A MODEL FOR INTELLIGENT PROTECTION OF CRITICAL COMPUTER SYSTEMS

FERNANDO GONZALEZ[1], MAREK RUSINKIEWICZ[2] AND JANUSZ ZALEWSKI[1]

[1] *Dept. of Software Engineering*
Florida Gulf Coast University
Ft. Myers, FL 33965, USA

[2] *Ying Wu College of Computing*
New Jersey Institute of Technology
Newark, NJ 07102, USA

(received: 6 June 2021; revised: 20 August 2021;
accepted: 22 August 2021; published online: 30 November 2021)

Abstract: We propose a unified model for the enforcement of safety and security of cyber-physical systems in critical applications. We argue that the need for resilience of a critical system requires simultaneous protection from hazards (safety) and from unauthorized access (security). We review how the critical system properties are handled and present a framework for their modeling. Then we present a model for the enforcement of critical system properties through situational awareness, including threat monitoring, data analysis and state prediction for decision making. We conclude by presenting a case study of a power grid simulation and advocate the ability to move from today's reactive approaches to proactive ones that aim at avoiding system failures.

Keywords: Computer Safety, Computer Security, Critical Systems, Intelligent Systems, Safety Security Modeling

DOI: https://doi.org/10.17466/tq2021/25.3/a

1. Introduction

Critical computer systems may cause significant damage to the environment if they fail. They form the nation's infrastructure for transportation, energy distribution, and telecommunication, so the primary concern in their design and use has been safety. Since most of these systems incorporate networked computers, informational resources, and physical components, such as sensors and actuators, they fall into the category of *cyberphysical systems*.

The increasing reliance on cyberphysical systems has been paralleled by a corresponding increase in the variety, frequency and impact of attacks from a

range of assailants. Both commercial companies and government agencies face continuous and increasingly more sophisticated cyber-attacks ranging from denial of service and data exfiltration to sophisticated worms and logic bombs. The targets include not only computer information systems, but also the network communication infrastructure and power grids. The main scope of cybersecurity is protection of networked infrastructure and information from unauthorized access and destruction. Thus, concerns about system safety and safety of the environment should be treated jointly with cybersecurity to the extent that they become indiscernible, as concerns about one imply the other.

Moreover, with the unprecedented expansion of the Internet, cyberphysical systems are quickly evolving towards the Internet of Things. In the Internet of Things (IoT) the emphasis is on exchange of data among entities that include sensors and actuators connected to a computer network (the Internet). Recently, a more general term, Internet of Everything (IoE), has been introduced to underline the fact that cyberphysical systems, in addition to Things and Data they exchange, include other important components, namely People and Processes.

Current estimates about the development of IoT technologies vary, but they all have one aspect in common, namely, that the enormous growth in the number of individually IP addressable devices on the Internet, the high rate of devices being added per week, the exponential increase of data created by devices vs. those created by people, as well as broadening the economic impact of IoT on all sorts of industries will continue [1]. It appears, however, that this fast pace of evolution in technology is not being met sufficiently quickly by the development of respective protection techniques required for safe and secure operation.

Since their role in the society and in the nation's well-being is essential, the cyberphysical systems must possess desired quality characteristics, in addition to proper functionality, and need to be strongly protected against the unauthorized use, to prevent the unwanted consequences of their potential failures. In particular, techniques for prevention and detection of attacks, as well as for recovery from attacks need to be developed. While our previous project was related to the analysis of safety and security in cyberphysical systems for attack prevention [2], this paper concerns the attack detection and recovery.

The rest of the paper is organized as follows. Section 2 presents a review of selected aspects of previous related work, Section 3 discusses a model of a cyberphysical system applicable to modeling the critical systems, and Section 4 outlines the organization and functionality of the safety security supervisor based on this model. A case study of a simulation of an electric power grid is presented in Section 5 and the paper ends with a conclusion.

2. Review of Related Work

Providing protection for all kinds of cyberphysical systems is a daunting task, so we focus here on one particular class of systems playing a crucial role in the nation's critical infrastructure, which is the electric power grid. Automation

for the power grid is normally provided by SCADA (Supervisory Control and Data Acquisition) technologies [3].

Therefore, to properly tackle the protection issues we need to combine the system safety concerns that were traditionally addressed in the context of SCADA, with the network and information security concerns addressed by the computer security approaches in Information Technology (IT). This is complicated by the fact that although the two approaches are slowly converging, there are still important differences between traditional IT and SCADA systems, some of which are listed below:

- The time horizon for deployment of IT systems is 3-5 years; for SCADA is more like 15-20 years
- Unlike many IT systems, the SCADA systems cannot be shut down for maintenance; they need to be up 24/7
- The consequences of a failure are bigger and more immediate in case of SCADA
- In many aspects, security of SCADA systems lags several years behind the best practices of information and networking security.

