

A Digital Twin Framework for Testing, Evaluation and Deployment of Resilient Cyber-physical Systems

RAYMOND CHARLES BORGES HINK¹, (Member, IEEE), MARK A. BUCKNER¹, (Member, IEEE), SUPRIYA CHINTHAVALI², CHRIS A. CRAIG³, TIMOTHY S. DANIEL⁴, JOEL A. DAWSON³, MILTON NANCE ERICSON⁵, (Senior Member, IEEE), JOSHUA C. HAMBRICK¹, (Member, IEEE), PHILIP IRMINGER¹, (Member, IEEE), RYAN A. KEREKES⁵, (Member, IEEE), JUAN LOPEZ JR.⁴, (Senior Member, IEEE), KALYAN S. PERUMALLA⁶, (Member, IEEE), NICHOLAS A. PETERS⁷, (Senior Member, IEEE), STACY J. PROWELL³, (Senior Member, IEEE), VARISARA TANSAKUL⁶, CURTIS R. TAYLOR³, (Member, IEEE), BAILU XIAO¹, (Member, IEEE), and SRIKANTH B. YOGINATH⁶

¹Power and Energy Systems Group, Electrical & Electronics Systems Research Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee USA 37831

²Computational Data Analytics Group, Computational Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee USA 37831

³Cyber Defense Group, National Security Sciences Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee USA 37831

⁴National Security Sciences Directorate, Oak Ridge National Laboratory, Oak Ridge, Tennessee USA 37831

⁵Sensors and Embedded Systems Group, Electrical & Electronics Systems Research Division, Oak Ridge, Tennessee USA 37831

⁶Discrete Computing Systems Group, Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee USA 37831

⁷Quantum Information Science Group, Computational Sciences and Engineering Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee USA 37831

Corresponding author: Juan Lopez Jr., e-mail: lopezj@ornl.gov).

This work was performed in part at Oak Ridge National Laboratory, operated by UT-Battelle for the U.S. Department of Energy under contract no. DE-AC05-00OR22725. Funding was provided by ORNL's Laboratory Directed Research and Development Program. Notice: This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>). This research was also supported in part by an appointment to the Oak Ridge National Laboratory ASTRO and HERE Programs, sponsored by the U.S. Department of Energy and administered by the Oak Ridge Institute for Science and Education.

ABSTRACT

As the level of automation in critical infrastructure increases, the ability to detect cyber intrusions becomes more crucial and extremely challenging. Recent cyber attacks demonstrate the devastating and widespread affects they can have on critical infrastructure. We describe an approach to detect and prevent cyber attacks by continuously comparing the infrastructure state with a real-time “digital-twin” simulation of it. Specifically, we describe and demonstrate a Digital Twin Framework (DTF) designed specifically to detect and eventually prevent such attacks. Our framework and models are validated against experimental data from two critical infrastructure experimental emulators, first for a canal lock system and second, an electric distribution system. These systems are chosen as they have very different dynamics. The canal lock system's digital twin uses a recurrent neural network trained from the experimental data collected via the DTF. A digital twin of the transmission system is created using a commercial real-time power systems simulator and integrated into our DTF along with the hardware, embedded controllers, and live sensor data using the Open Field Message Bus data model, and publish/subscribe communication protocols. A cyber attack is used on both systems to demonstrate the DTF's detection capability.

INDEX TERMS Digital Twins, Digital Ghosts, Cybersecurity, Real-time simulation, distribution, transmission, Electric grid, Cyber physical systems, modeling

I. INTRODUCTION

Critical infrastructures, especially the electric power grid, are increasingly integrating new technologies to improve system performance, reduce operational and maintenance costs, reduce environmental impact, improve the fidelity and accuracy of monitoring systems, and improve overall reliability, interoperability, security, and resilience. The inclusion of such new technology invariably leverages ubiquitous internet communications, transforming critical infrastructure into a collection of "cyber-physical systems" (CPS) (see for example, [8]). Thus, these improvements bring with them an increase in potential attack surfaces as evidenced by numerous cyber-attack events [15], [20], [26], [37], [43], [45], [47].

In particular, as electric grid modernization efforts are carried out, careful attention must be paid to cyber security. While electric grid recovery after natural disruption (e.g., storms) is a well understood process, recovery after cyber-induced failures poses additional, potentially more difficult, challenges. Cyber-induced failure can be difficult to identify, requires an entirely different set of recovery procedures (e.g., eradication of malware on otherwise functional equipment), and requires workforce skills which are largely non-existent in today's electric grid industry. Furthermore, the expansion of the Department of Homeland Security's (DHS) National Cybersecurity and Communications Integration Centers (NCCIC) to three locations nationwide coupled with an upward trend in incident command system (ICS) incident reports and related vulnerability reports, are strong indicators that CPSs require additional protection.

To provide such a defense, we propose to leverage the concept of a digital twin (DT) operated in a larger framework of functionality. DTs were featured among the top ten strategic technology trends for 2018 by Gartner [18], and have been applied in a wide variety of industries including aeronautics, robotics, manufacturing, informatics, and healthcare [19], [36]. A key capability of a DT is the prediction of a system's response to various events. Additionally, it is desirable for a DT to host a wide range of sensors that provide insight and system self-awareness. Combined with adaptive and predictive analytic options, the goal is to realize a Digital Twin Framework (DTF) that provides an asset owner a real-time method to thwart cyber attacks or optimize performance in the presence of disturbances.

With a DTF, we extend beyond current DT concepts to address the growing need for use-case flexibility, interoperability, modularity, and real-time functionality alongside a target CPS. In addition to real-time monitoring, the DTF can be applied for replaying scenarios for forensic analysis. Alternatively, a DTF can also be used to "see into the future" by analyzing multiple alternative scenarios faster than real-time for evaluating potential courses of action and their predicted statistical impacts. Finally, this technology also supports out-of-band monitoring of a CPS, which can sometimes be the only method to detect an anomaly (for example, an accounting system discrepancy led to the discovery of a cyber intrusion to steal computation time [38]).

In this paper, we describe and illustrate a DTF's applicability and versatility with two different use cases based on real-life CPSs. The first use case is a canal lock system and the second use case is focused on electricity distribution. However, the concepts and the novel framework presented here is more broadly applicable to other CPS infrastructures similar to those defined by DHS in [9].

A. BACKGROUND AND RELATED WORK

Before describing DTFs, we review one of the key components, the DT itself. The DT concept originated within the manufacturing sector's Product Lifecycle Management (PLM) sub-discipline in 2002 as "Information Mirroring", which tied the real and virtual spaces together [22]. Later, Grieves formulated the concept of a "Digital Twin" as a virtual representation of what has been produced [23]. He placed special emphasis on the need for synchronization between the corresponding data points of the physical asset (PA) and virtual entities in a Unified Repository. Although not truly a DT, NASA has adopted the PLM approach for its space systems development process [4]. As the DT concept originated with the manufacturing sector, much of the literature has focused on this domain [22]–[24], [32], [36]. Several prominent DT use cases have emerged: 1) health analyses for maintenance activity and planning, 2) mirroring PA status, and 3) supporting decision-making through engineering and statistical tools.

As the importance of DTs grew, other organizations with different applications have adopted formal definitions

One more formal definition for a DT is a digital equivalent of a PA, process, or system running in tandem with a PA [10]. One place where this is applied is in the aerospace sector. General Electric (GE) [6] has created an advanced and functional DT that integrates analytic models (physics-based models, artificial intelligence, and enabling sensor technology) for engine components. These measure asset health, wear and performance with customer defined Key Performance Indicators (KPIs) and business objectives. Researchers at NASA [19] and the US Air Force [41] have also investigated the implementation of DTs within the design and validation processes of their respective organizations. Other research in this area places a greater emphasis on hardware-in-the-loop (HIL) for testing, validation, and integration in a naval ship power system [11].

Research projects have also explored the use of DTs or DT-analogous simulation environments for cyber security simulation and experimentation. The National Cyber Range provides a simulation of the entire Internet for the purpose of scenario testing and cyber security research [16]. Another such effort for power systems focuses on simulation to perform fault analysis, testing of protection and control functions, and the evaluation of new technologies such as the Internet of Things (IoT) [39]. The testing and tuning of network configurations and analyzing the effects of physical

disturbances on the security controls of the substation network were considered in [39].

