

SOCCA: A Security-Oriented Cyber-Physical Contingency Analysis in Power Infrastructures

Saman Zonouz *Member, IEEE*, Charles M. Davis *Member, IEEE*, Katherine R. Davis *Member, IEEE*, Robin Berthier, Rakesh B. Bobba *Member, IEEE*, William H. Sanders *Fellow, IEEE*

Abstract—Contingency analysis is a critical activity in the context of the power infrastructure because it provides a guide for resiliency and enables the grid to continue operating even in the case of failure. In this paper, we augment this concept by introducing *SOCCA*, a cyber-physical security evaluation technique to plan not only for accidental contingencies but also for malicious compromises. *SOCCA* presents a new unified formalism to model the cyber-physical system including inter-connections among cyber and physical components.

The cyber-physical contingency ranking technique employed by *SOCCA* assesses the potential impacts of events. Contingencies are ranked according to their impact as well as attack complexity. The results are valuable in both cyber and physical domains. From a physical perspective, *SOCCA* scores power system contingencies based on cyber network configuration, whereas from a cyber perspective, control network vulnerabilities are ranked according to the underlying power system topology.

Index Terms—Cyber-physical systems, security, contingency analysis, situational awareness, state estimation.

I. INTRODUCTION

State estimation and contingency analysis are two of the most fundamental tools for monitoring the power system. State estimation is the process of fitting data from sensors in the field to a system model and determining an estimate of the power system state [1]. State estimation engines use techniques such as the weighted least squares (WLS) algorithm to determine the state of the system based on the measurements [2]. Once the state estimator program determines a system estimate, the estimate is used to run a series of “what if” scenarios referred to as contingency analysis. Contingency analysis performs a series of power flow studies with various pieces of equipment outaged in the model, allowing operators to predict the state of the system for such an event [3].

By its nature, state estimation depends on the communication infrastructure, commonly called the SCADA (supervisory control and data acquisition) system. These systems are currently undergoing many upgrades as part of the smart grid initiative. This initiative does not only affect the SCADA system but also brings the telecommunication revolution to the entire energy delivery infrastructure, from control centers to generation, transmission and distribution substations, and even to customer homes. A direct consequence has been a significant increase in the number of inter-connected cyber

components. Manufacturers commonly add ethernet or radio communication modules to controllers, relays, and sensors, along with information report and configuration functionalities such as embedded web servers.

The increased cyber connectivity of the infrastructure and the interdependency of cyber and physical components introduces a greater level of complexity, and securing power system operations against malicious compromise becomes more challenging. Indeed, contingency analysis in the highly interconnected grid should be expanded to include incidents of intentional nature such as cyber attacks.

While the problem of detecting and mitigating cyber intrusions has been extensively studied over the past two decades in the context of traditional IT systems, the requirements and constraints of the smart grid environment in terms of security are different and usually more stringent. For example, power grid components commonly have timing requirements that prevent traditional security solutions from being deployed. The dependencies due to cyber-physical interactions in the grid are not yet well understood. Recently, there have been several attempts to model and analyze the cyber-physical threats in an offline manner [4]–[7]. Zonouz et al. [8] proposed an online framework that fuses uncertain information from distributed power system meters and cyber-side intrusion detectors to detect malicious activities within the cyber-physical system. However, to the best of our knowledge, there has been no efficient online solution proposed for contingency analysis that considers cyber adversary induced physical contingencies.

This paper introduces a *cyber-physical contingency analysis* framework for analyzing the physical impacts resulting from compromise in the cyber network. In particular, we present *Security-Oriented Cyber-Physical Contingency Analysis* framework or *SOCCA*, which takes into account cyber- and power-side network topologies, malicious cyber asset compromises, and power component outages. During an offline process, *SOCCA* analyzes the cyber network topology and the network firewall rules to automatically generate a network connectivity map as a directed graph encoding accessibility among different hosts (*e.g.*, computers and other cyber components). Then, *SOCCA* uses the connectivity map to create a partial Markovian state-based model of the power-grid in an online manner. At any time instance, *SOCCA* estimates the current security state using the generated model and the set of cyber intrusion detection sensor alerts. Using a new cyber-physical security index and cyber-physical state notion, *SOCCA* measures the criticality level of each system state and produces a ranked list of potential cyber and/or physical contingencies that need to be addressed.

Saman Zonouz *s.zonouz@miami.edu* is an Assistant Professor in the Electrical and Computer Engineering Department at the University of Miami. Matt Davis and Kate Davis *{matt,kate}@powerworld.com* are with Power-World Corporation. Robin Berthier, Rakesh Bobba, and William H. Sanders *{rgb,rbobba,whs}@illinois.edu*, are Research Scientist, Assistant Research Professor, and Donald Biggar Willett Professor, respectively, in the Electrical and Computer Engineering Department at the University of Illinois.

The proposed framework does not require any additional measurements than those already present in modern control centers and does not impose additional communication requirements. Specifically, cyber intrusion detection systems are already in use in modern control networks, and the sensor alerts from those systems are typically available from a network Security Event Manager (SEM) used to aggregate and manage security alerts. Similarly, the needed power system state information is available from the Energy Management System (EMS).

The contributions of this paper are threefold:

- *Cyber-physical system formalism and automated model generation:* We propose a new formulation of cyber-physical failures and compromises and an algorithm to automatically generate the corresponding models.
- *Cyber-physical security index for the power grid infrastructures:* We present a novel and scalable security index for power grid contingency screening.
- *Cyber-aware contingency analysis:* We present a framework for power grid contingency screening that uses the security index to determine impact and criticality for each state and then ranks contingencies that could be caused by cyber attacks.

SOCCA is not designed to replace the traditional power system contingency analysis solutions, which analyze accidental failures that could occur at any part of the power system due to natural causes. Instead, SOCCA presents a complementary framework that concentrates on potential contingencies due to remote malicious attacks.

This paper is organized as follows. Section II reviews the high-level architecture of the SOCCA framework and how its components are logically interconnected. Section III discusses how SOCCA models various security incidents and their potential correlations. Section IV presents how the cyber-physical model probabilistically determines the current state given the real-time sensory information. Section V describes the way SOCCA explores and analyzes potential cyber-physical contingencies according to the current cyber-physical state estimate and provides a risk-based ranked list.

II. BACKGROUND

Power system modeling is used extensively in grid operations. The most common types of models used are steady state power flow models and sensitivities. These models are used to monitor the state of the system and predict the effects of changes.

Steady state power system modeling consists of enforcing the conservation of power. Given a set of power injections and withdrawals, the power flow finds the set of voltages and angles that satisfy power balance. The system state may be written as

$$\mathbf{x} = [\mathbf{V}, \boldsymbol{\theta}] \quad (1)$$

where \mathbf{V} is a vector of voltages, $\boldsymbol{\theta}$ is a vector of voltage angles. The vector of real power loads is \mathbf{P}_l and the vector of reactive power loads is \mathbf{Q}_l . Since generator outputs are controllable (within limits), they are collected separately in a vector of controls, \mathbf{u} .