On the other hand, there have been multiple attempts to describe problems and methods for cybersecurity protection. In the latest paper [4], Peng et al. review cybersecurity issues in smart grid with respect to attack scenarios, detection and protection methods, as well as estimation and control strategies from communication and control viewpoints. In conclusion, they advocate for giving more serious consideration to the following issues:

- Security detection combined with advanced analysis methods
- Modeling the attacks with more practical conditions
- Distributed detection and estimation of attacks
- Resilient control strategies.

In the current paper, we focus in particular on the first two recommendations.

Another recent paper, by Wadhawan et al. [5], addresses the protection of smart grid against cyberphysical attacks. The authors use the function-based methodology to evaluate smart grid and compute the likelihood of the compromise of grid's cyber components to assess the risk for these components, applying the technique of reinforcement learning, making the case for using intelligent techniques.

In yet another paper, Otuoze et al. [6] analyze the security challenges as well as the objectives of a smart grid, discuss countermeasures for grid protection and suggest a framework for achieving a secured grid. In their framework, the essential idea relies on taking a holistic approach, based on the sources of threats, for identification and clearance of detected threats.

Though a multitude of other articles have been published in the last decade on the subject of protecting and providing cybersecurity for the critical systems,

especially the power grid, for example [7, 8], the approaches described are more or less similar to those outlined above [4–6]. In summary, they call to address the following most important issues:

- Take a holistic approach to tackle the problem in a comprehensive manner
- Apply advanced analysis methods combined with modeling
- Use intelligent techniques to assist in reasoning about protection and cybersecurity.

The next section outlines a model of a cyberphysical system that can be used to analyze safety and cybersecurity in this context and provide protection for the critical computer systems.

3. Model of a Cyberphysical System

To address the quality characteristics of critical systems we need to build a conceptual model, which can be supported and developed further by adding respective assessment techniques and tools. In [9] we proposed a model of a cyberphysical system that evolved from previous research on real-time architectures. It consists of the core controller node providing required system functionality and four interfaces connecting it to the external process (plant), human operator, database and a network, which form an attack surface shown in Figure 1. This model addresses the basic questions that need to be asked: "What" is protected and against "Whom"? The "What" is reflected in the attack surface and the "Whom" is represented by the view of an attacker.

Furthermore, the model is based on the concept of observable behavior – the measurements that can be externally monitored and used for feedback. At the center of the model is the Protected Domain that consists of the cyberphysical system being protected, composed of the plant and the controller, which issues control commands and receives measured plant parameters coming from different sensor values. (In many large applications there is no centralized controller and the control functions are distributed among system nodes.) The controller compares these measurements with prescriptive and/or historical data from the database to ensure that the entire system is operating normally and is in a safe and secure state.

In addition to defining an attack surface and adding a view of an adversary, Figure 1 can also serve as a model that illustrates how to combine the safety and security concerns. This concept can be extended 10 to encompass the symmetry between safety understood as "freedom from hazards" that may activate system faults, and security understood as "freedom from threats" that may exploit system vulnerabilities. This leads, in turn, to the redefinition of both safety and security to expose their dual nature, that is, making a distinction between a system property and a system state.

On one side, a system can be claimed to reach (or not) a secure state and/or a safe state, in which case it ensures security and/or safety, respectively. This can be modeled in ways similar to describing the program state: by a collection of