Moving towards a DTF, there has also been research on DT environments that automatically generate a virtual machine-based system to create a DT that replicates its physical counterpart [12]. Security modules were also added in the environment, though no real-time monitoring of a PA was demonstrated. Further, they also assume that component specifications exist at a detailed level, which is not always valid. In [17], the authors describe a DT platform with an ability to combine physics-based models with the PA to predict battle space susceptibility and derive optimum performance. However, this work focused on providing only probabilistic information to support logistical planning, decision-making, and resource allocation.

Other platforms [40], [42], for example, use low-latency HIL data collection to provide high-fidelity simulations of devices for testing and validation. Other offerings are less closely coupled with hardware testbeds but are able to simulate larger aggregations of devices [21], [34]. Opal-RT, for instance, through their HYPERSIM platform, provides real-time simulation of large critical infrastructure networks through validated, object-based models [34]. However, they differ significantly from a DT. These simulation platforms are not intended to operate in tandem with an operational system, or to provide a persistent, long-term simulation against which real-world operation can be compared. Further, they often exclude factors that are important determinants of real-world system behavior, such as transmission latency, device dynamics, or communication channels. Our work overcomes these deficiencies.

B. NOVELTY

In this work, we focus on developing a DTF that can host multiple kinds of DTs. These DTs could be multiple threads of the same DT but with different parameters, or they could be distinct digital replicas of parts of a cyber-physical asset (e.g., a sensor, actuator, intelligent edge device (IED), master terminal unit (MTU), etc.). Our DTF provides advantages and uses cases that are not currently addressed by the previous literature. Specifically, our architecture provides the following benefits:

- (i) **Interoperability** - The ability to plug-in various types of simulation/emulation models interchangeably to create the DT, e.g., using simulation and HIL emulation software. Some of the software solutions can include, but are not limited to: physics-based models, machine learning algorithms, and purely statistical models. Some HIL options that can be used are Typhoon and OpalRT.
- (ii) **Intelligent Analysis** - The ability to apply meta-analysis or machine learning to DT and PA provides greater insights into the state of both the DT and the PA. This will enable better detection of anomalous behavior and also offers the potential to provide an automated response to correct the behavior.

- (iii) **Modularity and Scalability** - Enables experimenting with variants of each module, including different DTs, and it allows distributed, concurrent and resilient deployment. The described DTF is protocol and communication technology-agnostic, i.e., it can use various Pub/Sub protocols including RabbitMQ, ZeroMQ or NATS. It can also interface with many types of PA or PA devices that use different protocols such as DNP3, Modbus, IEC 61850, and compliance with rising grid interoperability standards such as OpenFMB.
- (iv) **Real-time operation** - The DTF can run at the same speed and in tandem with the PA.

C. ORGANIZATION

The rest of this paper is organized as follows: Section II gives an overview of our framework and architecture. Section III describes a DTF use case implementation on a canal lock system. Section IV details a DTF use case implementation of an electric transmission system. Section V discusses and compares the the two example use cases, and Section VI provides conclusions.

II. DIGITAL TWIN FRAMEWORK (DTF)

We developed our DTF to be a generalized, application-agnostic framework to incorporate resilient operation in CPSs. Our DTF includes a distributable module-based application that can potentially run anywhere there are available compute resources e.g., at the edge, in the fog/mist, or in the cloud. Our vision is a DTF that can rapidly and effectively introduce resilience to a variety of complex asset types and across different domains and applications. Our view of resilience spans operation under both malicious as well as unintentional negative events.

A. DEFINITIONS IN OUR DIGITAL TWIN FRAMEWORK

Digital Twin Framework: The DTF is the infrastructure necessary to enable all the features and benefits envisioned by this research in an asset-agnostic manner. The DTF provides an eco-system built using artificial intelligence, open and extensible data models and known stable interfaces that is reusable, repeatable, and amenable to rapid assembly. The DTF is shown in Figure 1.

B. FUNCTIONAL ARCHITECTURE OF THE DIGITAL TWIN FRAMEWORK

- **Physical Asset (PA):** This is a CPS that the DTF mirrors.
- **Publish-Subscribe Bus (PSB):** This enables indirect, asynchronous, and reliable communication among all the DTF modules. This bus relieves static and dynamic dependencies, provides robustness against interruptions or failures of individual modules at run time, and makes the DTF flexible and extensible to additional functionalities in the future.

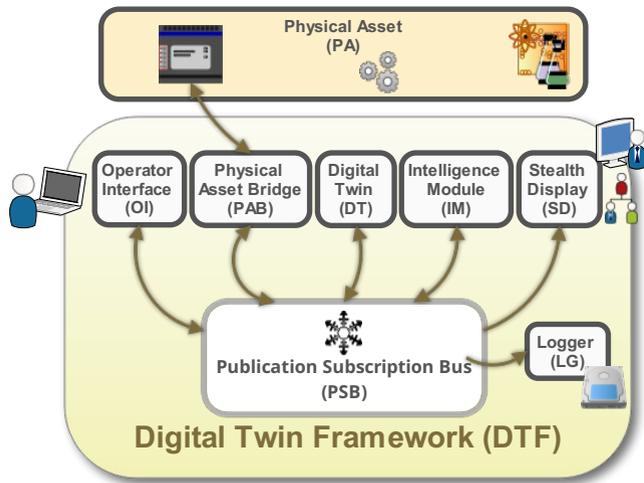


FIGURE 1: Functional Architecture of the Digital Twin Framework

- **Operator Interface (OI):** This interface provides the primary gateway by which the DTF operator launches and monitors the DTF modules.
- **Physical Asset Bridge (PAB):** The PAB provides the interface between the PA and the DTF. This can be configured in one of two ways. First, to protect the operational integrity of the PA, it could be configured to only pass information from the PA to the DTF. Alternatively, the PAB may allow bidirectional communications to enable automated control of the PA.
- **Digital Twin (DT):** An integrated multiphysics, multi-scale, and probabilistic simulation of an as-built system that uses the best available models, sensor information, and input data to mirror and predict activities and performance over the life of its corresponding PA [10].
- **Intelligence Module (IM):** The IM oversees the dynamic operation of the PA and DT in order to offer observations, alarms, triggers, and/or reaction alternatives.
- **Stealth Display (SD):** This is intended as a “God’s eye view” of the dynamic status of the PA, DT and IM, providing the stakeholders, management, and operators an idea of the dynamic status of both the PA and the DTF.
- **Logger (LG):** This is an internal functionality of the DTF that supports multiple purposes such as: (a) debugging and testing the DTF and its modules, (b) recording all transactions for service/legal purposes, and (c) after-action replays in which the DTF is exercised without the PA.

III. USE CASE: CANAL LOCK SYSTEM

We first describe application of the DTF to a physical emulation of a canal lock system. We chose this CPS because its design is conceptually simple but operates in nontrivial and nonlinear way. Canal systems have also been the focus of CPS attacks [1], [29] making it a realistic DTF use case.

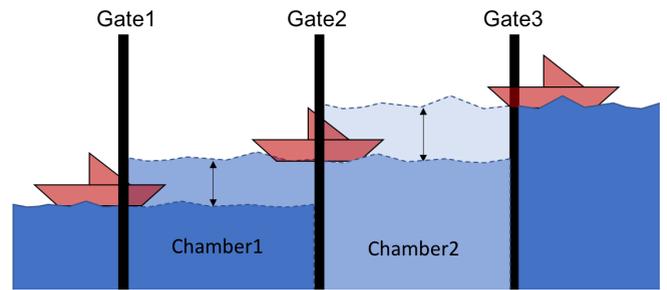


FIGURE 2: Canal lock systems allow ships to transition between waterways at different levels by adjusting water levels in chambers.

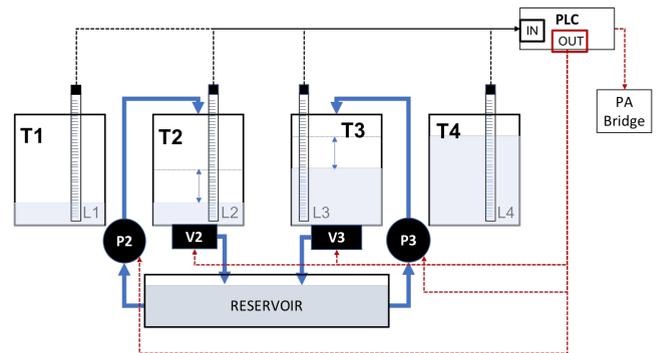


FIGURE 3: The PA for canal lock system consists of four tanks, two pumps, four depth sensors, and two valves. Three virtual gates exist in the PLC ladder logic to simulate gates opening and closing.