The power flow problem can now be written as

$$\mathbf{f}(\mathbf{x}, \mathbf{u}) = \mathbf{0} \quad (2)$$

where $\mathbf{f}(\mathbf{x}, \mathbf{u}) = \mathbf{0}$ is a complex vector representing the injection at each node in the system. The function $\mathbf{f}(\mathbf{x}, \mathbf{u})$ represents the system model. It encapsulates factors like line impedances and system topology. Breaking $\mathbf{f}(\mathbf{x}, \mathbf{u})$ into real and reactive parts gives

$$f_i^p = -P_i^g + P_i^l + \sum_{k \in C} |V_i| |V_k| (G_{ik} \cos \theta_{ik} + B_{ik} \sin \theta_{ik}) \quad (3)$$

$$f_i^q = -Q_i^g + Q_i^l + \sum_{k \in C} |V_i| |V_k| (G_{ik} \sin \theta_{ik} - B_{ik} \cos \theta_{ik}) \quad (4)$$

These equations represent the nonlinear problem that is commonly called the power flow in power systems literature. The power flow is at the heart of most power systems analysis. It provides the basis for many tools and sensitivities that are used to predict the state of the system in the event of an outage.

Because the power flow is a non-linear problem, it is typically solved using the Newton-Raphson method, an iterative technique that requires multiple evaluations and factorizations of a large sparse matrix of sensitivities, the Jacobian matrix \mathbf{J} . The repeated factorizations can be a time consuming process, so more efficient approximate methods have been developed. These methods involve applying assumptions to the power flow equations to arrive at a simplified system model [9].

A commonly used simplification reduces the power flow to a linear problem, commonly called the DC power flow in the power system literature [3]. A constant matrix relates system angles and power injections. The DC power flow is the basis for many sensitivities. For example, power transfer distribution factors (PTDFs) estimate the changes in flow due to a transfer across the power system, and line outage distribution factors (LODFs) estimate the changes in flow on a line caused by the outage of another line [3]. PTDFs and LODFs are frequently used to predict the state of the system after an outage [10], [11]. There are also efficient extensions of LODFs to calculate changes in flow due to multiple outages [12].

III. CYBER-PHYSICAL SYSTEM SECURITY FORMALISM

In this section, we explain how we model power grid security attacks using a stochastic Markovian mechanism, and how models are automatically generated given the power grid cyber-physical topology.

A. Cyber-Physical State Notion

Before discussing how the current state of the power grid will be represented in SOCCA, a concise *cyber-physical state* notion needs to be defined. The security state, d , is the set of privileges that an adversary (or group of adversaries) has obtained out of the domain of possible privileges, D , which encompasses all privileges on all cyber hosts in the system as well as all connected physical power system devices.

The security state is used to indicate whether or not a specific event (contingency¹) has occurred in the grid infrastructure. In particular, we consider two types of contingencies.

¹In this paper, we treat both cyber- and power side incidents as *contingencies*.

First, there are cyber-side vulnerability exploitations, which are carried out by an attacker to obtain specific privileges and improve his or her control over the power network. Therefore, the information in a state denotes the attacker’s privileges in that state, e.g., root access on a mission-critical host system in power control room. Those privileges are used to determine what further malicious damage the attacker can cause in that state. Second, there are consequences, which are caused by the adversary after he or she obtains the required privileges. Specifically, consequences are defined to be incidents which affect the physical operation of the underlying power system. As a case in point, a transmission line outage, whether due to lightning or a remote malicious “open” command to a power relay, is a power-side consequence which results in a redistribution of power flow. The intent is to always operate the system such that the redistribution of power flow does not affect the end-user consumers.

B. Modeling Power Failures and Cyber Compromises

Generally, every power grid attack path consists of an escalating series of malicious actions by the adversary. The system’s initial state is (\emptyset) , in which no contingency has occurred and the attacker does not yet have privilege in the system. Starting from this state, the adversary aims to gain the set of privileges required to reach his or her goals, e.g., causing a power transmission line outage by opening the corresponding relay.

More specifically, every cyber-physical attack is in a finite set of security states S that cover all possible security conditions that the system could be in. The system is in one of the security states s at each time instant. From the system’s current state s , there are two types of transitions, corresponding to 1) adversarial vulnerability exploitations, and 2) malicious power contingencies. Formally, the attacker can choose and take an adversarial action $a \in A$ admissible in s , resulting in a state transition to s' .

To enumerate all possible attack scenarios, we model the adversarial actions as a discrete Markov decision process (MDP) [13]. A discrete Markovian decision process Γ is defined as a tuple $(S, A, F(\cdot), P, \gamma)$ where S is the security state space, assumed to be an arbitrary non-empty set endowed with the discrete topology. A is the set of actions which consists of adversarial vulnerability exploitations. For every $s \in S$, $A(s) \subset A$ is the set of admissible actions at state s . The measurable function $F : S \rightarrow \mathbb{R}$ is the susceptibility measure to attacks calculated for each state, and P is the transition probability function. That is, if the present state of the system is $s \in S$ and the attacker takes an action $a \in A(s)$, resulting in state transition to state s' with probability $P(s'|s, a)$, he or she obtains an immediate reward of $F(s')$. The discounting factor is γ which is normalized, i.e., $0 < \gamma < 1$. The discounting factor in control theory is a coefficient that models the fact that a future reward is worth less than the same amount of immediate reward.

C. Automatic MDP Generation

SOCCA automatically generates the MDP model for the power network given the control network topology, access

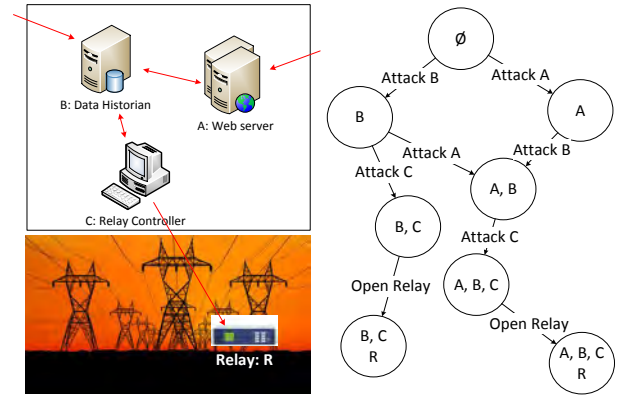


Fig. 1: A Power System Control Network and its Corresponding MDP

control policies, and cyber-physical interconnections within the power grid. The power network’s access control policies, such as firewall rulesets, are composed of rules about sources (IP/port addresses) that are either allowed or not allowed to reach a destination. SOCCA parses the rulesets and creates a binary network connectivity matrix that is a Cartesian product of host systems. The $[i, j]$ entry of the matrix takes on a true value if traffic from host h_i to host h_j is allowed, and a false value otherwise. The connectivity matrix always includes an Internet node representing a group of hosts outside of the network where attackers are assumed to initially reside.

