2016 IEEE Conference on Communications and Network Security (CNS): International Workshop on Cyber-Physical Systems Security (CPS-Sec)

Attacking Supercomputers Through Targeted Alteration of Environmental Control: A Data Driven Case Study

Keywhan Chung, Valerio Formicola,
Zbigniew T. Kalbarczyk, and Ravishankar K. Iyer
Coordinated Science Laboratory,
University of Illinois at Urbana-Champaign,
Urbana, Illinois, USA
{kchung10,valeform,kalbarcz,rkiyer}@illinois.edu

Alexander Withers, Adam J. Slagell
National Center for Supercomputing Applications,
University of Illinois at Urbana-Champaign,
Urbana, Illinois, USA
{alexw1,slagell}@illinois.edu

***Abstract*—In this paper, we show that a malicious user can attack a large computing infrastructure by compromising the environmental control systems in the facilities that host the compute nodes. Such violations cannot be easily recognized by the administrators who manage the cluster, because of limited observation of the events in the cyber-physical systems. We describe real cases of failures due to problems in the cooling system of Blue Waters, the petascale supercomputer of the University of Illinois at Urbana-Champaign. Blue Waters has cooling cabinets that use chilled water provided by the National Petascale Computing Facility (NPCF). We demonstrate, using real data, that the control systems that provide chilled water can be used as entry points by an attacker to indirectly compromise the computing functionality through the orchestration of clever alterations of sensing and control devices. In this way, the attacker does not leave any trace of his or her malicious activity on the nodes of the cluster. Failures of the cooling systems can trigger unrecoverable failure modes that can be recovered only after service interruption and manual intervention.**

I. Introduction

In this paper, we present a study of failures of the Blue Waters supercomputer at the University of Illinois at Urbana-Champaign that were caused by problems in the environmental temperature of the cabinets. Real-data show that the incidents were caused by anomalies in the cyber-physical control systems, and we claim that a malicious user can induce similar failures by penetrating the facility building automation systems. This kind of attack can be very hard to detect since it corrupts the computational functionality of the supercomputer without leaving a visible trace for the system administrators to follow. The attack exploits the limitations of the system monitoring the physical changes and their consequence on the cyber-environment, e.g., in providing fast mitigation for the temperature increase in the cabinets hosting the compute nodes. The attack takes advantage of the inertia in the cooling system (connected to the compute nodes) to ensure correct operational parameters when sudden changes affect the cooling sources.

Following an increased interest in cyber-physical systems (CPSes) and industrial control systems (ICSes), a number of researchers have studied design challenges of such systems in the context of ensuring security. Especially, in terms of chilled water-based cooling, [1], [2] presented a practical overview of liquid-based cooling systems, its design and installation. Patterson et. al., in [3], discusses the challenges of liquid based cooling in terms of cost and power efficiency by comparing the cooling systems of the top 15 super computers. In understanding the problems in developing current generation of CPSes, [4] highlights the poor knowledge on coupling and interdependencies among different CPSes and explores different techniques for modeling critical infrastructures. Addressing the security problems in the CPSes, [5] presents an overview of smart grid and CPS security, especially architectural vulnerabilities, and the root causes of security problems. In addition, [6] studies cascading failures as a vulnerability in power grid security.

In this paper, we demonstrate the possibility of attacking supercomputers through targeted alternation of environmental control. The attack is subtle as it exploits the coupling between two CPSes to alter the environmental parameters to cause failure of the compute nodes. In this context, this study presents a new attack vector which utilizes a vulnerability in a CPS to intrude the operation of a well-hardened computing infrastructure.

The contributions of this paper are: (1) Examples of failures of Blue Waters supercomputer that were caused by the deteriorated environmental conditions of the facility hosting the compute nodes; (2) An attack model that exploits knowledge on vulnerabilities of the environmental controls to cause failures (including system-wide outages) of a well-hardened supercomputing facility.