A. CANAL LOCK SYSTEM

Canal lock systems allow boats or ships to navigate waterways that reside on different elevations. Locks operate by gradually adjusting water levels in chambers, to either raise or lower ships. The entry and exit to a lock is controlled by a gate as shown in Figure 2. The ship either moves from a lower-level to a higher-level waterway (up scenario) or from a higher-level to lower-level waterway (down scenario). To complete each scenario, a ship travels between chambers through three different gates. The control system that ensures the safe passage of ships across the canal lock system is a CPS. The system performs the sequence of operations including opening and closing gates, and increasing and decreasing water levels in chambers.

B. IMPLEMENTATIONS

1) Physical Asset

We created a small-scale CPS system emulating the operation of canal lock system as shown in Figure 3. The emulation model of the canal lock system was comprised of four 14cm × 14cm × 30.48cm acrylic tanks. The first and last tanks in the canal lock system represent the lower and upper water bodies that the two intermediate tanks connect. The two intermediate water tanks (T2 and T3 in Figure 3) act as the locks and are filled from a water reservoir using two 5V

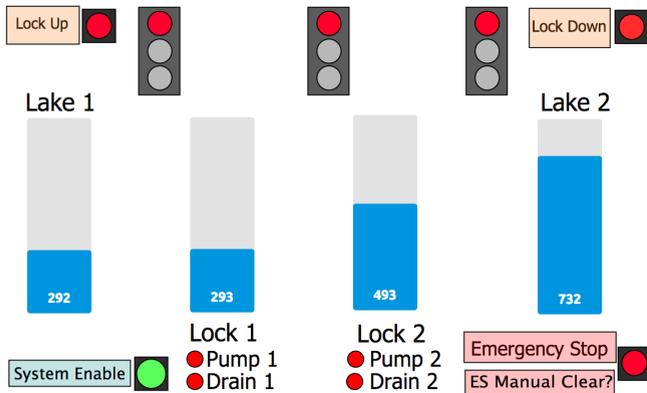


FIGURE 4: The HMI that controls the canal lock system is shown.

DC water pumps [46]. The 5V DC power supply is used to activate water pumps for tanks T2 and T3. Water in tanks T2 and T3 are drained to the reservoir using Schneider Electric modulating valves [13]. The drain rate of the valves is determined by a control voltage ranging from 0-10V DC. We determine the water level in the tanks using tape sensors that provide a resistance value proportional to the depth between 0-30.48cm. We simulate gate and ship movement between canals using logic programmed directly into the PLC.

We coordinate the control of the canal lock system using an Allen-Bradley Micrologix 1100 PLC [7] with an additional Allen-Bradley 1762-IF2OF2 extension slot. The tape sensors are configured as analog inputs to the PLC, the pumps as digital control outputs, and the valves as analog control outputs. We power all components in the canal lock system using a 24V power supply.

We use Rockwell Automation's RSLogix 1100 Micro Starter Lite software to develop the PLC Ladder Logic to control the canal lock system. To interface with the PLC, an operator uses a custom mySCADA human machine interface (HMI) [30]. A snapshot of the HMI is shown in Figure 4. The canal lock system's operation is driven by operator input to the HMI including the ability to (1) enable or disable control of the system, (2) start a Lock Up scenario, (3) start a Lock Down scenario, (4) issue an emergency stop, and (5) clear an emergency stop status notification. During the operation, the water level for each tank, the gate statuses, and the actuation statuses are relayed to the HMI. This information is received by the HMI via a Modbus connection to the PLC.

2) DTF

We now describe how we implement each component of the DTF for the canal lock system.

- **Publish-Subscribe Bus (PSB):** The Pub-Sub bus system was realized using the ZeroMQ library. The ZeroMQ daemon accepts connections from each of the DTF components and handles the passage of messages across different components.

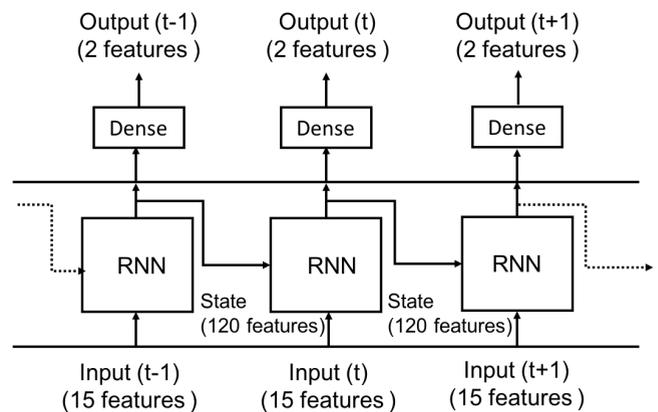


FIGURE 5: The RNN model for the canal lock system DT accepts 15 input features and predicts the water levels of T2 and T3.

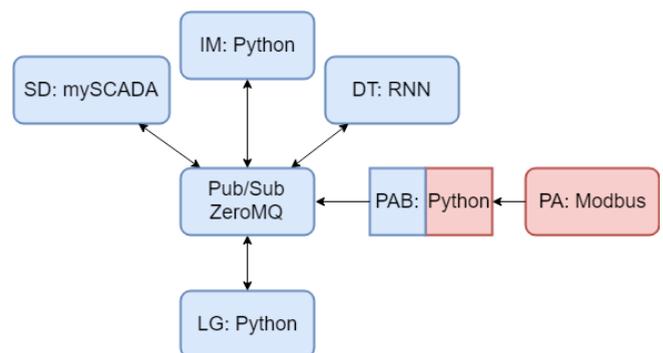


FIGURE 6: The DTF implementation for the electric grid case study is shown.

- **Physical Asset Bridge (PAB):** The PAB for canal lock system collects the system's state by polling the PLC registers over Modbus. When data is received from the PLC, the PAB generates messages conforming to defined topics and publishes them to the PSB. The following 15 features are collected by the PAB and timestamped before being transmitted to the PSB:
 - Tank levels: Water level of each tank in inches
 - Gate status: Status of each gate (opening, opened, closing, closed)
 - Pump status: Binary status of each pump (on, off)
 - Valve status: Binary status of each valve (on, off)
 - Valve voltage: Voltage level being applied to each valve (0-9 Volts)
 - Ship status: Ship position in canal/tank (T1, T2, T3, T4)
 - Direction: Up or down scenario

Additionally, the PAB listens to topics related to initiating up and down scenarios for data collection. However, it does not support all controls provided via the HMI (e.g., emergency stop).

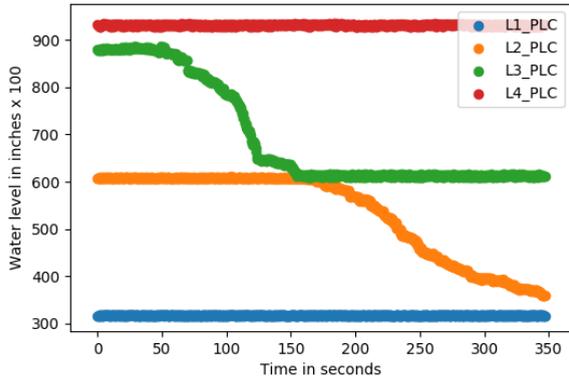


FIGURE 7: The water level for each tank under a normal down scenario is shown.

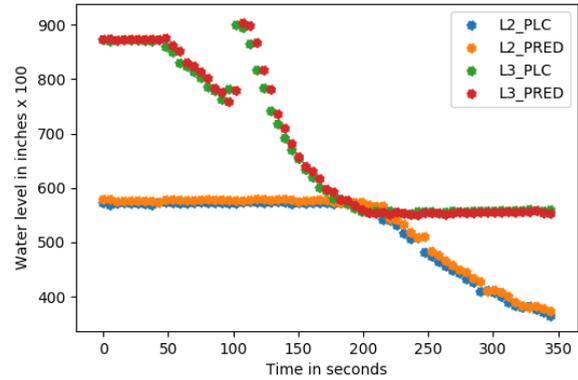


FIGURE 8: An attack scenario is shown. It results form an abrupt and unexpected increase in the water level of tank T3. This physical attack causes the DT to deviate from the Canal lock system PA.