The connectivity matrix incorporates *all* the possible accesses allowed by the global access policies. However, an attacker, with control over a particular host computer in the network, needs a vulnerability in one of the accessible hosts to exploit and improve his or her privileges. Therefore, SOCCA further refines the matrix to encode only the adversarial paths that can occur through vulnerability exploitations. In particular, SOCCA analyzes the power grid topology input to find the set of known system vulnerabilities in host systems. Given individual vulnerabilities, SOCCA determines their difficulty level through the web-based national vulnerability database [14] that uses the common vulnerability scoring system [15]. SOCCA converts the reported high, medium, low format to numerical 0.25, 0.5, 0.75 values that represent the attacker’s success rates. Consequently, the individual nonzero values in the connectivity matrix are updated accordingly, i.e., SOCCA replaces them either with zero if the corresponding hosts are not vulnerable or with relevant success rates otherwise.

Note that, while we use specific numerical values for transition probabilities in the rest of this work to demonstrate and evaluate SOCCA, a sensitivity analysis can be undertaken by varying the transition probabilities to identify structural weaknesses that are agnostic to transition probabilities. Similarly, while we focus on transition probabilities based on known vulnerabilities, the framework is flexible and zero-day attacks and vulnerabilities can be accounted for by using a small non-zero transition probability on all allowed connections.

To generate the MDP model, SOCCA traverses the connectivity matrix and concurrently updates the model following an incremental process. First, the MDP’s initial state (\emptyset) is created and the MDP generation starts with the network’s entry point (Internet) node in the connectivity matrix. Considering

the connectivity matrix as a weighted directed graph, a depth-first search (DFS) is run on the graph. While the search is recursively traversing the graph, it keeps track of the current state in the MDP, i.e., the set of privileges already gained through the path traversed so far. When the search meets a graph edge $[i, j]$ that crosses over privilege domains h_i to h_j , a state transition $a_a \in A$ in the MDP is created if the current state in the MDP does not include the privilege domain of the host to which the edge leads, i.e., h_j . The transition in the MDP is between the current state and the state that includes exactly the same privilege set as the current state plus the host h_j directed by the graph edge $[i, j]$. The transition created is labeled with the appropriate success rate (transition probability) from the connectivity matrix. The MDP's current state in the algorithm is then updated to the latter state, and the algorithm proceeds until no further updates to the MDP are possible according to the connectivity matrix. Figure 1 shows a simplified power network with vulnerabilities and its corresponding MDP model. Connectivity matrix elements are denoted as red arrows among pairs of network component.

It is noteworthy that SOCCA's MDP generation engine takes the known system vulnerabilities into consideration while generating the models. To address possible vulnerabilities targeted by zero-day attacks, the engine can incorporate those vulnerabilities to facilitate worst-case system contingency analysis. Additionally, the algorithm can handle generation of MDP models for systems where an adversary or group of adversaries starts from the Internet and can penetrate into the network from different vulnerable entry points. Several attackers would be modeled as a single "more powerful" attacker who can penetrate from any of those entry points.

Finally, SOCCA enhances the MDP model to also consider power contingencies using the power grid topology input that encodes the cyber-power interconnections, i.e., which power components are controlled by a particular host. In particular, considering every MDP state and the attacker's privileges in that state, SOCCA determines whether any malicious power contingency could occur and creates the required states and transitions accordingly. Those transitions will have the success rates of 1.0, because once the attacker gains the required privileges, he or she can cause a power contingency, by directly sending the corresponding command and without exploiting a vulnerability.

IV. CYBER-PHYSICAL SECURITY INDEX

SOCCA uses a new power grid security index to evaluate the security level of each MDP state. The proposed index takes into consideration the severity of the potential malicious physical consequences (i.e. percentage of line overload) and the difficulty to penetrate into the power network. In particular, using a defense-centric metric, SOCCA measures how susceptible the power system is to cyber attack induced contingencies, e.g., line outages, at every MDP state. As explained in the previous section, states capture the set of contingencies that can occur due to malicious actions. Then, SOCCA makes use of an adversary-driven metric to quantify how the attackers can obtain the required privileges to cause

those physical contingencies. Note that the physical contingencies considered here are steady-state contingencies. Although not discussed, SOCCA could potentially be modified to study transient contingencies as well.

Here, to measure the power system's susceptibility to cyber attack induced contingencies (e.g., line outages) for each state, SOCCA updates the admittance matrix according to the line outages encoded in that state, and solves the AC power flow equations using the iterative Newton-Raphson algorithm to calculate the line flows. Alternatively, an approximate DC model may be used, requiring only a linear solution, but potentially sacrificing detail [16]. SOCCA estimates the susceptibility degree using a modified version of the performance index [12] that assigns 0 to a state if there are no line flow violations and a positive value otherwise, computed using the following equation:

$$F(s) = \sum_{l \in L} [\max\{\frac{f_s(l)}{f^{MAX}(l)} - 1, 0\}]^2. \quad (5)$$

Here, L is the set of all lines, $f_s(l)$ denotes flow on line l in state s , and $f^{MAX}(l)$ denotes the maximum flow allowed on line l . In the event that a power flow fails to converge, a severe physical impact can be assumed which justifies setting $F(s)$ to a large number, the outage severity. The outage severity should be much larger than any of the line severity measures.

To calculate the overall cyber-physical security index for each state $I : s \rightarrow \mathbb{R}$, SOCCA uses the above power system performance index in the following dynamic programming equation:

$$I(s) = \max_{a \in A(s)} \{\gamma \cdot \sum_{s' \in S} P(s'|s, a) [\Delta F(s, s') + I(s')]\}, \quad (6)$$

where $I(\cdot)$ (as the MDP's value function) is formulated as a function of difficulty levels of various cyber attack paths $P(s'|s, a)$ and their physical impact $\Delta F(s, s') = F(s') - F(s)$. To solve Equation (6), SOCCA implements the value iteration algorithm. As formulated, proximity of the attacker to a physical component, existence of easy-to-traverse attack paths, and high final physical impacts of the attacks improve the fitness of a system state from the point of view of the adversary.

V. CONTINGENCY SCREENING

Algorithm 1 shows the pseudo-code for the algorithm that SOCCA employs to evaluate and rank individual cyber-physical contingencies. Briefly, SOCCA calculates an ordered list of single contingencies first, and investigates multiple contingencies ordered by adversarial preferences until either all the contingencies are analyzed or a predefined deadline has passed.

The cyber-physical contingency selection algorithm receives the generated MDP, the current state of the power grid and a hard deadline for the online analysis, and returns the contingency list (Algorithm 1). SOCCA starts off with initializing a temporal buffer list and a FIFO (first in, first out) queue (Lines 1, 2). Each state is assigned an initial color value (white, indicating the state has not yet been checked), as well as the corresponding physical performance and cyber-physical security indices as discussed in Section IV (Lines 3-6). At any