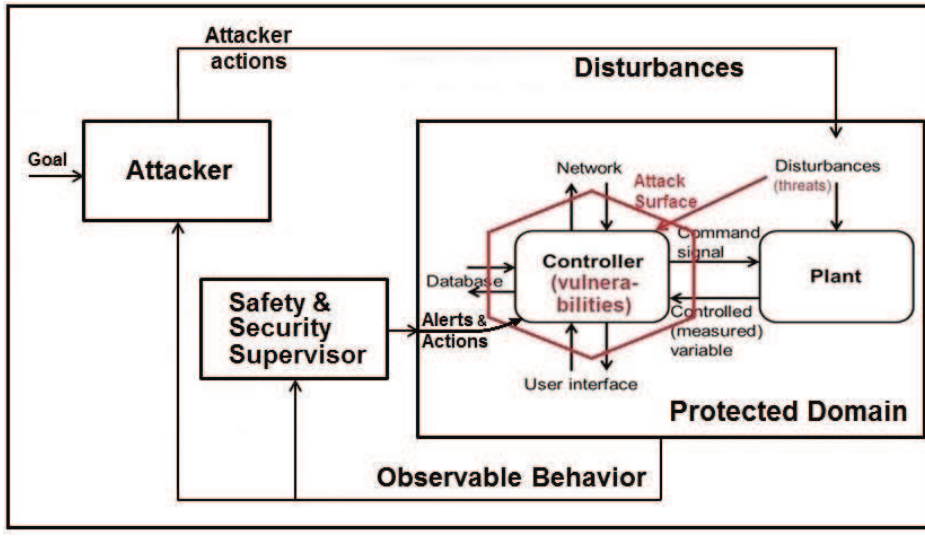

Figure 1. Model of a cyberphysical system and attacker

values of all variables at any given instant, and represented by tools such as state transition diagrams, Petri nets, etc. On the other hand, one can ask to what extent the system as a whole can exhibit any of these properties, that is, how secure is the system or how safe it is, and consequently how resilient it is? This view leads to the need of measuring each of these respective properties, or at least evaluating them quantitatively. Given that the current level of mathematical theory to evaluate system state and properties is not adequate, one is left with two options: experimental evaluations, that is, measurements, or using computational tools (simulations). This paper puts emphasis on the use of the latter, which has been previously advocated for, e.g., in [10].

4. Safety and Security Supervisor

The role of the Safety and Security Supervisor shown in Figure 1 is to monitor continuously the behavior of the Protected Domain, which consists of states of the Plant as well as the Attack Surface (formed by all the interfaces of the controller), to conduct the analysis and detect suspicious events that can be symptoms of potential attacks. It is consistent with other known approaches to cybersecurity of control systems, for example [11].

The attacker actions affecting the Protected Domain may result in externally visible observables that are their direct or indirect effect. This observable behavior is monitored by the Safety and Security Supervisor to compute and predict system states, in order to protect it. The monitoring of the observable behavior and identification of potential threats provides the basis for Situational Awareness and serves as an input to the decision making process applying the protection mechanisms/algorithms to compute relevant security actions, which is illustrated in Figure 2. We should note that (some elements of) the system's ob-

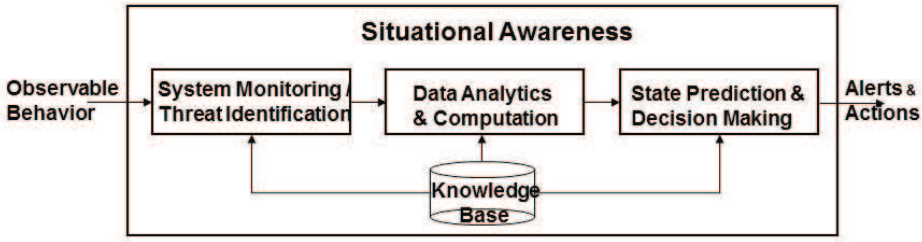

Figure 2. Structure of safety and security supervisor

servable behavior may be also visible to the Attacker and can be used to adjust the attack profile.

As illustrated in Figure 2, the holistic view of cybersecurity to provide Situational Awareness takes into account the following aspects:

- System Monitoring and Threat Identification
- Data Analytics and Computation
- State Prediction and Decision Making

all making use of the Knowledge Base that stores and explores previous states of interest. Roles of each specific component are explained in the following subsections.

4.1. System Monitoring and Threat Detection

Threat detection involves processing large amounts of data under near-real time constraints. The data may come from heterogeneous information sources including sensors, text messages, system logs, still images, and videos. The very useful paradigm is that of perpetual queries, in which once defined query is continuously evaluated over streaming data.

The queries may range from simple, e.g., "value of a variable exceeded the threshold" to quite complex "gradient of the variable change exceeded the threshold twice during a specified time window". The queries may be evaluated simultaneously at different spatial and temporal granularities. Finally, the detection system must be capable of incorporating ad hoc, opportunistic information sources [12].

4.2. Data Analytics and Computations

Continuous monitoring of streaming data and detection of patterns that may indicate possible violations of safety and security generate very large amounts of data. The raw data may be useful in generating alerts, but only statistical data analysis can detect more subtle phenomena needed for decision making. One of very useful techniques is multi-scale data analysis.

A cybersecurity system must be able to define composite events of interest using temporal event integration and then to continuously monitor and detect events of interest, including:

- Simple events (specific tasks started, data element updated, message received)
- Temporal events (clock or timer related)
- Application-specific events (image change, face recognition, trip-wire)
- Composite events (patterns consisting of primitive events combined using logical and algebraic operators such as conjunction and disjunction or temporal operators: before, after, between)
- Data Mining events (e.g., trends detected).