- **Digital Twin (DT):** We use a recurrent neural network (RNN) [14] model as the DT for the canal lock system. We used 120 units resulting in an RNN model with learning parameters of 16,320 $((15 \text{ features} + 120 \text{ units}) \times 120 + 120 \text{ bias})$. The RNN layer is followed by a fully-connected dense layer with 120 inputs and two outputs, which learns 242 parameters $(240 \text{ weights} + 2 \text{ bias})$. The RNN’s prediction is a single time step ahead of the PA allowing us to determine the trajectory of the PA’s state (i.e., the dynamically changing tank level). The RNN is re-synced using PA data after each prediction. We collected and normalized the training data. The RNN model was trained on normal operation data with 15 features for 100 epochs resulting in errors less than 0.1 percent in prediction. A high-level overview of the RNN is shown in Figure 5.
- **Stealth Display (SD):** The stealth display is a web page that visualizes real-time data from both PA and DT using JavaScript. The SD visualizes water level for each tank from the PA and the DT.
- **Operator Interface (OI):** The operator interface is written as a Python script that is used to initiate up and down scenarios on the PA by publishing messages to the PAB. The OI was also used to automate the data collection process for training the RNN.
- **Logger (LG):** The logger is implemented as a Python script that uses a wildcard to subscribe to all topics. The logger supports writing to both a CSV log file as well as storing data in a database.
- **Intelligence Module (IM):** As a proof of concept, the IM calculates the difference in water levels for tanks in the PA and DT. The IM publishes an alert along with the timestamp to the PSB indicating that an anomaly has occurred if the difference surpassed a predetermined threshold. The alert is received and visualized by the SD.

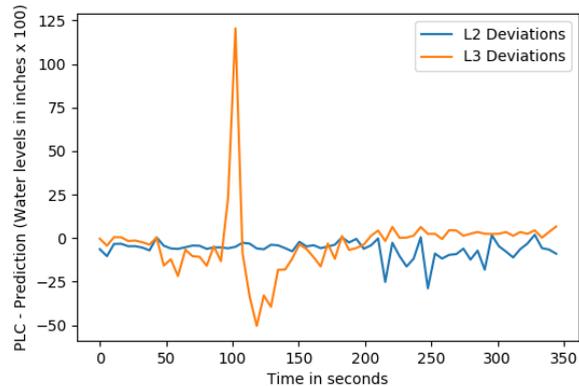


FIGURE 9: Deviation between DT and PA as calculated by the IM, in the Canal lock system

C. USE CASE TAKEAWAY: CPS ATTACK DETECTION

Applying the DTF to the canal lock system allows us the capability of detecting attacks against the CPS. As a simple illustration to demonstrate the DTF’s ability to detect anomalies, we performed a physical attack on tank T3 during a down scenario. The water level in tank T3 during a normal down scenario is shown in Figure 7. In Figure 7, we note that variations in sensor data and calibration cause fluctuation in the visualization of the data. In this scenario, the water levels in both tanks T2 and T3 drop from their initial water levels to accomplish the movement of the ship from tanks T4 to T1. We created an anomaly in the operation of the canal lock system by introducing a physical attack that abruptly increases the water level in tank T3 by approximately a liter while tank T3 is being drained. The single-point RNN prediction model momentarily deviates as seen in Figure 8. This deviation is captured by the IM as shown in Figure 9 and is detected as an anomaly.

IV. USE CASE : ELECTRIC GRID

The second use case demonstrates application of the DTF to a power system asset. The intent of this case study is to explore the flexibility of the DTF for power systems application and to demonstrate the use of commercial, off-the-shelf modeling tools as a method for DT implementation. Potential uses for this is detecting abnormal operating conditions including both maintenance issues and cyber-physical attacks, acting as a platform to understand the vulnerability of the grid to cyber-physical attack, real-time analysis with predicted statistical impacts to determine alternative courses of action and, finally, providing a platform for integration to a larger scale model.

In this case study, the PA consisted of a power source, a distribution line, a controllable resistive load, and two smart relays. The PA was implemented using results from the development of ORNL's Software-defined Intelligent Grid Research Integration and Development (SI-GRID) platform [33]. SI-GRID is a low voltage grid emulator that is open, scalable, extensible, dynamic, and reconfigurable. SI-GRID originally consisted of multiple programmable power electronic sources and programmable loads that can be reconfigured by opening and closing any of the multiple contactors used for system reconfiguration by physically connecting any device to any microgrid. An example of the SI-GRID platform overall topology can be found in [33].

In order to investigate the feasibility of using a commercial simulation platform as the DT, Typhoon HIL [42] was chosen. Typhoon HIL is a commercially available HIL system with a native Python interface. This interface will be described in detail below.

Two separate implementations of NATS were used by having one for the PA, and one for the DTF. NATS is a PSB that uses a central exchange server to facilitate and route communication between publishing data and subscribing clients [31]. NATS was chosen due to ease of implementation and performance requirements for modeling and simulating power systems. The specific details of each NATS implementation will be discussed in relevant sections.

A. OPERATIONAL SCENARIO

The operational scenario for the electric grid case study was a radial distribution feeder serving a varying load. In a radial feeder, power is supplied from a single source which is fed via distribution lines to downstream loads. A typical distribution line has a number of switches for isolation and reconfiguration in response to abnormal conditions such as faults or maintenance outages. These can be either mechanical switches that require human intervention, or protective relays which will operate to protect the system. This type of configuration is common for electrical distribution systems in the United States. Load on the system is controlled to match a pre-defined profile to enable exploration of the system under different load conditions.

B. IMPLEMENTATIONS

1) Physical Asset

In this use case, a simple 60 Hz AC electrical system was constructed using SI-GRID components. The system consisted of a 3-phase 24 V_{L-L} source, three 3-phase line emulators, a controllable resistive load bank, and two smart relays, as shown in Figure 10.

A National Instruments (NI) single board RIO (sbRIO) provided control and data acquisition for the load bank. Each load bank has a maximum load of 300 W per phase at 24 V_{L-L} and is controllable down to 1 W. The sbRIO connects directly to the relaying and measurement board through a General Purpose Inverter Controller (GPIC) [5]. This provides high-sample-rate of voltage and current measurements at the load. The load bank also supports individual phase control for testing of unbalanced systems and unbalanced load profile playback. The control and monitoring interface used for the load bank was Modbus TCP.

Similarly, the control and instrumentation for each switch is provided by an NI sbRIO. The sbRIO provided highly sampled voltage and current readings on both sides of the switch, as well as local and remote control of the switch state (open/close). Like the load bank, the control and monitoring interface the switch was Modbus TCP.

The three 3-phase, 5-wire line emulators adds resistive and inductive impedance to the system. Each line emulator consists of 10 series inductors per phase, totalling 220 μH at an X/R ratio of approximately 8. The line emulator allows for accurate emulation of balanced and unbalanced conditions. Each line emulator represents approximately 1,000 ft of medium-voltage distribution line. Line and phase capacitance can also be added, but none were used in this scenario.

An OpenFMB-based NATS pub/sub protocol provided data acquisition and control for the PA system. OpenFMB is an emerging data standard developed by utilities for distributed communication and control of power systems assets [35]. Translators polled Modbus data from different PA assets, loaded the data into the appropriate OpenFMB data structures, and published the OpenFMB data over the NATS pub/sub protocol. Similarly, translators subscribed to relevant OpenFMB commands on the NATS network and translated them into appropriate Modbus commands for the devices. The translators were built into Docker [28] containers which were hosted on SEL 3360 hardened computers running Linux. An overview of the overall system and communication architecture is shown in Figure 10.

2) Digital Twin Framework

Due to changes in physical system characteristics, measurement sampling rates, data volume, latency requirements, and other considerations, the DTF implementation for the electric grid has a number of differences compared to the previous case study. The specific implementation of the DTF for the electric grid case study is described below. An overview of this DTF implementation is shown in Figure 11.

- **Publish-Subscribe Bus (PSB):** For the performance and configuration reasons described above, a NATS

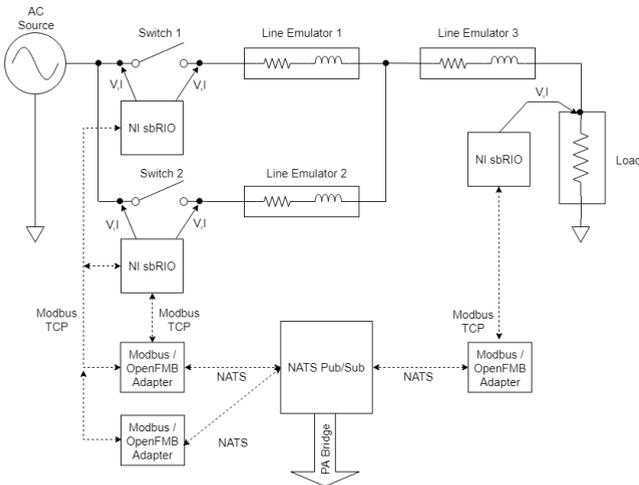


FIGURE 10: A schematic of an electric grid physical asset is shown.