Algorithm 1: Cyber-Physical Contingency Selection

Input: MDP, current_state, deadline
Output: [ContingencyList]

```

1 List  $B \leftarrow \emptyset$ ;
2 Queue  $Q \leftarrow \emptyset$ ;
3 for  $s \in S$  do
4   Color[s]  $\leftarrow$  White;
5    $F(s) \leftarrow \sum_{l \in L} [\max\{\frac{f_s(l)}{f_s^{MAX}(l)} - 1, 0\}]^2$ ;
6    $I(s) \leftarrow \max_{a \in A(s)} \{\gamma \cdot \sum_{s' \in S} P(s'|s, a) [\Delta F(s, s') + I(s')]\}$ ;
7 end
8 Color[current_state]  $\leftarrow$  Gray;
9 Enqueue( $Q$ , current_state);
10 while (get_time()  $\leq$  deadline) and ( $Q \neq \emptyset$ ) do
11    $s \leftarrow$  Dequeue( $Q$ );
12   for  $a \in A(s)$  do
13      $R(s, a) \leftarrow \sum_{s' \in S} P(s'|s, a) [\Delta F(s, s') + I(s')]$ ;
14     Insert( $B$ , [ $R(s, a)$ ,  $s, a$ ]);
15   end
16   Sort( $B$ );
17   Concatenate(ContingencyList,  $B$ );
18   for  $b \in B$  do
19     if Color[ $s'_{b,s,b,a}$ ] = White then
20       Color[ $s'_{b,s,b,a}$ ]  $\leftarrow$  Gray;
21       Enqueue( $Q$ ,  $s'_{b,s,b,a}$ );
22     end
23   end
24    $B \leftarrow \emptyset$ ;
25   Color[s]  $\leftarrow$  Black;
26 end

```

time instant, the algorithm keeps several states active that are denoted by the gray color and stored in the initialized queue, where the power grid's current state is the first one (Lines 8,9). SOCCA updates the queue with respect to the most important state from the attacker's point of view, i.e., the worst-case scenario, which is the state with the highest expected degree of cyber-physical damage. From any state, $I(s)$ calculates the overall cyber-physical security index based on $F(s)$.

During an iterative process, the algorithm removes one state s at a time from the queue (Lines 11) and explores all possible immediate contingencies from that state, identified by $a \in A(s)$ (Lines 12). In particular, for each immediate contingency, SOCCA calculates an expected value of the state from the point of view of the adversary, assuming that all his or her actions in the future are optimal except the immediate next one. The results of these predictions are saved in the buffer list (Lines 13, 14). $R(s, a)$ indicates the benefits of taking malicious action a from state s . Once done with all the immediate contingencies from the state s , SOCCA sorts the results according to $R(s, a)$ and appends a copy of the ordered buffer to the contingency list output (Lines 16, 17). To update the data structures, SOCCA checks the individual elements in the buffer, and colors the destination states (assuming successful transitions) $s'_{b,s,b,a}$ as gray if they are still white and adds them to the FIFO queue (Lines 18-21). Finally, to prepare

for the next iteration, SOCCA clears the temporal buffer list and colors the analyzed state s as black so that it will not be checked again if encountered in future analysis due to the directed loops within the MDP (Lines 24, 25).

It is important to clarify that SOCCA can consider multiple power line outages. In particular, as explained, the proposed framework considers all possible attack paths, each of which often consists of many attack steps, namely several cyber asset vulnerability exploitations as well as one or more maliciously induced physical contingencies such as a power line outage. However, the attack steps cannot grow arbitrarily and are limited by the initial attack point (where the attacker initially resides), i.e., the Internet in our implementations, as well as the network global access policy rules and the system vulnerabilities.

Algorithm 1 can also be employed as an online solution to provide power grid security officers with predictive situational awareness capabilities. Indeed, SOCCA can monitor how future actions by attackers could globally impact the power grid given the current system state. This information is extremely valuable for proactive intrusion prevention systems which reconfigure the system such that the maximum possible damage to the system caused by the attackers' potential next action is minimized. SOCCA enables the officers to decide which critical components should be monitored more closely in order to detect potential exploitations of known or unknown vulnerabilities. As the algorithm focuses on contingencies originating from malicious cyber attacks, initial contingencies are usually all remote cyber-side vulnerability exploitations. This is because physical devices are almost never directly accessible from a remote machine, unless the attacker has already penetrated deep into the control network. SOCCA takes into consideration possible physical contingencies once states with required set of privileges (compromised host systems) for physical consequences have been reached by the algorithm.

The algorithm considers the path traversal difficulty and the final impact in calculation of the ultimate security index. Therefore, the less likely, yet large impact contingencies may not be fully evaluated during the analysis of Algorithm 1 as they would be ordered further behind while considering the deadline limit of on-line application. To address this point, the algorithm can be updated to always consider the contingencies with high impact and low likelihood. In particular, if the difficulty level of all the cyber-side attacks (i.e., vulnerability exploitations) is set to 0, SOCCA sorts the contingencies only based on their impact level. Consequently, the modified algorithm could partially devote its time to analyze the contingencies based on only their final physical impact.

The description of the contingency screening algorithm so far assumes that its input set including the current system MDP state (set of compromised hosts) and power system state are correct. However, an attacker with the required privilege levels could potentially mislead SOCCA by corrupting the input system state that results in incorrect overall contingency rankings. As a cyber-physical solution against the false data injection attacks against power system state, we have proposed security-oriented cyber-physical state estimation (SCPSE) in [17]. SCPSE takes into account both cyber-side intrusion

detection system (IDS) alerts as well as power system sensor measurements to identify corrupted measurements and ignore them during the power system state estimation process to obtain an accurate system state despite false data injection attacks. SOCCA can be extended to leverage SCPSE algorithms to deal with false data injection attacks on power system state. For false data injection against MDP state, our cyber-side system state (MDP) estimation algorithm can handle bad or false IDS alarms to the extent captured by their false positive and negative rates. However, verifying IDS alerts themselves is outside the scope of this work. That said, the issue of IDS alert verification (*e.g.*, [18]–[20]) and IDS trust management has received considerable research attention over the years (*e.g.*, [21]–[24]).

VI. EVALUATION

In this section, we discuss the implementation of SOCCA and present experimental evaluation results. All our experiments were performed on a 32-bit system with an Intel Core 2 2.16 GHz CPU, 3.00 GB of memory, and the Windows 7 Professional operating system.

Implementations. A unified XML [25] format was used to describe the power system control network topology and network access policy rules (*e.g.*, firewall rules). During the offline phase, SOCCA uses the NetAPT tool [26] to perform a comprehensive security analysis of the access policy rules and to produce the network connectivity matrix according to the control network topology input. The matrix is later translated to the corresponding MDP model. On the power side, we used PowerWorld Simulator [27] to simulate the underlying power system and solve the power flow equations to calculate the cyber-physical security index. In particular, we used the SimAuto toolbox to set up a real-time connection to PowerWorld.

In our experiments, we evaluated SOCCA on a simulated power grid infrastructure. The underlying power system was the IEEE 24-bus reliability test system [28] (Figure 2(a)). The power system consisted of 38 transmission lines, and was monitored and controlled by two control networks with identical network topologies and access control policies. The control network models were built based on topology of a real power control network which is kept anonymous due to non-disclosure agreement. Figure 2(b) shows the topology of a single control network that has 59 nodes, *e.g.*, host systems and firewalls. The first control network monitors and controls buses 1–12 in the power system (Figure 2(a)), and the second network monitors and controls buses 13–24. In particular, each power bus is monitored and controlled by a single host system in the corresponding control network.