II. Chilled Water System from Campus to Blue Waters

Figure 1 shows the Blue Waters supercomputer with the cyber-physical control systems that maintain the operational environment. Blue Waters is a petascale supercomputer managed by the National Center for Supercomputing Applications

978-1-5090-3065-1/16/$31.00 ©2016 IEEE

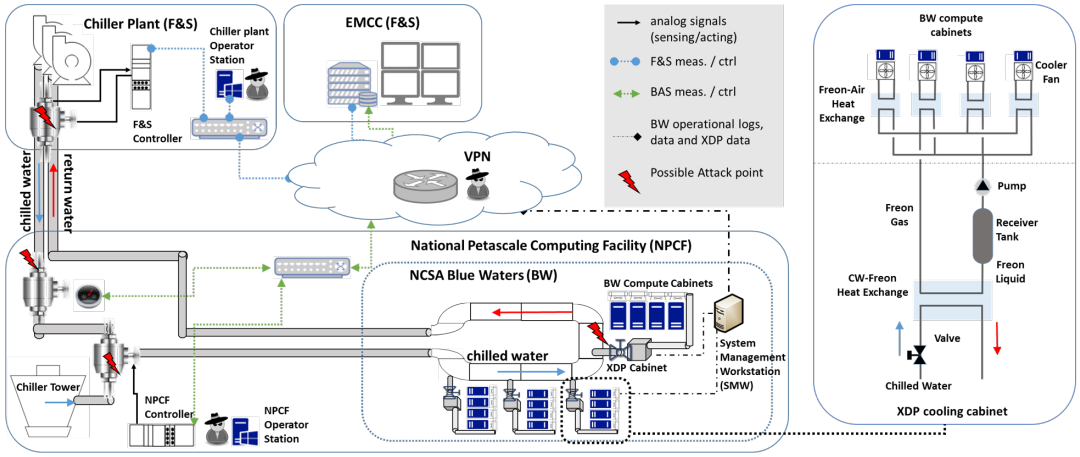

Fig. 1: Overview of the CPS architecture from the campus to the computing system, and details of the cooling cabinet.

(NCSA) at the University of Illinois. Blue Waters is located at the National Petascale Computing Facility (NPCF), an 88,000-square-foot building that has adopted a state-of-the-art building automation system (BAS). The BAS regulates the environmental parameters of the server room, as well as the offices hosted in the building.

Cooling the hardware to a predefined temperature is required for optimal operation and to guarantee the availability and reliability of computational resources. The computing cabinets of Blue Waters are cooled through dedicated Freon-based cooling cabinets (Liebert Cray XDP) provided by the vendor, Cray. There are 72 XDP cabinets that control 288 computing cabinets. NPCF provides the chilled water to condense the Freon through the main water loop installed in the building. The loop is designed to target energy efficiency and blending two water sources: the campus chilled water provided by the chiller plant of the campus Facilities and Services (F&S) unit, and a set of water towers chilled by nature. The mix of waters and the flow pressure are regulated by the NPCF BAS.

A. *Cooling Systems in the Campus and in Blue Waters*

The chilled water arriving at the Cray XDP system is initially regulated through two different control systems: the campus ICS (F&S) and the NPCF control system. The campus ICS prepares and delivers the chilled water to campus buildings, including the NPCF. F&S manages the campus chilled water plant (flow regulators, pumps for pressures, and tower cooling fans), and the distribution network (mostly distribution valves), over a virtually isolated network for remote control and measurements. Actuators adjust the control values based on measurements and commands from the remote engineering workstations at F&S.

NPCF's second source of chilled water is its on-site cooling towers. To efficiently control the cooling system (and other environmental parameters), NPCF has deployed a BAS that operates on a dedicated VPN. Measurement points (chilled water temperature, pressure, and flow) and actuators (valves to regulate the input source, pumps to keep the pressure consistent, and fan for cooling at the cooling towers) are used to provide the chilled water to the XDP cabinets within a predefined range of pressure and temperature. As we show, violation of the predefined conditions can eventually lead to the failure of the computing nodes.

The right side of Figure 1 illustrates the cooling system within the XDP cabinets. The system includes two heat exchange loops. In the bottom loop, the chilled water absorbs the heat from the Freon gas and converts the gas to liquid (chilled water). The liquid-form Freon is then distributed to the four computing cabinets on the top, where the Freon in each cabinet chills the air that flows in the computing cabinet. When Freon absorbs the heat from the air, it is transformed into a gas, which is then returned to the heat exchanger in the XDP cabinet. Each XDP cabinet includes an independent iCOM controller, which automatically controls the cabinet parameters. The controller manages the valve actuator to adjust the amount of water required to condense the Freon. Also, the controller regulates the Freon pressure on the top loop through a pump. Measurements of the refrigerant temperature and pressure are used as a reference. Thus, the heat-removal capacity of Freon depends on its pressure and temperature.

