware by closely following the proposed mechanism from earlier. Indeed, this malware employs replay attacks in an attempt to replicate previously-recorded sensor measurements to a system operator, while also injecting damaging inputs to the system. It is well known that this type of attack is stealthy to any passive

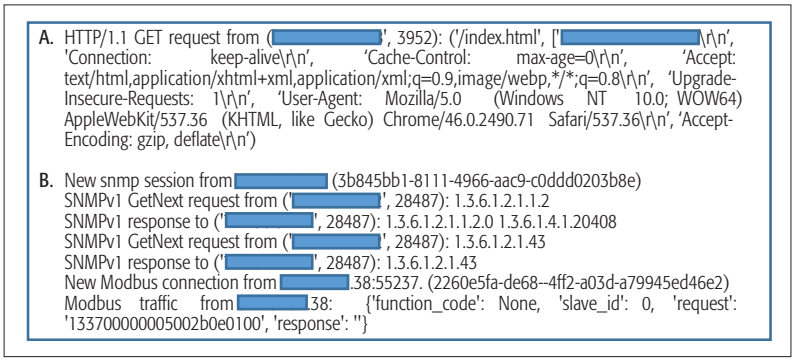


**Figure 6.** Snapshot of the requests generated by the external CPS attacks as inferred by the CPS honeypot.

physical detector because the resulting CP flows cannot be distinguished from benign CP flows characterizing an attack-free scenario. Nevertheless, in this case, the CP threat detector readily inferred the occurrence of an ongoing malicious activity in the physical realm since it captured the behavior and semantics of the malware as it interacted with the CPS. Furthermore, the proposed architecture was capable of attributing such activity to the exact sample of the Stuxnet malware, namely, Worm.Win32.Stuxnet.b . We also note the captured/modeled behavior of this malware, which included unauthorized reading and writing of the sensors' measurements as well as unauthorized writing toward CPS actuators' channels.

Indeed, the aforementioned case studies demonstrate the effectiveness of the proposed architecture in disclosing attackers' strategies and actions for both internal and external CPS attacks. More imperatively, the proposed approach also provides invaluable, tangible forensic evidence that can be employed for attribution in the cyber realm and resiliency in the physical realm.

## Concluding Remarks and Future Research Directions

Following our vision to diminish the gap between highly theoretical solutions and practical approaches for securing CPS, this article uniquely proposed a federated approach for resilient CPS by exploiting real CTI. The core rationale behind

The core rationale behind the proposed architecture is to derive tangible CPS attack models from empirical measurements, which can be employed to infer and attribute real CPS attacks. An innovative CP threat detector amalgamated CP data flows from the physical realm with attack signatures from the cyber realm.

the proposed architecture is to derive tangible CPS attack models from empirical measurements, which can be employed to infer and attribute real CPS attacks. An innovative CP threat detector amalgamated CP data flows from the physical realm with attack signatures from the cyber realm to infer and score real attack scenarios, as well as provide evidence of attribution and a means to generate CPS resiliency mechanisms.

This effort presents a solid basis from which to expand into other directions in CPS security. Foremost, one thrust will be dedicated to automating the tasks associated with creating resiliency countermeasures and algorithms from the captured CTI. Another aim is to investigate and develop additional CP threat detector designs and types and evaluate their detection tradeoff and efficiency under various realistic attack scenarios. Finally, an open source utility is also being designed to incorporate the overall notions and ideas behind the proposed architecture. This latter deliverable will help facilitate the incorporation, rapid prototyping, and much-needed evaluation of this solution in real-world operational CPS environments.

## References


[1] Xu Li *et al.*, "Smart Community: An Internet of Things Application," *IEEE Commun. Mag.*, vol. 49, no. 11, 2011, pp. 68–75.

[2] F. Kargl *et al.*, "Insights on the Security and Dependability of Industrial Control Systems," *IEEE Security & Privacy*, vol. 6, 2014, pp. 75–78.

[3] The Industrial Control Systems Cyber Emergency Response Team (ICS-CERT). https://ics-cert.us-cert.gov/.

[4] A.A. Cárdenas, S. Amin, and S. Sastry, "Research Challenges for the Security of Control Systems," *Proc. 3rd Conf. Hot Topics in Security*, Berkeley, CA, USA, 2008, pp. 61–66.

[5] Y. Mo, S. Weerakkody, and B. Sinopoli, "Physical Authentication of Control Systems: Designing Watermarked Control Inputs to Detect Counterfeit Sensor Outputs," *IEEE Control Systems*, vol. 35, no. 1, Feb. 2015, pp. 93–109.

[6] C. Neuman, "Challenges in Security for Cyber-Physical Systems," *DHS: S&T Wksp. Future Directions in Cyber-Physical Systems Security*, vol. 7, Citeseer, 2009.

[7] A. Teixeira *et al.*, "A Secure Control Framework for Resource-Limited Adversaries," *Automatica*, vol. 51, 2015, pp. 135–48.

[8] S. Weerakkody and B. Sinopoli, "Detecting Integrity Attacks on Control Systems Using a Moving Target Approach," *54th IEEE Conf. Decision and Control (CDC)*, 2015, pp. 5820–26.

[9] H. Fawzi, P. Tabuada, and S. Diggavi, "Secure Estimation and Control for Cyber-Physical Systems Under Adversarial Attacks," *IEEE Trans. Automatic Control*, vol. 59, no. 6, June 2014, pp. 1454–67.

[10] M. Caselli, E. Zambon, and F. Kargl, "Sequence-Aware Intrusion Detection in Industrial Control Systems," *Proc. 1st ACM Wksp. Cyber-Physical System Security*, 2015, pp. 13–24.