MDP generation. Given the power network topology and the access policy rules, *i.e.*, about 100 firewall rules, SOCCA constructed the network connectivity matrix and generated the corresponding MDP model. It is noteworthy that because the MDP models may not be scalable specially for large-scale power-grid infrastructures, SOCCA makes use of the envelope algorithm [29], where the MDP is generated partially, and hence, not every individual state needs to be enumerated and analyzed. More technically, given the current system state,

only reachable states up to some finite horizon are explored and used for the contingency analysis. Figure 3 illustrates a simplified version of the generated MDP in which states with contingencies that are exclusively cyber are drawn in white, while states with physical consequences are in gray. The first number on each state represents its ID. Table I maps each state ID to the IDs of the compromised assets (shown in Figure 2) in each state. As shown on the generated MDP, the attacker initially resides remotely in the internet with no privilege on the power network (MDP’s state 0) and can then traverse different attack paths to access a particular host in the power network. Each MDP edge represents an access (*i.e.*, possibly a vulnerability exploitation) allowed from a source to a destination host in the power network.

Performance. To validate SOCCA’s efficiency on various networks with different sizes and topologies, we measured how long it takes to generate the MDP model for randomly-generated power networks. One important parameter that affects the model generation complexity is the *vulnerability factor*. This factor is defined as the number of host computers that could be accessed *and* compromised from a particular host in the power control network. Figures 4(a) and 4(b) show the MDP generation time requirement and the model’s size for vulnerability factor 1, *i.e.*, once the attacker compromises any host system, he or she can *always* find one other vulnerable machine to compromise. The results were averaged over 1000 runs. As illustrated, for large-scale power networks with 330K host computers, SOCCA analyzed the inputs and generated the MDP model within 24 milliseconds.

For cases with the vulnerability factor of 2, SOCCA generated the corresponding MDP graph within 400 milliseconds for a network with 37 nodes (Figures 4(c) and 4(d)). Figure 5 shows the time taken to solve the security index of individual states. A vulnerability factor of two means that from *every* host, two other host computers are both accessible (*i.e.*, access control policies allow that) and vulnerable to exploitations. We believe this is very pessimistic and not a very common scenario in industrial control infrastructures due to the extremely strict global access control policies and system security patching within real-world power control networks. The main intention in doing this experiment in the paper was to evaluate, under such pessimistic assumptions, how the increasing network size affects size of the power grid’s MDP model. We continued the experiment for cases with even higher (and much less realistic) vulnerability factors. The model generation for a network with 18 and vulnerability factor of 3 (4) nodes took 74 (5521) milliseconds on average. As expected, for a fixed network topology, the generated model size, and hence the overall performance overhead increases exponentially with the increasing vulnerability factor, which acts as a graph-theoretic branching factor in the MDP model generation procedure (Section III-C). To increase the SOCCA’s scalability for such cases, we implemented the *envelope* approximation algorithm [29] that, briefly, is based on the control theoretic finite look-ahead optimization technique and considers the next finite set of contingencies while generating the corresponding MDP model. Figures 4(e) and 4(f) show, respectively, how long SOCCA takes, using the

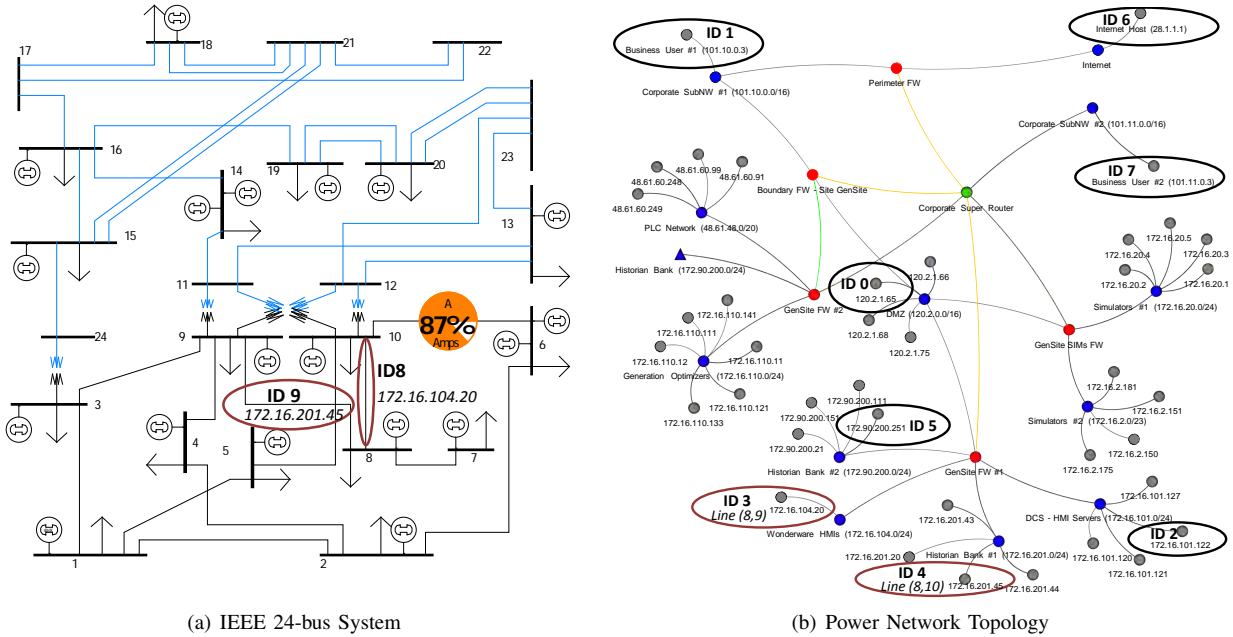


Fig. 2: Experimental Power Grid Testbed Architecture (with Compromised Asset IDs)

envelope algorithm, to generate the MDP model of the power grid network of different sizes with the vulnerability factor of 4, and the generated MDP model size.

Metrics. In our experiments, we pessimistically assumed that all the hosts include security vulnerabilities in order to perform a worst case performance analysis. SOCCA calculates the performance and security indices for individual states in the generated MDP (i.e., shown as second and third number on each state in Figure 3, respectively). In an MDP, there are usually many states with an identical set of physical contingencies that result in equal performance index values. To accelerate the metric calculation and minimize the number of connections to and calculations by PowerWorld, which is a time-consuming step due to the AC power flow solution procedures, SOCCA employs a caching solution to calculate the performance index value for each physical contingency set only once.

Contingency Ranking. SOCCA implements Algorithm 1 to rank various security incidents that could occur according to the system’s current state and the generated MDP model once the performance and security indices are calculated for the power grid’s corresponding MDP model. Table II shows the ranked list of cyber-physical contingencies for each state in our case study power grid. It is important to mention that the reported results are for the case in which the attacker has not yet caused any contingency in the power grid, i.e., the current state is $s_0 = \emptyset$ with ID 0. As shown, the edge $s_0 \rightarrow s_{74}$ is ranked as the most critical contingency as it allows the attacker to get to the most impactful physical consequence with the least amount of cyber exploitation effort.

VII. RELATED WORK

We review the related literature and highlight particular aspects wherein they fall short. Furthermore, we discuss how

SOCCA addresses those issues.