A remote system management workstation (SMW) queries every cabinet to collect the measurements and transfer them to dedicated storage. A dedicated board in each XDP pulls the iCOM controller at any SNMP request from the SMW.

B. *CPS Monitoring and Data Sources*

The measurements and control packets at F&S are mirrored for monitoring and archived in a database at the Energy Management Control Center (EMCC), which manages the generation and distribution of energy over the campus. NPCF is in charge of the ICS and BAS in the hosting facility. XDP data is monitored by the vendor, Cray, and collected on the SMW node of Blue Waters, managed by NCSA. Monitoring of the BAS and the XDP cooling cabinets relies on periodic manual inspection and on-site check upon alarms delivered by email to the operators in the different organizations.

For analytics, in collaboration with NCSA (which manages the Blue Waters computing services), NPCF, and F&S, we

have deployed a SPAN port on each control network switch to collect packet data in addition to decoded parameters. The environmental data sources can be categorized into three subsets: (1) XDP cooling cabinet (1 per min), providing chilled water temperature, room temperature/humidity, dew point, valve percentage; (2) NPCF BAS (1 per min), providing chilled water flow/temperature/pressure, motor control, valve percentage, pump control, fan speed; (3) campus F&S ICS (1 per 5 sec), providing heating capacity (MTBU), chilled water speed/temperature/pressure, fan speed, vibration.

III. Analysis of Blue Waters Supercomputer Failures due to Problems in the Cooling System

In this section, we discuss incidents that occurred in the cooling system of Blue Waters drawn jointly from the failure reports and ICS monitoring data for a period spanning from June 2013 to June 2016. Failures in the cooling system of Blue Waters are less frequent than logical failures discussed in [7], accounting for 2.73% of total incidents (i.e., 148 out of 5419). The failures are based on the logs maintained by system engineers who maintain the compute nodes. The failing components can be categorized as being within the XDP cooling cabinets (valve, pump, gasket, temp. sensor), the XE[1] computing cabinets (fan shutoff), or the BAS. Failures that affect Blue Waters, and are visible to the administrators (e.g., that propagate to system-wide alarms) occur mostly for unexpected behavior of the XDP cabinet. Less frequent are the BAS failures that propagate to the computing cabinets, as they have more chance to be mitigated by the CPS controls in the XDP cooling cabinets.

A. Scenario 1: Loss of Freon Pressure

Pump gasket failures and Freon leakage accounted for 39.73% of the failures in the cooling system. Unlike problems in the chilled water loop, problems in the Freon loop have direct effect on the computing cabinets. If there are problems in the water loop provided by the campus or the NPCF building, each XDP cooling cabinet can adjust its actuators to mitigate the changes. Specifically, when the heat-removal capacity of the water decreases (e.g., because of a problem in the building's water-flow pressure or the water temperature), the XDP valve in the cooling cabinet is automatically opened to increase the amount of water that exchanges heat with the Freon. Despite the possibility of compensating for such anomalies, two factors limit the heat-removal capacity of the XDP cabinet: the maximum amount of water that can enter the XDP when the valve section is 100% open, and the level and flow pressure of Freon liquid present in the XDP when there is low pressure of Freon. Thus, the chances of mitigating BAS failures upon a gasket failure and Freon leakage are limited.

Figure 2 presents a temporal sequence of chilled water temperatures, refrigerant temperatures, and valve opening percentages in an XDP cabinet with a gasket failure. A major failure happened around 09:00. A disconnection of the measurements

[1]XE is the name of the Cray architecture for supercomputers used in Blue Waters

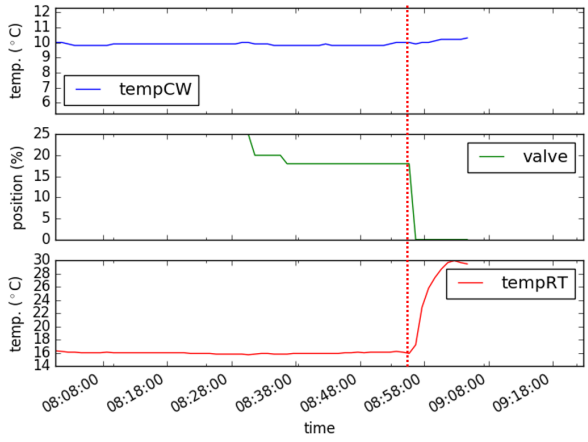

Fig. 2: XDP measurements at the time of gasket failure and its repair: chilled water temperature (tempCW), valve open % (valve), refrigerant temperature (tempRT)