The other important capability of a situational awareness system based on event detection is the ability to produce and deliver targeted situational alerts. This is based on recursively composing events and delivering alerts to systems or persons who expressed interest in an event – either by explicitly subscribing to it or because of the role they play in the organization. Finally, the awareness system should be able to respond to queries over information stored in the awareness repository, e.g., for explanation.

The data analysis techniques may involve clustering, graph mining, and deep-machine learning to identify potential threats and to mitigate them.

4.3. State Prediction and Decision Making

Then standard analytical techniques, such as anomaly and outlier detection or nearest neighbor's determination could be applied. For the evaluation of what-if scenarios machine learning approaches, such as neural nets can be applied to use the historical data for building models that can be used to predict events based on the inputs patterns.

One of the important functions of the situational awareness system is the ability to distinguish between attacks and failures. We understand that failures cause the symptoms, but symptoms are what we can observe. Hence one must proceed from symptoms to failures using unsound rules and abductive reasoning. As an example, let us consider Communication Network Security.

Network elements (hosts, routers, wireless radio links) can fail in unpredictable ways. Usually one cannot distinguish between the primary fault (root cause) and secondary faults, since the symptoms of fault and performance problems overlap and interfere with each other. We need to reason over observations to correlate the observations about state changes and thereby diagnose the root cause of a problem. Once the root cause is diagnosed, one can reconfigure/repair the network appropriately.

Because the network attacks are orchestrated by a human adversary, their detection is complicated by the fact that the adversaries may change the attack signatures to look different each time. Attacks may be disguised to look like normal operations and mimic normal communication patterns or spread out across time. Since the attacks may look very similar to performance problems their detection may require reasoning based on history and/or the temporal dimension. To do this one may need to selectively maintain history of "normal" operations. Then one can detect sequences of events that may signify the outset of an attack.

An important consideration is that insider attacks may often look normal up to a point. Therefore, one has to balance the impediments that frequent false positives alert may impose on the normal operation of the system versus the desired degree of security. All these assumptions have to be incorporated in the state prediction algorithm.

4.4. Knowledge Base

The role of the Knowledge Base in Figure 2 is to serve as an autonomous unit of a primary storage of previous behaviors that represent both normal and abnormal (suspicious) states. As an example of its typical function let us consider the results of network message delay measurements over a period of time, presented in Figures 3 and 4.

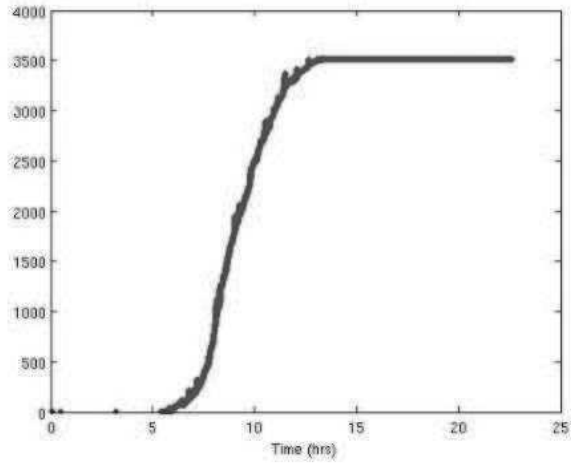

Figure 3. Effects of network message (ms) delay vs. time

Figure 3 shows a phase transition that occurred between 10th and 15th hour of observation, when the network moved from a normal state (no delays) to a fully congested state (no throughput).

Figure 4 presents the same measurements on a much finer scale. One can see that there were actually two spike events (at 10.5 and 10.6 hrs), which may have caused the failure. Therefore, a network fault analysis system needs to calculate multi-scale profiles of streaming data, so that clusters at different scales can be detected and characterized.

The Knowledge Base ensures that the records of known behavior are stored and their format ensures immediate access on per need basis. It also autonomously detects discrepancies in system states, which require attention.

5. Case Study

To illustrate the application of the concept of a Supervisor, this section presents a case study of an electric power grid simulation, which does the Data

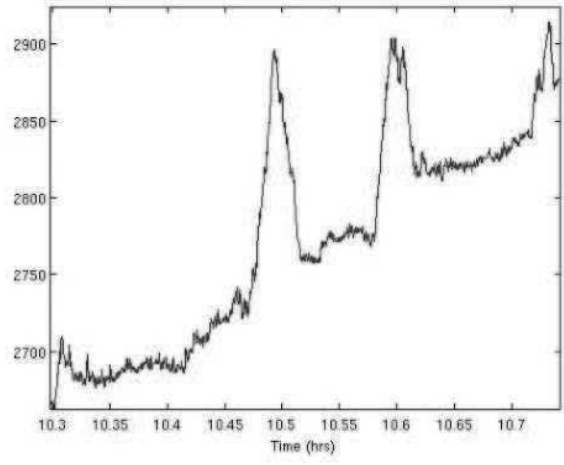

Figure 4. Message delays (ms) in condensed time

Analytics part shown in Figure 2 to determine the effects of abnormal behavior, which can then be used to assess and predict the impact of potential threats on system state and develop respective actions in real time.