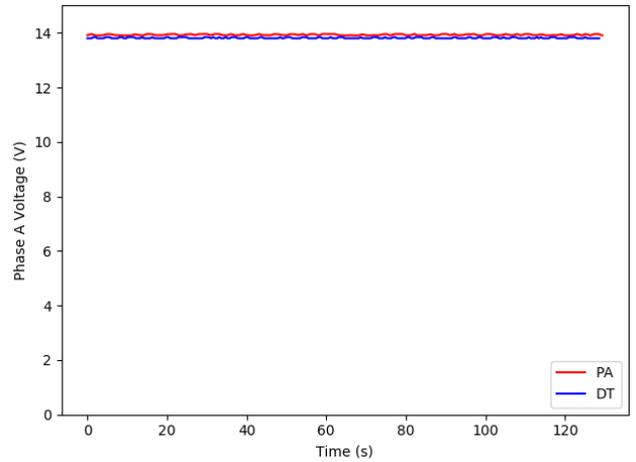


FIGURE 12: Phase A Voltages for the DT and PA are shown.




FIGURE 11: The DTF implementation for the electric grid case study is shown.

pub/sub protocol was also used for communication within the DTF. The DTF NATS network was isolated from the PA NATS pub/sub protocol. Relevant information to the DTF was facilitated through the PAB.

- Physical Asset Bridge (PAB):** For this implementation, the PA Bridge subscribed to topics on the PA NATS network, translated the OpenFMB data into a DTF specific data model and published the data to the DTF NATS protocol. The PAB was implemented to work as a NATS client for both PA and DTF. As a client to PA NATS, the PAB was responsible for subscribing data from PA for OpenFMB messages. As a client to DTF NATS, it generated new message based on specified DTF data model after parsing OpenFMB messages. The DTF data model is a heavily stripped down version of OpenFMB data model and uses Protobuf [44] to generate code to communicate messages across the internal DTF NATS protocol.
- Digital Twin (DT):** Owners and operators of physical systems such as the electric grid often build and maintain models of their infrastructure. One important

consideration for this case study was demonstrating that the DTF was flexible enough to incorporate these models without need for extensive development of custom tools or simulation engines. The DT for this use case was built using the Typhoon HIL commercial modeling environment.

Real-time simulators such as Typhoon HIL are used extensively in the power industry for design and evaluation of electric systems. Typhoon HIL was selected for this application because of the extensive and well-documented Python API. Using existing Python NATS libraries, the API allowed for simple integration with the Python-based DTF implementation.

- Stealth Display (SD):** The Stealth Display shows graphs of voltage and current from the DT and PA, as well as differences between DT and PA values in real-time as streaming displays using Bokeh library. Bokeh library is a Python-based visualization software suite which can render given parameters on an interactive web page [3].
- Logger (LG):** The Logger, as in the previous case, encapsulates the events published to the DTF, and stores them to assist in development of the DT by tuning the parameters of the model and perform incident-response analysis. The data set is stored in chronological order for debugging and scenario playback.
- Intelligence Module (IM):** The Intelligence Module was implemented as a script that applies an arithmetic model to derive the difference between the voltage and current values produced by the PA and DT. The IM then publishes this data to the DTF for display or logging purposes.

C. DT TAKEAWAY: ENTERPRISE INTEGRATION

Integrating novel detection software is a difficulty for existing power systems. A new intrusion detection mechanism commonly requires trained technicians for testing, configuration,

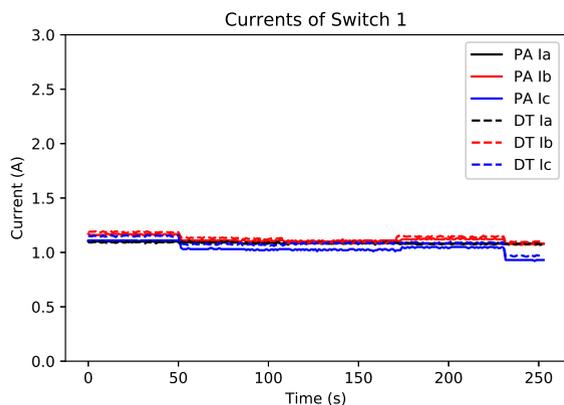


FIGURE 13: Current of Switch 1

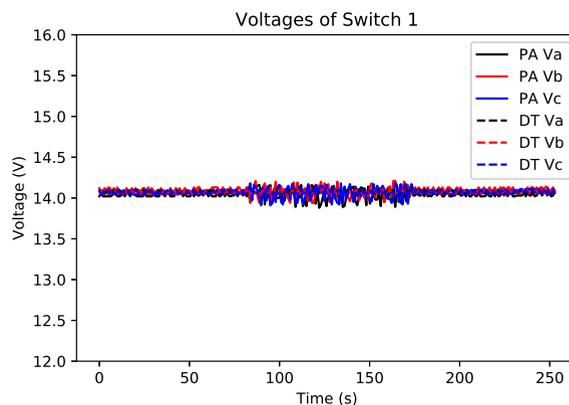


FIGURE 15: Voltage of Switch 1

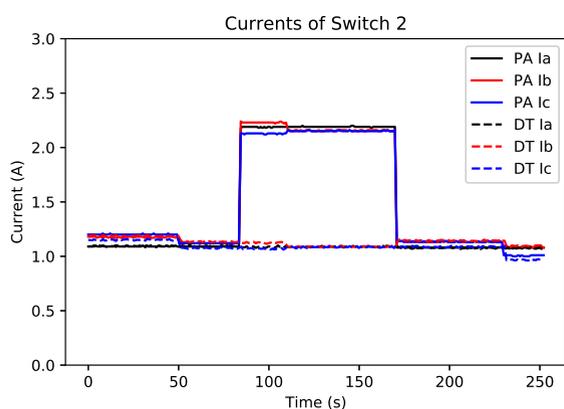


FIGURE 14: Current of Switch 2

and installation. The DTF eases this burden by using modular components. As previously mentioned, Typhoon was used as the DT in place of novel software. One of the benefits of using hardware of this caliber is its accuracy. Figure 12 shows that Typhoon closely models our existing PA environment with difference of 0.1V under normal operations. While some configuration effort was required for it to communicate to the NATS server, the effort level was far less than that of developing and validating modeling software.

One potential barrier to industry adoption of DTs is the time and expense needed to create models sufficient for use in DT applications, particularly for more complex systems. By using existing commercial tools and models developed by industry subject matter experts as the DT, this burden can be reduced.

D. USE CASE TAKEAWAY: CYBER-ATTACK DETECTION

We chose a false data injection attack to demonstrate the capabilities of the DTF. This is consistent with the type of attacks discussed in [25], [27] and [2] as particularly hard to detect. In this specific example attack, current was increased as can be seen in Figure 14, yet the attacker makes the voltage appear nominally unchanged as seen in Figure 15.

By monitoring the the entire system, one is able to identify the false view of the voltage data form looking out-of-band at the current data.

V. DISCUSSION

Our two case studies demonstrate the generality and modularity of the designed conceptual DTF. We applied the DTF to a canal lock system using real-world CPS sensors and actuators. The DTF enabled us to observe the canal lock systems operation and build an RNN-based DT of the PA that accurately model the system's interactions with high fidelity. While operating, we performed an attack on the canal lock system by unexpectedly increasing the water level in a draining canal. Indeed, the IM detected the deviation, thus detecting an attack on the CPS.

In the second case study, we adapted the DTF to support a low voltage, re-configurable electric grid using OpenFMB for communicating control and state information. For adaptation, no alterations were needed to the design of the DTF. The flexibility of the framework enabled rapid adaptation to a different cyber-physical domain. Given the industry support for modeling in the grid space, we chose to use Typhoon HIL, a commercial hardware-in-the-loop modeling solution, and replaced the PSB with NATS. With these alterations, the DTF was again able to properly support an electric grid PA. This case study demonstrates that the DTF described in this paper is flexible enough to incorporate industry standard tools in DT applications. Furthermore, utilities can leverage investments in simulation hardware, such as HIL or real-time simulators, for use as DTs.