indicates that the cabinet was brought down for maintenance and then was returned to service after 30 hours. While the pump compressed the Freon to compensate for the loss of pressure, the decreasing density eventually led to the loss of cooling capacity. Even if the leakage had been detected, repair of failing parts and recharge of the gas would still have required a shutdown of the cooling cabinet. As an XDP cabinet is the only cooling source for a set of four computing cabinets, a failure can cascade to four or more computing cabinets. This scenario shows that (1) the loss of cooling capacity in the XDP directly affects the compute nodes, despite the heat removal provided by the water in the building; and (2) the dynamics of some events can be so fast that there is not enough time to recover and compensate for the variation locally, thus causing the compute nodes to experience an emergency power off (EPO) within a short amount of time.

B. Scenario 2: Loss of Control on Water Valve Actuator

The most common failures in the cooling system affect the water valve actuator. Depending on the measurements of the refrigerant temperature, XDP determines the valve opening level to regulate the cooling capacity required to keep the set point. This operation is frequent during the intermediate seasons (spring and fall), as the BAS performs more transitions because of the blending operations between the tower water and the F&S chilled water. These transitions trigger more XDP reactions that can lead to anomalous behaviors of the valves. Figure 3 shows the loss of control in the valve actuator in one of the XDPs. Because of a variation in the cooling capacity (chilled water temperature is changing), the valve opens and then closes around 15:40. A sudden decrease in the CW temperature initiates a closure of the valve to ensure that the refrigerant temperature remains at the set point. However, with the valve stuck at "full closure,"the cooling cabinet loses control over the cooling capacity, which results in an increase in the refrigerant temperature. This scenario shows that (1) again, the time needed to react to a sudden change in the heat-

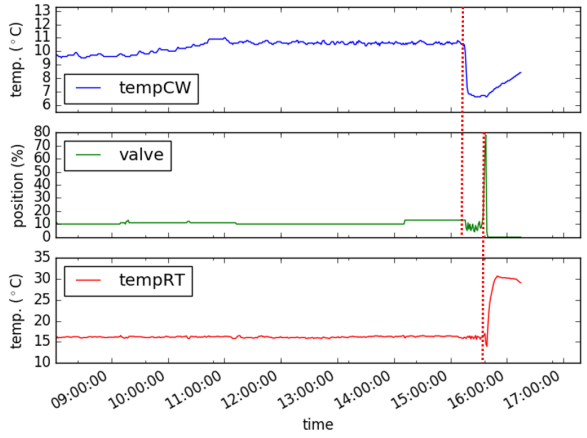

Fig. 3: XDP measurements at the time of valve actuator failure and its repair

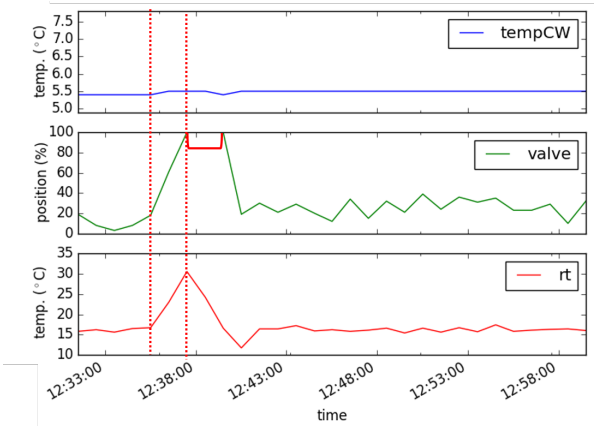

Fig. 4: XDP measurements at the time of NPCF Chilled Water pressure loss

removal capacity of liquids here refrigerant temperature rise induced by a fully closed valve can be critical to recover; and (2) complete XDP valve closure is a risky operation induced by the NPCF to operate the transitions for water blending.