There have been several research efforts for computer network contingency analysis, which could be used to enumerate the set of next possible adversarial actions. Static adversary-driven security assessment techniques [30]–[32] explore potential malicious technical actions (contingencies) for every system state before the system goes operational. For instance, ADVISE [30] creates an executable state-based security model of a system and an adversary that represents how the adversary is likely to attack the system and the results of such an attack. As modeling and accurate prediction of the attacker’s behavior are very hard if not impossible in practice, defense-centric security assessment approaches have also been explored recently [33]–[36]. These techniques use manually filled knowledge bases of alert applicability, system configuration, or target importance to associate a context with each alert and to provide security assessment accordingly. The main barriers for real-world deployment of those techniques are the high human involvement required and the lack of awareness for future adversarial actions and social factors. However, static exploration of the whole state space is infeasible in practice due to the state explosion problem and furthermore, physical system contingencies are not accounted for. By design, SOCCA considers attacks consisting of power and/or cyber contingencies and explores a remarkably reduced search space by dynamic generation of the next possible states given the current probabilistic state estimate.

Contingency analysis in power systems has been explored by many researchers in the past (see [37] for a comprehensive survey). The initial efforts were based on first-order performance index sensitivities to rank contingencies, Ejebe et al. [38]. There have been several follow-up attempts to improve the ranking quality by considering higher order sensitivities [39], [40]. Furthermore, there has been an increasing interest

TABLE I: Mappings between MDP State IDs (Figure 3) and Compromised Asset IDs

State	Assets	State	Assets	State	Assets	State	Assets	State	Assets	State	Assets	State	Assets	State	Assets	State	Assets
0	6	1	61	2	610	3	6102	4	61023	5	610234	6	6102349	7	61023495	8	610234957
9	6102349578	10	610234958	11	61023497	12	610234978	13	61023498	14	610235	15	6102354	16	61023547	17	610235478
18	61023548	19	6102357	20	61023578	21	6102358	22	610237	23	6102374	24	61023748	25	6102378	26	610238
27	6102384	28	6103	29	61034	30	610349	31	6103495	32	61034957	33	610349578	34	61034958	35	6103497
36	61034978	37	6103498	38	61035	39	610354	40	6103547	41	61035478	42	6103548	43	610357	44	6103578
45	610358	46	61037	47	610374	48	6103748	49	610378	50	61038	51	610384	52	6104	53	61049
54	610492	55	6104925	56	61049257	57	6104927	58	610495	59	6104957	60	610497	61	6105	62	61052
63	610524	64	6105247	65	610527	66	61054	67	610547	68	61057	69	6107	70	61072	71	610724
72	61074	73	617	74	60	75	61024	76	602	77	6023	78	60234	79	602349	80	6023495
81	60234957	82	602349578	83	60234958	84	6023497	85	60234978	86	6023498	87	60235	88	602354	89	6023547
90	60235478	91	6023548	92	602357	93	6023578	94	602358	95	60237	96	602374	97	6023748	98	602378
99	60238	100	602384	101	603	102	6034	103	60349	104	603495	105	6034957	106	60349578	107	6034958
108	603497	109	6034978	110	603498	111	6035	112	60354	113	603547	114	6035478	115	603548	116	60357
117	603578	118	60358	119	6037	120	60374	121	603748	122	60378	123	6038	124	60384	125	604
126	6049	127	60492	128	604925	129	6049257	130	604927	131	60495	132	604957	133	60497	134	605
135	6052	136	60524	137	605247	138	60527	139	6054	140	60547	141	6057	142	607	143	6072
144	60724	145	6074	146	67												

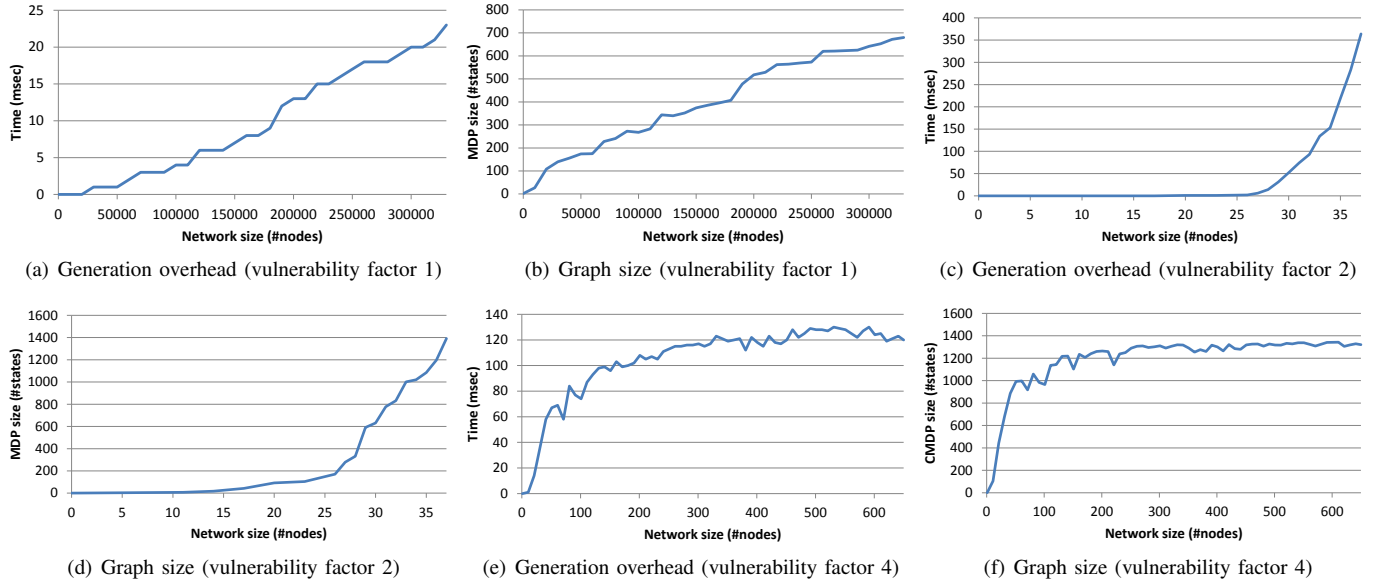


Fig. 4: Offline Automated MDP Graph Generation

in the analysis of multiple contingencies [41], [42] after the introduction of new NERC standards [43]. Davis et al. [44] propose a linear sensitivity-based approximate measure of how close the power system is brought to islanding by a particular outage contingency. The authors use the metric to categorize various line outages and show that it outperforms similar metrics because of taking care of precisely islanding singularities.

Almost all of the past contingency analysis techniques consider natural incidents to be root causes of the power system contingencies. As a result, they ignore cyber side events and, in particular, contingencies due to deficient or compromised cyber components. To deal with those issues, SOCCA introduces a cyber-physical security formalism that takes into account failure scenarios due to compromised cyber and/or power components.