VI. CONCLUSION

DTs have proven promising in the design phase of PAs as well as in their preventive maintenance and in optimizing performance throughout the product life-cycle. Here we have shown it is possible to rapidly apply the framework to different types of CPSs while leveraging existing Commercial-Off-The-Shelf (COTS) hardware and software. We then demonstrate an application that seeks to improve the cyber

resilience of CPSs. We have also introduced a framework that enables DT scalability and rapid innovation through modularity, allowing interchange of the technologies used for creating the digital copy as well as those used for analytics and anomaly detection. Finally, we have shown how a DT can be used as a compensating control mechanism for CPS security while running in real-time alongside the CPS parent.

The most important aspect of our work that will be the focus future efforts is the IM. Much of the work supporting this paper necessarily focused on the development of other components and the IM represented here does not indicate what is possible. The IM within the architecture will support the inclusion of multiple DTs without overwhelming the PA operators with data. From a machine learning perspective, this component has the potential to demonstrate artificial intelligence applied to a real-world problem. It may also provide a platform for the machine learning algorithms to potentially respond to PA anomalies in real-time. This will greatly impact the resilience of the critical infrastructure systems.

We anticipate many additional use cases can be implemented using our approach. These new uses will include an automated learning/training mode for the machine learning models within both the DT and the IM. In addition, we envision a related “tuning” mode that allows the IM to modify the DT over time to better match the PA. We also anticipate a “faster-than-real-time” advisory mode that can provide what-if scenario analysis. This analysis would be useful for providing suggested courses of action for specific events, and when collected, will provide a “playbook” for asset operators to respond to incidents.

Additionally, a new, enhanced, OI could help commission or provision a DTF. The logistics and human factors related to the deployment of a complex DTF within a CPS must be done consistently and correctly and this interface will support this process. Also, the large number of available communication protocols and system architectures within CPSs make the PAB one of the most critical modules for the adoption of our DTF. Tools to assist in modifying the PAB to communicate with a new PA need to be developed. These tools should improve the security of the PAB module as well as reduce the cost of deploying the DTF. Ideally, the PAB will minimize the amount of protocol-specific customization which is needed and it will maximize the variety of components so that it more generally applies to other PAs. One way this may be accomplished is through the incorporation of industry data model standards such as OpenFMB.

The proposed DTF can be refined and extended in multiple ways. For example, it can be applied to other use cases such as digital forensics and “what-if” scenario system exploration. We expect that the DTF will be helpful for demonstrating interoperability within the electric grid space as well as in other industrial control systems, especially those using the OpenFMB standard. We also expect the execution of multiple DTs simultaneously to be useful and that the IM will play a key role in reconciling predictions from each DT.

Finally, we hope extensions we have described will enable more thorough detection and differentiation of cyber-attacks from natural physical phenomena occurring to CPSs.

ACKNOWLEDGMENT

We acknowledge Dwayne Bradley, Stuart Laval, David Lawrence, Alexander Melin, Bruce Warmack, and Alex Williams for technical support at various stages of this research.

REFERENCES

- [1] Saurabh Amin, Xavier Litrico, S Shankar Sastry, and Alexandre M Bayen. Stealthy deception attacks on water scada systems. In Proceedings of the 13th ACM international conference on Hybrid systems: computation and control, 2010.
- [2] J. M. Beaver, R. C. Borges-Hink, and M. A. Buckner. An evaluation of machine learning methods to detect malicious scada communications. In 2013 12th International Conference on Machine Learning and Applications, volume 2, pages 54–59, Dec 2013.
- [3] Welcome to Bokeh — Bokeh 0.13.0 documentation, Jun 2018. [Online; accessed 12. Sep. 2018].
- [4] Pamela Caruso, Daniel Dumbacher, and Michael Grieves. Product lifecycle management and the quest for sustainable space exploration. In AIAA SPACE 2010 Conference & Exposition, 2010.
- [5] Save time and resources with the ni compactrio general purpose inverter controller (gpic) national instruments. <http://www.ni.com/white-paper/14163/en/>. Online; accessed 05-September-2018.
- [6] General Electric Company. GE digital twin: Analytic engine for the digital power plant. White paper, 2016.
- [7] MicroLogix 1100 Programmable Controllers. Rockwell automation. http://literature.rockwellautomation.com/idc/groups/literature/documents/um/1763-um001_-en-p.pdf. [Online; accessed 24-August-2018].
- [8] Critical infrastructure security and resilience. <https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil>.
- [9] Critical infrastructure sectors | Homeland Security. <https://www.dhs.gov/critical-infrastructure-sectors>. (Accessed on 09/13/2018).
- [10] DoD digital engineering definitions. https://www.acq.osd.mil/se/initiatives/init_de_def.html.
- [11] Christian Dufour, Zareh Soghomonian, and Wei Li. Hardware-in-the-loop testing of modern on-board power systems using digital twins. In 2018 International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM), pages 118–123. IEEE, 2018.
- [12] Matthias Eckhart and Andreas Ekelhart. Towards security-aware virtual environments for digital twins. In Proceedings of the 4th ACM Workshop on Cyber-Physical System Security, pages 61–72. ACM, 2018.
- [13] Schneider Electric. Erie™ VM Series Poptop™ Series Modulating Valves. <https://iportal2.schneider-electric.com/Contents/docs/F-26801-8>. PDF. [Online; accessed 24-August-2018].
- [14] Jeffrey L Elman. Finding structure in time. *Cognitive science*, 14(2):179–211, 1990.
- [15] Nicolas Falliere, Murchu, and Eric Chien. W32.Stuxnet dossier. Symantec Security Response online report, February 2011.
- [16] Bernard Ferguson, Anne Tall, and Denise Olsen. National cyber range overview. In Military Communications Conference (MILCOM), 2014 IEEE, pages 123–128. IEEE, 2014.
- [17] Thomas C. Fu. Navy platform digital twin. url: <http://onlinepubs.trb.org/onlinepubs/mb/2017Spring/fu.pdf>, May 2017.
- [18] Inc. Gartner. Gartner identifies the top 10 strategic technology trends for 2018. <https://www.gartner.com/en/newsroom/press-releases/2017-10-04-gartner-identifies-the-top-10-strategic-technology-trends-for-2018>, October 2017. Online; accessed 06-September-2018.
- [19] Stargel Glaessgen. The digital twin paradigm for future nasa and u.s. air force vehicles. In 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference. American Institute of Aeronautics and Astronautics, April 2012.
- [20] Andy Greenberg. ‘Crash Override’: The malware that took down a power grid. <https://www.wired.com/story/crash-override-malware/>. Online; accessed 11 Sep 2018.