[11] S. Zonouz, J. Rrushi, and S. McLaughlin, "Detecting Industrial Control Malware Using Automated PLC Code Analytics," *IEEE Security Privacy*, vol. 12, no. 6, Nov. 2014, pp. 40–47.

[12] S. Zonouz *et al.*, "SCPSE: Security-Oriented Cyber-Physical State Estimation for Power Grid Critical Infrastructures," *IEEE Trans. Smart Grid*, vol. 3, no. 4, Dec. 2012, pp. 1790–99.

[13] E. Bou-Harb, M. Debbabi, and C. Assi, "Cyber Scanning: A Comprehensive Survey," *IEEE Commun. Surveys & Tutorials*, vol. 16, no. 3, 2014, pp. 1496–519.

[14] S. V. N. Vishwanathan *et al.*, "Graph Kernels," J. Mach, *Learn, Res.*, vol. 11, Aug. 2010, pp. 1201–42.

[15] A. Teixeira *et al.*, "Min-Hash Fingerprints for Graph Kernels: A Trade-Off Among Accuracy, Efficiency, and Compression," *J. Information and Data Management*, vol. 3, no. 3, 2012, pp. 227.


## Biographies


ELIAS BOU-HARB (ebouharb@fau.edu) is currently an assistant professor in the Computer Science Department at Florida Atlantic University, where he directs the Cyber Threat Intelligence Laboratory. Previously, he was a visiting research scientist at Carnegie Mellon University. He is also a research scientist at the National Cyber Forensic and Training Alliance (NCFTA) of Canada. He holds a Ph.D. degree in computer science from Concordia University, Montreal, Canada. His research and development activities and interests focus on the broad area of operational cyber security, including attack detection and characterization, Internet measurements, cyber security for critical infrastructure, and big data analytics.

WALTER LUCIA is an assistant professor at the Concordia Institute for Information Systems Engineering, Concordia University, Canada. He received an M.Sc. in automation engineering and a Ph.D. in systems and computer engineering from the University of Calabria, Italy. He was a visiting research scholar in the ECE Department at Northeastern University, USA, and a visiting postdoctoral researcher in the ECE Department at Carnegie Mellon University, USA. His current research interests include control of unmanned vehicles, switching systems, fault-tolerant control, model predictive control, and resilient control of cyber-physical systems.

NICOLA FORTI received his B.Sc. degree in mechanical engineering and a M.Sc. degree in automation engineering from the University of Florence, Florence, Italy, in 2009 and 2013, respectively. He is currently pursuing his Ph.D. degree in information engineering with the Automatic Control group at the University of Florence. In 2015 he visited the Department of Electrical and Computer Engineering at Carnegie Mellon University in Pittsburgh, PA, USA for one year. His main research interests include networked control systems, data fusion, secure estimation and control of cyber-physical systems, and multi-target tracking.

SEAN WEERAKKODY received the B.S degree in electrical engineering and mathematics from the University of Maryland, College Park, USA, in 2012. He was awarded the National Defense Science and Engineering Graduate fellowship in 2014. He is currently currently pursuing the Ph.D. degree in electrical and computer engineering at Carnegie Mellon University, Pittsburgh, PA. His research interests include secure design and active detection in cyber-physical systems and estimation in sensor networks.

NASIR GHANI [SM] is a professor in the Department of Electrical Engineering at the University of South Florida (USF), and also a research liaison for the state-wide Florida Center for Cybersecurity (FC2). He has over 20 years of research and development experience in the academic and industry sectors. Earlier he was associate chair of the Electrical and Computer Engineering Department at the University of New Mexico, and prior to that, a faculty member at Tennessee Tech University. He has also spent several years working in industry (Nokia, IBM, and Motorola) and several hi-tech startups. Dr. Ghani has co-authored over 200 publications, and also has three highly-cited U.S. patent inventions. His research interests include high-speed networks, cybersecurity, cyberphysical systems, cloud computing, and disaster recovery. Currently he is involved in various projects on cybersecurity for the Internet of Things (IoT) and software defined networks. His work has been funded by the National Science Foundation (NSF), Department of Energy, Department of Education, Qatar Foundation, and Sprint-Nextel Corporation. He also received the NSF CAREER award in 2005. Dr. Ghani has served as an associate editor for *IEEE Communications Letters*, *IEEE/OSA Journal of Optical and Communications and Networking*, and *IEEE Systems*. He has also guest-edited special issues of *IEEE Network* and *IEEE Communications Magazine* and has co-chaired numerous symposia for IEEE GLOBECOM, IEEE ICC, and IEEE ICCCN. He also served as chair of the IEEE Technical Committee on High Speed Networks (TCHSN) from 2007-2010. He received his bachelor's degree from the University of Waterloo, his master's degree from McMaster University, and his Ph.D. degree from the University of Waterloo.

BRUNO SINOPOLI received the Dr. Eng. degree from the University of Padova in 1998, and his M.S. and Ph.D. in electrical engineering from the University of California at Berkeley, in 2003 and 2005, respectively. After a postdoctoral position at Stanford University, Dr. Sinopoli joined the faculty at Carnegie Mellon University, where he is a professor in the Department of Electrical and Computer Engineering with courtesy appointments in mechanical engineering and in the Robotics Institute and co-director of the Smart Infrastructure Institute, a research center aimed at advancing innovation in the modeling analysis and design of smart infrastructure. Dr. Sinopoli was awarded the 2006 Eli Jury Award for outstanding research achievement in the areas of systems, communications, control, and signal processing at U.C. Berkeley, the 2010 George Tallman Ladd Research Award from Carnegie Mellon University, and the NSF Career award in 2010. His research interests include the modeling, analysis, and design of secure by design cyber-physical systems with application to interdependent infrastructures, Internet of Things, and data-driven networking.