During the last decade, several researchers have approached the problem of hybrid cyber-physical security modeling for the power-grid from different angles. Mo et al. [4] reviews a series of security challenges and possible intrusions in power-grid infrastructure. However, a unified modeling framework is not provided. Pasqualetti et al. [6] model a power system under cyber-physical attack as a linear time-invariant descriptor

system with unknown inputs, and design a dynamic detection and identification scheme using geometric control theoretic tools. It is not clear from the paper how the cyber network topology affects the way attacks occur in the modeled cyber-physical system. Sridhar et al. [7] review how traditional intrusion tolerance techniques could be applied in cyber-physical settings, and introduce a layered approach to evaluate risk based on the current state of the power-grid. The authors do not discuss how the state is determined, and additionally, accidental failures are not accounted for. Zonouz et al. [8] proposed a framework that fuses uncertain information from different types of distributed sensors, such as power system meters and cyber-side intrusion detectors, to detect the malicious activities within the cyber-physical system. Specifically, they presented a security-oriented cyber-physical state estimation (SCPSE) system, which, at each time instant, identifies the compromised set of hosts in the cyber network and the maliciously modified set of measurements obtained from power system sensors. However, to our knowledge SOCCA is the first framework to consider cyber adversary induced physical contingencies.

VIII. CONCLUSIONS

In this paper, we presented SOCCA, a security-oriented cyber-physical contingency analysis framework that identifies

TABLE II: Ranked (R) List of Cyber-Physical Contingencies (MDP edges)

R	Edge	R	Edge	R	Edge	R	Edge	R	Edge	R	Edge	R	Edge	R	Edge	R	Edge		
0	0.74	1	0.1	2	0.146	3	74.125	4	74.101	5	74.2	6	74.134	7	74.142	8	74.76	9	1.2
10	1.73	11	146.142	12	146.73	13	125.126	14	101.123	15	101.102	16	101.111	17	101.119	18	101.28	19	101.77
20	2.52	21	2.28	22	2.3	23	2.61	24	2.69	25	134.139	26	134.111	27	134.61	28	134.135	29	134.141
30	142.145	31	142.119	32	142.69	33	142.143	34	142.141	35	76.77	36	73.69	37	126.53	38	126.127	39	126.103
40	126.131	41	126.133	42	123.99	43	123.124	44	123.118	45	123.122	46	123.50	47	102.103	48	111.118	49	111.112
50	111.116	51	111.38	52	111.87	53	119.122	54	119.120	55	119.46	56	119.116	57	119.95	58	28.50	59	28.29
60	28.38	61	28.46	62	28.4	63	77.99	64	77.78	65	77.95	66	77.4	67	77.87	68	52.53	69	52.75
70	52.29	71	52.66	72	52.72	73	3.75	74	3.4	75	3.62	76	3.70	77	61.66	78	61.38	79	61.62
80	61.68	81	69.72	82	69.46	83	69.70	84	69.68	85	139.131	86	139.66	87	139.136	88	139.112	89	139.140
90	135.136	91	135.87	92	135.62	93	135.138	94	141.140	95	141.116	96	141.68	97	141.138	98	145.133	99	145.72
100	145.144	101	145.120	102	145.140	103	143.144	104	143.95	105	143.70	106	143.138	107	53.54	108	53.30	109	53.58
110	53.60	111	127.79	112	127.54	113	127.128	114	127.130	115	103.79	116	103.104	117	103.108	118	103.30	119	103.110
120	131.58	121	131.128	122	131.104	123	131.132	124	133.60	125	133.130	126	133.108	127	133.132	128	99.100	129	99.94
130	99.98	131	99.26	132	124.110	133	124.100	134	124.115	135	124.121	136	124.51	137	118.94	138	118.115	139	118.117
140	118.45	141	122.98	142	122.121	143	122.117	144	122.49	145	50.26	146	50.51	147	50.45	148	50.49	149	112.104
150	112.115	151	112.113	152	112.39	153	112.88	154	116.117	155	116.113	156	116.43	157	116.92	158	38.45	159	38.39
160	38.43	161	38.14	162	87.94	163	87.88	164	87.14	165	87.92	166	120.108	167	120.121	168	120.113	169	120.47
170	120.96	171	46.49	172	46.47	173	46.43	174	46.22	175	95.98	176	95.96	177	95.92	178	95.22	179	29.30
180	29.51	181	29.39	182	29.47	183	29.5	184	4.26	185	4.5	186	4.22	187	4.14	188	78.79	189	78.100
190	78.96	191	78.5	192	78.88	193	75.54	194	75.5	195	75.63	196	75.71	197	66.58	198	66.63	199	66.39
200	66.67	201	72.60	202	72.71	203	72.47	204	72.67	205	62.63	206	62.14	207	62.65	208	70.71	209	70.22
210	70.65	211	68.67	212	68.43	213	68.65	214	136.128	215	136.63	216	136.88	217	136.137	218	140.132	219	140.137
220	140.113	221	140.67	222	138.137	223	138.92	224	138.65	225	144.130	226	144.71	227	144.96	228	144.137	229	54.6
230	54.55	231	54.57	232	30.6	233	30.31	234	30.35	235	30.37	236	58.55	237	58.31	238	58.59	239	60.57
240	60.35	241	60.59	242	79.80	243	79.84	244	79.6	245	79.86	246	128.55	247	128.80	248	128.129	249	130.57
250	130.84	251	130.129	252	104.80	253	104.105	254	104.31	255	104.107	256	108.35	257	108.84	258	108.105	259	108.109
260	110.86	261	110.107	262	110.109	263	110.37	264	132.59	265	132.129	266	132.105	267	100.86	268	100.91	269	100.97
270	100.27	271	94.91	272	94.93	273	94.21	274	98.97	275	98.93	276	98.25	277	26.27	278	26.21	279	26.25
280	115.107	281	115.114	282	115.42	283	115.91	284	121.109	285	121.114	286	121.48	287	121.97	288	51.37	289	51.27
290	51.42	291	51.48	292	117.93	293	117.114	294	117.44	295	45.21	296	45.42	297	45.44	298	49.25	299	49.48
300	49.44	301	113.105	302	113.114	303	113.40	304	113.89	305	39.31	306	39.42	307	39.40	308	39.15	309	88.80
310	88.91	311	88.15	312	88.89	313	43.44	314	43.40	315	43.19	316	92.93	317	92.89	318	92.19	319	14.21
320	14.15	321	14.19	322	47.35	323	47.48	324	47.40	325	47.23	326	96.84	327	96.97	328	96.23	329	96.89
330	22.25	331	22.23	332	22.19	333	5.6	334	5.27	335	5.23	336	5.15	337	63.55	338	63.15	339	63.64
340	71.57	341	71.23	342	71.64	343	67.59	344	67.40	345	67.64	346	65.64	347	65.19	348	137.129	349	137.89
350	137.64	351	6.7	352	6.11	353	6.13	354	55.7	355	55.56	356	57.11	357	57.56	358	31.7	359	31.32
360	31.34	361	35.11	362	35.32	363	35.36	364	37.13	365	37.34	366	37.36	367	59.56	368	59.32	369	80.81
370	80.7	371	80.83	372	84.11	373	84.81	374	84.85	375	86.83	376	86.85	377	86.13	378	129.56	379	129.81
380	105.32	381	105.81	382	105.106	383	107.83	384	107.106	385	107.34	386	109.85	387	109.106	388	109.36	389	91.83
390	91.18	391	91.90	392	97.85	393	97.24	394	97.90	395	27.13	396	27.18	397	27.24	398	93.90	399	93.20
400	21.18	401	21.20	402	25.24	403	25.20	404	114.106	405	114.41	406	114.90	407	42.34	408	42.41	409	42.18
410	48.36	411	48.41	412	48.24	413	44.20	414	44.41	415	40.32	416	40.41	417	40.16	418	89.81	419	89.90
420	89.16	421	15.7	422	15.18	423	15.16	424	19.20	425	19.16	426	23.11	427	23.24	428	23.16	429	64.56
430	64.16	431	7.8	432	7.10	433	11.8	434	11.12	435	13.10	436	13.12	437	56.8	438	32.8	439	32.33
440	34.10	441	34.33	442	36.12	443	36.33	444	81.8	445	81.82	446	83.82	447	83.10	448	85.82	449	85.12
450	106.82	451	106.33	452	18.10	453	18.17	454	90.82	455	90.17	456	24.12	457	24.17	458	20.17	459	41.33
460	41.17	461	16.8	462	16.17	463	8.9	464	10.9	465	12.9	466	33.9	467	82.9	468	17.9		