- [21] GridLAB-D. <https://www.gridlabd.org>. Online; accessed 11-September-2018.
- [22] Michael Grieves. Product lifecycle management. www.egr.msu.edu/classes/ece480/goodman/PLMinEngineering.ppt, 2002.
- [23] Michael Grieves. Digital twin: Manufacturing excellence through virtual factory replication. White paper, 2014.
- [24] Michael Grieves and John Vickers. Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In Franz-Josef Kahlen, Shannon Flumerfelt, and Anabela Alves, editors, *Transdisciplinary Perspectives on Complex Systems: New Findings and Approaches*. Springer, Cham, 08 2016.
- [25] R. C. Borges Hink, J. M. Beaver, M. A. Buckner, T. Morris, U. Adhikari, and S. Pan. Machine learning for power system disturbance and cyber-attack discrimination. In 2014 7th International Symposium on Resilient Control Systems (ISRCs), pages 1–8, Aug 2014.
- [26] Robert A. Lee, Michael J. Assante, and Tim Conway. Analysis of the cyber attack on the ukrainian power grid: Defense use case. Technical report, Electricity Information Sharing and Analysis Center (E-ISAC), Washington, DC, United States, 2016.
- [27] Beibei Li, Rongxing Lu, Kim-Kwang Raymond Choo, Wei Wang, and Sheng Luo. On reliability analysis of smart grids under topology attacks: A stochastic petri net approach. *ACM Transactions on Cyber-Physical Systems*, 3(1), August 2014.
- [28] Dirk Merkel. Docker: Lightweight linux containers for consistent development and deployment. *Linux J.*, 2014(239), March 2014.
- [29] Bill Miller and Dale Rowe. A survey scada of and critical infrastructure incidents. In *Proceedings of the 1st Annual conference on Research in information technology*, 2012.
- [30] myDESIGNER - project creation tool. <https://www.myscada.org/mydesigner>. Online; accessed 24-August-2018.
- [31] NATS. Nats introduction. <https://nats.io/documentation/>. Online; accessed 11-September-2018.
- [32] Elisa Negri, Luca Fumagalli, and Marco Macchi. A review of the roles of digital twin in cps-based production systems. *Procedia Manufacturing*, 11:939–948, 2017.
- [33] B. Ollis, P. Irminger, M. Buckner, I. Ray, Dan King, A. Herron, B. Xiao, R. Borges, M. Starke, Yaosuo Xue, and B. Macleery. Software-defined intelligent grid research integration and development platform. In 2016 IEEE Power Energy Society Innovative Smart Grid Technologies Conference (ISGT), pages 1–5, Sept 2016.
- [34] OPAL-RT. Opal-rt. <https://www.opal-rt.com>. Online; accessed 11-September-2018.
- [35] OpenFMB™ - Collaboration Site. <https://openfmb.github.io/#community>, Jun 2017. [Online; accessed 13. Sep. 2018].
- [36] Procedia Manufacturing. Micro Manufacturing Unit and the Corresponding 3D-Model for the Digital Twin, Stockholm, Sweden, May 2018. Elsevier.
- [37] Michael Riley, Jennifer A. Dlouhy, and Bryan Gruley. Russians are suspects in nuclear site hackings, sources say. <https://www.bloomberg.com/news/articles/2017-07-07/russians-are-said-to-be-suspects-in-hacks-involving-nuclear-site>, July 2017. Online; accessed 11 Sep 2018.
- [38] Clifford Stoll. *The Cuckoo's Egg*. Knopf Doubleday Publishing Group, 1989.
- [39] Eniye Tebekaemi and Duminda Wijesekera. Designing an iec 61850 based power distribution substation simulation/emulation testbed for cyber-physical security studies. In *Proceedings of the First International Conference on Cyber-Technologies and Cyber-Systems*, 2016.
- [40] RTDS technologies. Rtds. <https://www.rtds.com/>. Online; accessed 11-September-2018.
- [41] Eric Tuegel, Anthony Ingrassia, Thomas Eason, and Michael Spottswood. Reengineering aircraft structural life prediction using a digital twin. *International Journal of Aerospace Engineering*, 2011:154798, 2011.
- [42] Typhoon HIL. <https://www.typhoon-hil.com>. [Online; accessed 13. Sep. 2018].
- [43] US-CERT. Russian government cyber activity targeting energy and other critical infrastructure sectors. <https://www.us-cert.gov/ncas/alerts/TA18-074A>, March 2018.
- [44] Kenton Varda. Protocol buffers: Google's data interchange format. Technical report, Google, 6 2008.
- [45] Danielle Walker. <https://www.scmagazine.com/home/news/havex-malware-strikes-industrial-sector-via-watering-hole-attacks/>, June 2014. Online; accessed 11-September 2018.
- [46] https://www.marinedepot.com/Reef_Octopus_Diablo_Variable_Speed_Water_Pump_DC_3500_Over_1000_Gallons_Per_Hour_Submersible_Aquarium_Pumps-Reef_Octopus-CV52043-FIWPSBTO-CV52029-vi.html. Online; accessed 24-August-2018.
- [47] K. Zetter. Inside the cunning, unprecedented hack of Ukraine's power grid. <https://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/>, March 2016. Online; accessed 11 Sep 2018.

BIOS



RAYMOND C. BORGES HINK (M'09) received the B.S. degree in computer engineering from Universidad del Turabo, Gurabo, Puerto Rico, in 2010 and the M.S. in computer science from the Polytechnic University of Puerto Rico, in 2012, during which he was a Nuclear Regulatory Commission fellow.

In 2013 he became a Research Associate in the Computer Data Analytics group at Oak Ridge National Laboratory and in 2017 he was hired into the Electrical and Electronics Systems Research Division as a Power Systems Cyber Security Researcher. Here he has co-authored several papers in the detection and classification of natural and cyber-disturbances in cyber-physical testbeds, including "Machine learning for power system disturbance and cyber-attack discrimination" and "An evaluation of machine learning methods to detect malicious SCADA communications".

He holds and has taught courses for the following certifications: CompTIA A+, CompTIA Network+, CompTIA Security+, Original Equipment Manufacturer (OEM) and holds the Microsoft Certified Technology Specialist (MCTS) and Certified Scrum Master (CSM) certifications. He is currently a PhD student at West Virginia University and was an SREB-State Doctoral Scholars Program fellow.



MARK A. BUCKNER (M'11) Mark received his B.A. degree in physics and psychology from Carson-Newman College and his M.S. and Ph.D. in nuclear engineering/applied artificial intelligence from the University of Tennessee, Knoxville. His main areas of research are: resilient cyber-physical systems; machine learning and artificial intelligence; and signal processing. He is currently a Senior Research Scientist at the Oak Ridge National Laboratory and Group Leader of

the Power and Energy System Group. Dr. Buckner has over 31 years of experience in signal processing, machine learning and computational intelligence. He is also a Certified Scrum Professional®, Certified ScrumMaster®, Certified Scrum Product Owner®, Certified Scrum In Hardware Trainer, a member of the OpenFMB™ Cyber Security Working Group, and the Industrial Internet Consortium.



SUPRIYA CHINTHAVALI Supriya Chinthavali received the B.S. degree in Electronics and Communications from Vishweshwaraiah Technological University, India, in 2005, an M.S. in Automotive Embedded Systems from Manipal Institute of Higher Education, India in 2008 and an M.S. in Computer science and Engineering from Georgia Tech Institute of Technology, Atlanta, Georgia, USA in 2015. She is an R&D Staff Member at Oak Ridge National Laboratory in the Computational

Data Analytics group. She has served as a principal investigator on several projects related to Energy sector's situational awareness, policy analysis and in understanding cascading impacts on critical infrastructures during extreme events. Her research interests include application of machine learning, data analytics and visualization techniques in various domains such as power system, healthcare etc. She is an ACM member and a Certified Scrum Master (CSM).



CHRISTOPHER A. A. CRAIG received a B.S. and M.S. degree cum laude in computer science at University of Tennessee, Knoxville, Tennessee in 2010 and 2018 respectively.

In 2010, he entered the work force as a system analysts for Unum programming bill reconciliation software. In 2011, he took his interests to Cisco Systems as a software security engineer. After five years of security experience, he spent his early career performing penetration testing and protocol analysis for a variety of Cisco products. His expertise extended into cloud networks, distributed computing, and trusted computing environments. In 2015, his talents took him to Oak Ridge where he's offered his services in distributed system computing for GPU-enabled 3D point cloud rendering and adaptive network security solutions for the grid.



JOEL A. DAWSON received the B.A. degree magna cum laude in Communication from Messiah College in 2008, and the M.S. degree in Computer and Information Sciences from the University of South Alabama in 2017.

Prior to graduation, he interned at ICS-CERT at Idaho National Laboratory in 2016. After graduation, he went to work as a Post-Master's Research Associate at Oak Ridge National Laboratory in the Cyber Defense Science Group. While there, he has contributed to projects in malware characterization and detection, digital twins, and IoT security.

Mr. Dawson is a member of Upsilon Pi Epsilon, Lambda Pi Eta, and is a graduate of the NSF CyberCorps Scholarship for Service program. He is the author of six publications and a Certified Scrum Master (CSM).



M. NANCE ERICSON (M'87, SM'14) received the B.S. degree from Christian Brothers University in Memphis, TN, USA, in 1987, and the M.S. and Ph.D. degrees from the University of Tennessee, Knoxville, TN, USA, in 1993 and 2002, respectively, all in electrical engineering. He is a Distinguished R&D Staff Member at the Oak Ridge National Laboratory in Oak Ridge, TN, USA, and serves in Adjunct Faculty and Joint Faculty Appointment roles with the Department of Electrical Engineering and Computer Science at the University of Tennessee, Knoxville. His research interests include low power, mixed-signal integrated circuits and systems, and harsh environment electronics including SiC gate drivers and radiation hardened circuits. He has over 130 technical publications and holds 13 patents in the area of electronics and electronics systems including high-temperature measurement, integrated photo-spectrometers, and implantable sensing methods.



JOSHUA HAMBRICK (M'07) is an R&D Staff Member at Oak Ridge National Laboratory in the Power and Energy Systems group. His research focuses on power systems modeling and simulation, power systems protection and microgrid design and control. He received his BS, MS and PhD in Electrical Engineering from Virginia Tech.