C. Scenario 3: Change in Chilled Water

In June 2016, the XDP experienced a failure during a water-blending operation at NPCF and F&S. The situation began when the campus F&S performed an operation on the tank attached to the campus chilled water system. The operation led to an increase in the pressure of the NCPF chilled water input, which the BAS was able to mitigate. However, when the operation ended, the chilled water pressure returned to normal very rapidly, and the BAS experienced a latency in mitigating for the change. The result was a loss of pressure in the chilled water loop. Figure 6 illustrates the XDP monitoring results on the day of the failure in one of the cabinets affected. To adjust for the pressure drop, the XDP opened the valve to increase the flow of chilled water. Once the valve reached 100%, XDP reached the maximum cooling capacity, leading to an increase in refrigerant temperature because of temporary pressure decrease. An analysis of the XDP data shows that at least seven cooling cabinets (XDP) had reached 30.59 °C, which is 14.5 °C above the set point. Combined with the heavy workload on the computing cabinets (XE), five compute nodes had EPOs as a result of exceeding the temperature threshold.

This scenario provides another failure mode and shows that 1) operations in the F&S can have impact directly on the compute nodes; 2) a system failure can be induced also by the sudden change in the water-flow pressure in the NPCF loop, and 3) the workload on the compute nodes chilled by one of the XDPs accelerates the increase of refrigerant temperature and the probability of EPO.

D. Discussion

The three failure scenarios indicate that the fast change of the refrigerant pressure or temperature is the most critical events in the CPS that can induce failures of the supercomputer. These events can be directly triggered in the XDP cabinets, or induced by the NPCF or F&S operations. Failures occur when the drastic deviation of cyber-physical parameters is not effectively mitigated before impacting the cooling of the compute nodes. In particular, XDP cooling cabinets are not able to provide timely mitigation of problems in the chilled water because of a significant difference in the transition times of BAS and XDP needed to restore the set points. If the workload on the compute nodes is significant, the cabinets suffer in EPO. Note that the problems that caused the failures in Blue Waters have been addressed by the operation team.

IV. Exploiting the CPS Mechanisms to Attack the Supercomputer

In this section, we describe an example attack using Blue Waters data. The attack is highly unlikely to succeed on Blue Waters, but the knowledge from the data helps us imagine possible attacks on general systems with similar scale and setting.

A. CPS Security policies

Similar to the CPS setup described in section II, isolation of control network through virtualization has become a common practice. The VPN technology provides a certain level of isolation without the need for installing a dedicated physical network. In addition, without the airgap, VPN allows remote access to the control system from an external network. While the remote access enables immediate responses to problems despite physical distance, it represents a risk since attackers can gain an entry point to the ICS by compromising the operator machines. In Figure 1, possible entry points for a cyber-attack are indicated by the black hats. The most direct attack on environmental controls would be via the XDP cabinets, accessed for instance from the bastion hosts. In that case the attacker might observe, or even alter, the cooling system parameters transmitted to the SMW. Additionally, he/she might violate the bridging board on the XDP to update the firmware on the iCOM, or stress the controller resources to create an interruption of the controller. A less invasive attack is possible

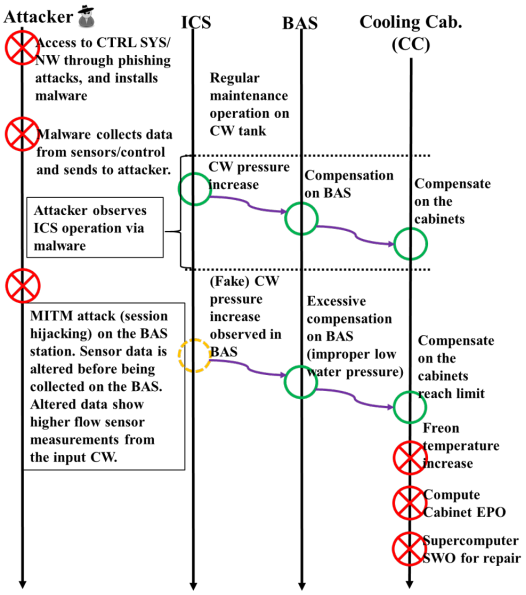

Fig. 5: Attack sequence based on failure scenario 3.

if the attacker is able to access the building automation system. In that case he/she has to compromise the chilling pumps in the tower and in the inlet water pipe. Finally, an attacker might be able to use the operator console of the chiller plant at F&S, where he/she can act on the chiller pumps.