5.1. System Monitoring and Threat Detection

The simulation tool developed previously[13] is given the model of the power system, the list of parameters it has in its control domain and a heuristic function that tells how the system is performing. The intelligence is based on Monte Carlo simulations where for each possible control strategies the tool sets the parameter to a corresponding control law, runs several nondeterministic simulations and collects statistical data. Then the heuristic function returns a value indicating how the system is operating under the influence of the control law in question. The heuristic function may include financial costs associated with events such as shedding loads, generating power, and importing power.

After evaluating all the laws the best control law is the one that the simulations estimate will make the system have the best performance according to the formulas in the heuristic function. This is the method that history has shown to be most effective in optimizing manufacturing plants. It is important to note that this controller finds values to parameters that tend to make the power system perform better. The simulation model itself is considered to be input to the system as is the heuristic function.

The simulation of each individual subsystem is performed using the commonly used discrete-event simulation method which of course uses only a single model. In order to make the system scalable to arbitrarily large systems, each sub power system is modeled individually with its own unique control model. The simulation is designed to accommodate a collection of individual models and treat them as a single software model all communicating over a simulated network. This

simulated communications network is what ties together all the individual models. It allows simulations of any subset of models to be executed without additional modeling effort. It also allows the fidelity of the simulation model to be adjusted by selecting the models to include. For example, reducing fidelity by omitting models representing low level details is more appropriate for making fast high-level decisions.

An executive function manages the simulated time, event list and the list of resources. As it pulls of the next event in the chronologically ordered event list it calls its simulation model (the single compound model) to execute the event. The simulated time is incremented to the time of the next event in a single discrete jump. All of these components, the event list, the list of resources, current time and the model, are encapsulated into a single object, called the simulation object.

The intelligent controller involves the collection of simulation objects operating as controllers using actual Internet communication and running in real time. The intelligence of the controller involves running simulations of the distributed system or subsystem concurrently with the controller. The intelligent control system has a recursive property that allows it to form a control hierarchy. Each controller in the hierarchy has the capability to execute simulations to optimize its control domain. In the power system case study, each of the three sub-systems optimizes its own system and the supervisory controller overseeing the complete system also performs optimization to improve system wide operations. This allows a failure in a subsystem to be noticed by a supervisory controller that can in turn request support from another subsystem in order to achieve a system wide recovery effort.

Research shows that the 2003 North-East Blackout [14], that resulted in over 10 Billion loss and affected 10's of millions customers could have been prevented with a fast system wide recovery effort. The plant that experienced the initial failure was incapable of recovering and relied on support from other plants requiring a system wide coordinated effort. The failure only took 3 minutes to propagate from its source plant to 21 additional power plants requiring real-time intelligent control.

5.2. Experimental Results

Sample runs of the simulator for a specific example of an electric power grid are shown in this section. The initial condition of the grid, Fig. 5, known to be under unknown threat, is as follows: Generator G4-2, bottom middle, is shown at capacity but that in itself is not a problem. Generators G5-1, G5-2, and G5-3, top left, are all on, however G5-1 is less efficient and is meant to only be used as a backup in high demand situations. In Figures 5 – 9, the circles represent generators, triangles represent loads, and the lines represent power transmission lines. The color green means the unit is operating within its limits, yellow means the operations is close to it limit and red means the unit is operating beyond its limit and must be remedied to avoid failure. The number below each unit is the operating current either output for generators or input for loads. The current

capacity of each unit is enclosed in brackets, either maximum current production for generators or maximum current consumption for loads. The needle in the dials represent the percent capacity the unit is operating at and its color reflect this percent as with the colors of the units themselves. When the simulation runs, the solid black circles move along the transmission lines with its speed proportional to the flow of power. With these notations, including the black solid circles moving along the transmission limes the dials, the user has a graphical display showing the instantaneous state of the complete system. Whenever the optimizer makes a change to a generator or load, it encloses the changed unit in a black circle to show the user where it made a change.