PHILIP IRMINGER (M'15) is a member of the Power and Energy System Group at ORNL. He received a B.S. in Electrical Engineering from the University of Tennessee. He has been working at ORNL for over 5 years performing research in different areas of power systems analysis including energy storage, microgrids, hardware testing, and protection modeling. This includes being on the integration, design, and build team for the 1st and 2nd version of the SI-GRID platform. This

includes working with a first of its kind power flow control device which utilizes a saturable core reactor to regulate the power flow. Also working with a novel secondary-use energy storage system through a partnership with General Motors and ABB. He has also contributed to an overall systems integration project which incorporates the use of control hardware in the loop to test and validate control schemes by incorporating many real-time systems, and then helping deploy these control schemes using the ORNL developed microgrid controller (CSEISMIC).



RYAN A. KEREKAS (M'18) leads the RF Communications and Intelligent Systems group at Oak Ridge National Laboratory in the Electrical and Electronic Systems Research Division. He received the Ph.D. degree in Electrical and Computer Engineering from Carnegie Mellon University in 2007 and subsequently joined Oak Ridge National Laboratory as a postdoc. Since 2009, he has been a research staff member at ORNL. His research interests include signal and image

processing, machine learning, and systems integration.



JUAN LOPEZ, JR. (M'04-SM'13) received the B.S. degree from the University of Maryland, College Park, MD, USA, in 2001, the M.S. degree from the Air Force Institute of Technology (AFIT), Wright-Patterson AFB, OH, USA, in 2005, and the Ph.D. degree in computer science from AFIT in 2016. He is a retired Cybersecurity Chief from the U.S. Marine Corps. He is currently the Cyber-Physical Research and Development Program Manager with the Oak Ridge National Laboratory, Oak Ridge, TN, USA. His research interests include critical infrastructure protection, supervisory control and data acquisition (SCADA) systems, industrial control systems (ICS) security, and electromagnetic pulse (EMP) effects to SCADA/ICS. His professional activities include Technical Lead for SCADA/ICS Research at the Air Force Cyberspace Technical Center of Excellence, AFIT from 2008 to 2016, and is the Co-Chair for the Industrial Society of Automation (ISA), Work Group 4, Task Group 7 (Security of ICS Sensors), since 2016. His professional credentials are Certified Information Systems Security Professional (CISSP), Certified SCADA Security Architect (CSSA), Certified Scrum Master, Lean Six Sigma Green Belt, and holds an FCC extra class amateur radio license.



NICHOLAS A. PETERS (M'16-SM'17) received the B. A. degree summa cum laude in physics and mathematics minoring in computer science, from Hillsdale College, Hillsdale, Michigan, in 2000 and the M.S. and Ph.D. degrees in physics from The University of Illinois Urbana-Champaign, in 2002 and 2006, respectively. After a short postdoc at the University of Illinois, in 2006, he became a Senior Research Scientist in Telcordia Technologies' Applied Research Organization, which later became Applied Communication Sciences. In 2012, he was promoted to Senior Scientist at Applied Communication Sciences. In 2015, he joined Oak Ridge National Laboratory as Senior Research and Development Staff Member. In 2016, he was granted a joint appointment to the faculty of The Bredesen Center at the University of Tennessee Knoxville. In April 2017, he was appointed to lead ORNL's Quantum Communications Team. He is an Associate Editor for Optics Express, the author of more than 60 journal and conference papers, and has been issued 11 US patents.

Dr. Peters' is a Senior Member of The Optical Society (OSA) and a member of the American Physical Society. His awards and honors include a Thomas Alva Edison Patent Award in Emerging Technology, a National Intelligence Meritorious Unit Citation, invitation to the National Academy of Engineering's (NAE) 18th annual U.S. Frontiers of Engineering symposium, two Corporate CEO awards, an ORNL significant event award, an ORNL Technology Commercialization Award, and the University of Illinois Scott Anderson Outstanding Physics Department Teaching Assistant Award.



KALYAN S. PERUMALLA is a Distinguished Research Staff Member and Manager at the Oak Ridge National Laboratory (ORNL). Dr. Perumalla founded and currently leads the Discrete Computing Systems Group in the Computer Science and Mathematics Division at ORNL. He is as an Adjunct Professor in the School of Computational Sciences and Engineering at the Georgia Institute of Technology (Georgia Tech) and a Joint Professor in the Department of Industrial and Systems Engineering at the University of Tennessee, Knoxville.

He served as a Fellow of the Institute of Advanced Study at the Durham University, UK, 2015, and on the National Academy of Sciences' Technical Advisory Boards on Computational Sciences and Information Science at the U.S. Army Research Laboratory, 2015-2017. Dr. Perumalla is among the first recipients of the U.S. Department of Energy Early Career Award in Advanced Scientific Computing Research. Over the past 20 years, he has served as a principal investigator or co-principal investigator on several research projects sponsored by federal agencies and industry. Dr. Perumalla earned his PhD in Computer Science from Georgia Tech in 1999. He co-authored a book, three book chapters, and over 100 articles in peer-reviewed conferences and journals. Five of his co-authored papers received the best paper awards, in 1999, 2002, 2005, 2008, and 2014.

Dr. Perumalla is a Certified Scrum Master and Certified Scrum Product Owner.



STACY J. PROWELL serves as the Chief Cyber Security Research Scientist at Oak Ridge National Laboratory and is the Program Manager for the lab's Cybersecurity for Energy Delivery Systems program. Dr. Prowell's research focuses on exploiting physical sensors and "side channel" information to detect and prevent intrusion, on deep semantic analysis of compiled software, and on the security of critical infrastructure. Dr. Prowell's work on a system for analysis of compiled software led to the Hyperion system, which received a 2015 R&D 100 award and two awards for technology transfer.

Previously, Dr. Prowell worked in the CERT Program of the Software Engineering Institute on automating the analysis of malware. Dr. Prowell is an IEEE Distinguished Lecturer for the Transportation Electrification Community, an Associate Professor of Electrical Engineering and Computer Science at the University of Tennessee, an Associate Professor of Computer Science at Tennessee Technological University, a member of Sigma Xi, and a Certified Scrum Master.



VARISARA TANSAKUL received the B.S. and M.S. degrees in industrial engineering from The University of Tennessee, Knoxville, Tennessee, USA, in 2016 and 2018, respectively. She previously worked as a graduate teaching assistant and graduate research assistant at the University of Tennessee, and is currently an ASTRO intern at Oak Ridge National Laboratory in the Discrete Computing Systems Group. Her research focuses on machine learning and data analytics.



CURTIS R. TAYLOR (M'14) is Cyber Security Research Scientist at Oak Ridge National Laboratory in the Cyber Defense group. His research interests include networking, network security, and systems security. He received his BS in computer science from the University of Tennessee - Knoxville and his MS and PhD in computer science from Worcester Polytechnic University. He is also a research associate at Worcester Polytechnic Institute and a co-founder of ContextSure

Networks.

Dr. Taylor is a graduate of the NSF CyberCorps Scholarship for Service (SFS) program and is a Certified Scrum Master (CSM).



BAILU XIAO (S'09-M'15) received the B.S. and M.S. degrees in electrical engineering from Huazhong University of Science and Technology, Wuhan, China, in 2006 and 2008, respectively, and the Ph.D. degree in electrical engineering from the University of Tennessee, Knoxville, TN, USA, in 2014. She joined Oak Ridge National Laboratory (ORNL) in 2014 as a Postdoctoral Research Associate, and she is currently a Research and Development Staff at ORNL. Her current areas

of interest include multilevel converters, power converters for distributed energy resources, and microgrid control.



SRIKANTH B. YOGINATH is a Research Staff Member at Oak Ridge National Laboratory in the Discrete Computing Systems Group. He received his Bachelor of Engineering degree in Telecommunications from The Bangalore University, India, in 1999. In the year 2003, he received his MS in Computer Science from Illinois Institute of Technology, Chicago and joined Oak Ridge National Laboratory as a Post Master's student.

While working at ORNL, he joined the PhD program at Georgia Institute of Technology (Georgia Tech), Atlanta, in the year 2008. He was awarded PhD in Computational Sciences and Engineering from Georgia Tech, in 2014. His research interests include parallel and distributed computing, virtual machine-based network simulations, parallel discrete event simulations, large-scale simulation frameworks and, scaling machine learning and deep learning algorithms.

...