B. Attack on a supercomputer via CPS

Taking advantage of our knowledge about CPS-related failures, we have built an attack scenario that (1) can be implemented as malware on CPS systems; (2) is hard to examine forensically to discover the root cause; and (3) can result in a system-wide outage (SWO), i.e., such that system services in the cluster must be interrupted to restore operation.

In this attack we assume the attacker obtains access to the ICS via phishing. The attacker first monitors the maintenance operations by installing a malware that exports sensor parameters from the BAS to a remote server. He/she uses this malware to study the target system. Then the attacker executes a Man in The Middle (MITM) to counterfeit the pressure values sent from the building input CW sensors. The values in the altered packets contain higher pressure than their actual value. The blender at the facility housing the supercomputer is induced to reduce the pressure to compensate the change. This recreates the conditions on the failure scenario 3, resulting in EPO and SWO. More details are available in Figure 5.

C. Possible Mitigation Methods

While we will leave the design and deployment of the attack for future work, here we list possible solutions. A common practice is rigorous control of policies on remote base workstations and laptops. In addition, the deployment of multi-factor authentication methods for access to the control systems and networks has effectively secured the Blue Waters. Also, system-level security monitoring can be deployed in the ICS, as in power systems [8], to validate physical aspects of measurements.

V. Conclusion

In this paper, we have demonstrated the feasibility of an attack that targets a well-hardened supercomputer infrastructure. The attack is based on a malware that can be installed by an intruder on the control network of one of the facility cooling systems. The intruder bypasses security barriers and exploits vulnerabilities and/or weak security policies in the CPS on which the computing infrastructure relies. The malware attacks the environmental control system, to indirectly target the supercomputer relying on those services; the attack is subtle because it exploits the characteristics of the dependence between the two systems during different operations to generate specific alteration of the cyber-physical parameters. As a continuation to this work, we are working on executing and studying the attacks and the effectiveness of the mitigation methods.

Acknowledgment

This work is partially supported by the National Science Foundation grants CSN 13-14891, ACI 15-35070, CNS-1513051, National Security Agency grant 2014-03124, Department of Energy grant 2015-02674 and Air Force Research Laboratory and Air Force Office of Scientific Research, under agreement number FA8750-11-2-0084. We thank, Thomas Durbin, Joseph Fullop , Jeremy Enos, Mike Showerman, Brett Bode at NCSA, Rodger Ford, Dave Dalton from Cray, and David Hardin, Warren Platt, Bruce Mikos, Randy Pankau at F&S for providing us the data and having many insightful conversations.

References

[1] B. Draney, J. Broughton, T. Declerk, and J. Hutchings, "Saving energy with freecooling and the cray xc30 " in *Cray User Group*, 2014.

[2] C. McMurtrie, L. Gilly, and T. Belotti, "Cray hybrid xc30 installation - facilities level overview," in *Cray User Group*, 2014.

[3] M. K. Patterson, S. Krishnan, and J. M. Walters, "On energy efficiency of liquid cooled hpc datacenters," in *2016 15th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, May 2016, pp. 685–693.

[4] N. HadjSaid, C. Tranchita, B. Rozel, M. Viziteu, and R. Caire, "Modeling cyber and physical interdependencies - application in ict and power grids," in *Power Systems Conference and Exposition, PSCE IEEE/PES*, March 2009, pp. 1–6.

[5] I. Dumitrache and D. I. Dogaru, "Smart grid overview: Infrastructure, cyber-physical security and challenges," in *20th International Conference on Control Systems and Computer Science*, May 2015, pp. 693–699.

[6] Y. Zhu, J. Yan, Y. Sun, and H. He, "Revealing cascading failure vulnerability in power grids using risk-graph," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 12, pp. 3274–3284, Dec 2014.

[7] C. D. Martino, Z. Kalbarczyk, R. K. Iyer, F. Baccanico, J. Fullop, and W. Kramer, "Lessons learned from the analysis of system failures at petascale: The case of blue waters," in *44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, June 2014, pp. 610–621.

[8] H. Lin, A. Slagell, Z. Kalbarczyk, P. W. Sauer, and R. K. Iyer, "Semantic security analysis of scada networks to detect malicious control commands in power grids," in *Proceedings of the First ACM Workshop on Smart Energy Grid Security, SEGS '13*. New York, NY, USA: ACM, 2013, pp. 29–34.