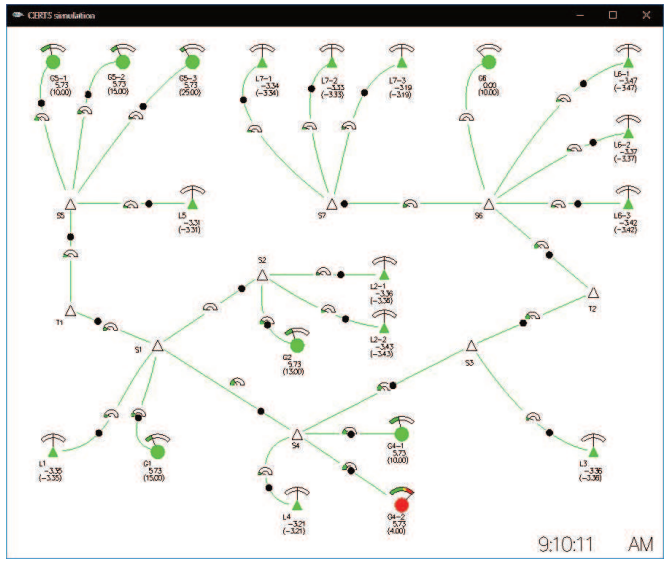

Figure 5. Initial state of the grid

When the optimizer was executed it determined that node G5-1 can be turned off (circled in Figure 6), which in security terms corresponds to reducing attack surface.

Next, we induced a failure and observed the system's response. The load L7-1, top center (circled) in Figure 7, was deliberately shed, but when the optimizer detected no threats and determined it's safe to reconnect load L7-1.

The next possible threat (event) is the disconnection of a transmission line, as shown in Figure 8 (bottom left, dotted line). The failed lines are modeled as physical failures needing repair and have negative effect on energy delivery. Thus, the optimizer determined the set of loads that must be shed in order to minimize the total cost of the failure. The respective loads, L6-1, top right, L3, bottom right, and L4, bottom center, were all shed by the controller in order to prevent the failure from cascading and causing additional damage.

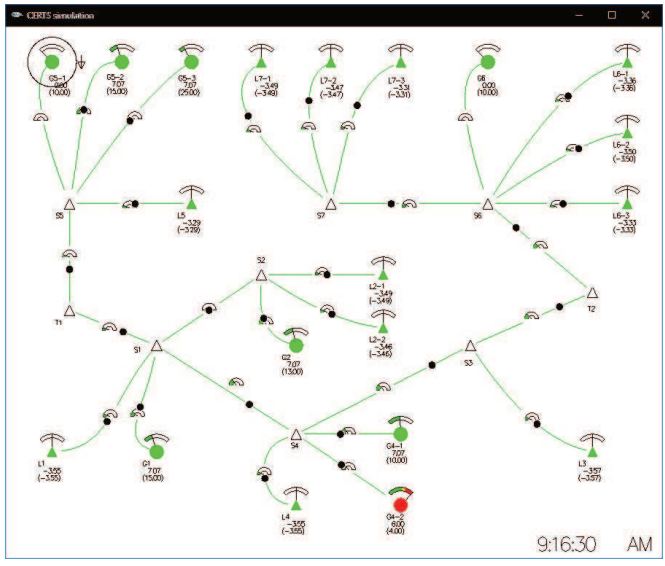

Figure 6. Optimizing the state of the grid

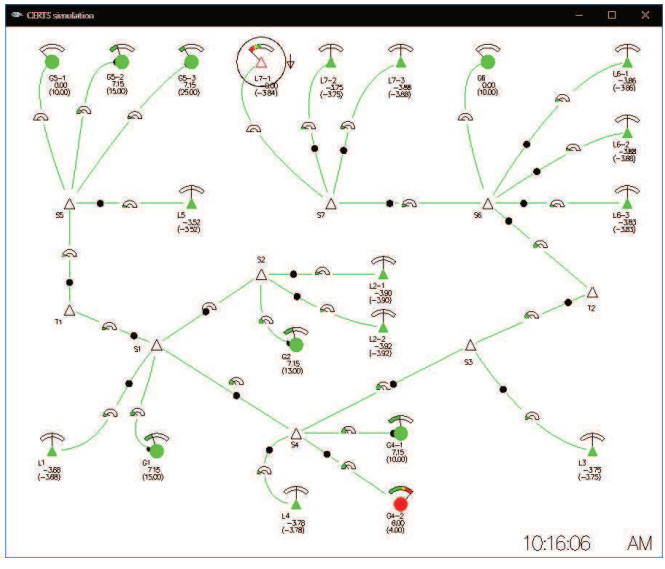

Figure 7. Reinstating the state of the grid

This exercise not only shows that the system can handle the failure by preventing it from cascading while minimizing total cost but also shows us the impact associated with an attack on that particular transmission line. This tool is also used to measure the vulnerability of each component in the system.

Unexpected turning off any component in the power grid, whether it's a generator, transmission line, transformer, load or anything else, may cause significant problems to the grid's operation and its customers, so it is considered

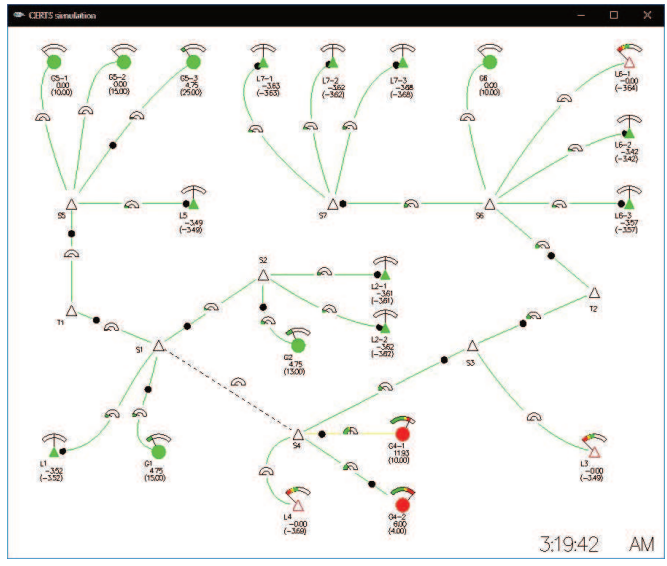

Figure 8. Effect of a failed line on the state of the grid

a major security and safety concern. However, the optimizer determined that generator G6, top center-right (circled) in Figure 9, is not being used. With the decision to turn on generator G6, optimized loads L3, L4, and L6-1 have been reconnected (all shown circled in this figure).

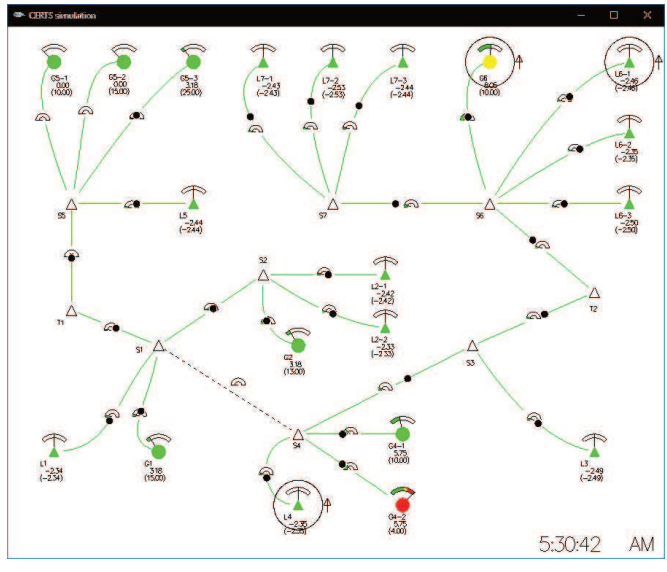

Figure 9. Reinstating the state of the grid

Overall, the simulator aids in making an intelligent determination whether or not a security attack is imminent or may have occurred. It creates a mapping

of the power system's infrastructure to its vulnerability to an attack. The vulnerability of each component in the model is assessed by inducing a failure in that component (imitating an attack) and recording the estimated optimized recovery and its associated cost.

6. Conclusion

In this paper we presented a cybersecurity model for critical computer systems based on careful examination of what needs to be protected and how this can be accomplished. The model forms a unified framework to represent both the safety and security aspects of cyberphysical systems. We then discussed the three main components of the cybersecurity, providing the situational awareness, namely system monitoring, data analytics, and state prediction. The proposed model supports cognitive approaches to safety and security based on generating hypotheses, evidence-based reasoning, and near real-time decision recommendations. It is applicable to counteracting Distributed Denial of Service (DDS) attacks, as well as to the detection of risky user behaviors, potential malware, and data exfiltration attempts. It can assist system analysts through interactive vulnerability analysis, risk assessment, and possible attribution. Using this model, one can consider alternative outcomes and explore what-if scenarios, leading to better decision making. This, in turn, would permit improvement in today's reactive approaches to address critical system properties more proactively, tomorrow.

References

[1] Ekholm B 2019 *President CEO: Ericsson Mobility Report. Special Edition*, *World Economic Forum*

[2] Subramanian N and Zalewski J 2016 *Quantitative Assessment of Safety and Security of System Architectures for Cyberphysical Systems Using the NFR Approach*, *IEEE Systems Journal* **10** (2) 397

[3] Thomas M S and McDonald J D 2015 *Power System SCADA and Smart Grids*, *CRC Press, Florida*

[4] Peng C. et al. 2019 *A Comprehensive Analysis of Smart Grid Systems against Cyber-Physical Attacks*, *IEEE Trans. Systems, Man and Cybernetics*

[5] Wadhawan Y, AlMajali A and Neuman C 2018 *A Survey on Security Communication and Control for Smart Grids under Malicious Cyber Attacks*, *Electronics (Switzerland)* **7** (249)

[6] Otuoze A O, Mustafa M W and Larik R M 2018 *Smart Grids Security Challenges: Classification by Sources of Threats*, *J. of Electrical Systems Information Techn.* **5** 468

[7] Wang W and Lu Z 2013 *Cyber Security in the Smart Grid: Survey and Challenges*, *Computer Networks* **57** 1344

[8] Sridhar S, Hahn A and Govindarasu M 2012 *Cyber-Physical System Security for the Electric Power Grid*, *Proc. of the IEEE* **100** (1) 210

[9] Sanz R and Zalewski J 2003 *Pattern Based Control Systems European Control Conference (ECC 2013)Engineering*, *IEEE Control Systems* **23** (3) 43

[10] McDonald M J et al. 2010 *Modeling and Simulation for Cyber-Physical System Security Research, Development and Applications. Report SAND2010-0568*, Sandia National Laboratories

[11] Sandberg H 2013 *Secure Control and Applications in Power Systems. Tutorial*, European Control Conference (ECC 2013)

[12] Bayardo R J Jr et al. 1997 *InfoSleuth: Agent-Based Semantic Integration of Information in Open and Dynamic Environments (in:) Proceedings SIGMOD'97*, ACM International Conference on Management of Data 195

[13] Gonzalez F G 2013 *An Intelligent Controller for the Smart Grid*, *Procedia Computer Science* **16** 776

[14] *Public Safety and Emergency Preparedness Canada. Ontario–U.S. Power Outage – Impacts on Critical Infrastructure* 2006, Incident Analysis Report IA06-002

[15] Kornecki A and Zalewski J 2015 *Aviation Software: Safety and Security (in:) Wiley Encyclopedia of Electrical and Electronics Engineering*, John Wiley Sons, New York

Dr Fernando Gonzalez is an Associate Professor of Software Engineering at Florida Gulf Coast University in Fort Myers, Florida, USA. He earned his Ph.D. degree in Electrical Engineering from the University of Illinois in 1997. He earned his Master's degree in Electrical Engineering and his Bachelor's degree in Computer Science from Florida International University, Miami Florida, USA in 1992 and 1989. Dr. Gonzalez has worked at Texas AM International University in Laredo, Texas, the U.S. Department of Energy at Los Alamos National Laboratory in Los Alamos, New Mexico, and at the University of Central Florida in Orlando, Florida. Dr. Gonzalez research interest includes intelligent agents for the Semantic Web, the intelligent control of large scale autonomous systems, autonomous vehicles, discrete-event modeling and simulation and human signature verification

Marek Rusinkiewicz is a computer scientist, an educator, and a former research executive. He recently retired from the position of Dean of the College of Computing Sciences at New Jersey Institute of Technology. He also held management positions as a Senior Group Vice President and the General Manager of Applied Research Laboratories at Telcordia Technologies (formerly Bell Communication Research). Before joining Telcordia, Rusinkiewicz was the Vice President for Information Technology Research at the Microelectronics and Computer Technology Corporation (MCC) in Austin, Texas. Rusinkiewicz has held academic positions at the University of Glasgow, the University of Michigan, and the University of Houston, where he was a Professor of Computer Science until 1999. His research interests include heterogeneous database systems, distributed computing, workflow management, and agent-based systems. He has consulted extensively for numerous industry and government organizations in the USA, Japan, Taiwan and Europe.

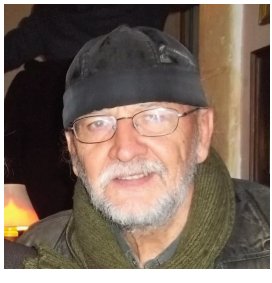

Janusz Zalewski is a retired professor of Computer Science and Software Engineering at Florida Gulf Coast University, in Ft. Myers, Florida, USA. He obtained his Ph.D. in Computer Science and Engineering in 1979 from the Faculty of Electrical Engineering at Warsaw University of Technology. He previously held academic positions at Embry-Riddle Aeronautical University and University of Central Florida. In the past, he worked on projects for the Superconducting Super Collider, Lawrence Livermore National Laboratory, FAA, and United States Air Force Academy, as well as consulted for a number of private companies, including Lockheed Martin, Har-

ris, Boeing, and others. He also had fellowships at NASA and Air Force Research Labs. His research interests include real-time embedded and cyberphysical systems, safety and security of complex computer systems and networks, and software engineering education. He is a member of the IFAC TC 3.1 on Computers for Control.