contingencies possible through cyber attacks (e.g., malicious vulnerability exploitations), given the current cyber security state of the power system control network. SOCCA provides the power grid security officers with predictive situational awareness capabilities to assess the global impact of different adversarial actions on the power grid. Thus, it enables operators to decide upon appropriate deployment of proactive intrusion prevention solutions. Our experimental results show that SOCCA complements the traditional power contingency analysis methods that consider physical power component failures due only to accidental failures and other natural causes.

ACKNOWLEDGMENTS

This material is based upon work supported by the Department of Energy under Award Number DE-OE0000097.

REFERENCES

[1] A. Monticelli, *State Estimation in Electric Power Systems: A Generalized Approach*. Kluwer Academic Publishers, 1999.

[2] O. Alsac, N. Vempati, B. Stott, and A. Monticelli, "Generalized state estimation," *IEEE Transactions on Power Systems*, vol. 13, no. 3, pp. 1069–1075, Aug. 1998.

[3] J. J. Grainger and W. D. Stevenson, *Power System Analysis*. McGraw Hill, 1994.

[4] Y. Mo, T. H.-J. Kim, K. Brancik, D. Dickinson, H. Lee, A. Perrig, and B. Sinopoli, "Cyber-physical security of a smart grid infrastructure," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 195–209, Jan. 2012.

[5] T. M. Chen, S. Member, J. C. Sanchez-aarnoutse, and J. Buford, "Petri net modeling of cyber-physical attacks on smart grid," *Smart Grid IEEE Transactions on*, vol. 2, no. 99, pp. 1–9, 2011. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5967924

[6] F. Pasqualetti, F. Dörfler, and F. Bullo, "Cyber-physical attacks in power networks: Models, fundamental limitations and monitor design," *CoRR*, vol. abs/1103.2795, 2011.

[7] S. Sridhar, A. Hahn, and M. Govindarasu, "Cyber-physical system security for the electric power grid," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 210–224, 2012.

[8] S. Zonouz, K. Rogers, R. Berthier, R. Bobba, W. Sanders, and T. Overbye, "Scpsp: Security-oriented cyber-physical state estimation for power grid critical infrastructures," *Smart Grid, IEEE Transactions on*, vol. 3, no. 4, pp. 1790–1799, 2012.

[9] B. Stott and O. Alsac, "Fast decoupled load flow," *IEEE Transactions on Power Apparatus and Systems*, vol. 93, no. 3, pp. 859–869, May 1974.

[10] NERC, *Transmission transfer capability: a reference document for calculating and reporting the electric power transfer capability of interconnected electric systems*. NERC, 1995.

[11] —, (2008) Standard mod-030-1 flowgate methodology. [Online]. Available: www.nerc.com/files/MOD-030-1.pdf

[12] C. M. Davis and T. J. Overbye, "Multiple element contingency screening," *IEEE Transactions on Power Systems*, vol. 26, no. 3, pp. 1294–1301, Aug. 2011.

[13] J. Filar and K. Vrieze, *Competitive Markov Decision Processes*. Springer-Verlag, 1997.

[14] "NULL pointer dereference in Linux Kernel 2.6.0," National Vulnerability Database, August 2009, available at <http://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2009-2692>. [Online]. Available: <http://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2009-2692>

[15] P. Mell, K. Scarfone, and S. Romanosky, "Common vulnerability scoring system," *Security Privacy, IEEE*, vol. 4, no. 6, pp. 85–89, nov.-dec. 2006.

[16] B. Stott, J. Jardim, and O. Alsac, "Dc power flow revisited," *Power Systems, IEEE Transactions on*, vol. 24, no. 3, pp. 1290–1300, 2009.

[17] S. Zonouz, K. Rogers, R. Berthier, R. Bobba, W. Sanders, and T. Overbye, "Scpsp: Security-oriented cyber-physical state estimation for power grid critical infrastructures," *IEEE Transactions on Smart Grid*, vol. 3, no. 4, pp. 1790–1799, 2012.

[18] C. Kruegel and W. K. Robertson, "Alert verification determining the success of intrusion attempts," in *DIMVA*, 2004, pp. 25–38.

[19] C. Mu, H. Huang, and S. Tian, "Intrusion detection alert verification based on multi-level fuzzy comprehensive evaluation," in *Computational Intelligence and Security*. Springer, 2005, pp. 9–16.

[20] C. Kruegel, W. Robertson, and G. Vigna, "Using alert verification to identify successful intrusion attempts," *Praxis der Informationsverarbeitung und Kommunikation*, vol. 27, no. 4, pp. 219–227, 2004.

[21] S. Ruohomaa and L. Kutvonen, "Trust management survey," in *Trust Management*. Springer, 2005, pp. 77–92.

[22] D. Frincke, D. Tobin, J. McConnell, J. Marconi, and D. Polla, "A framework for cooperative intrusion detection," in *Proceedings of the 21st national information systems security conference*, 1998, pp. 361–373.

[23] F. Cuppens, "Managing alerts in a multi-intrusion detection environment," in *acsac*, vol. 1, 2001, p. 22.




Fig. 3: Automatically Generated MDP for the Power Grid

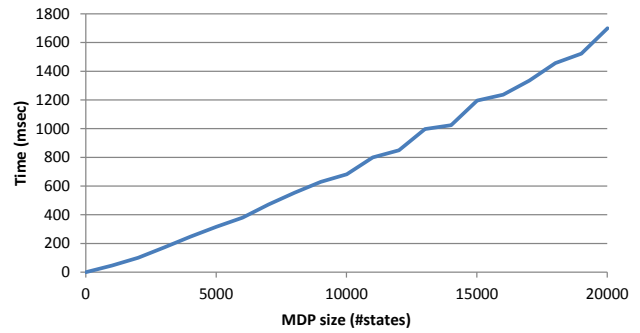


Fig. 5: Online Index Calculation

[24] C. J. Fung, J. Zhang, I. Aib, and R. Boutaba, "Robust and scalable trust management for collaborative intrusion detection," in *Integrated Network Management, 2009. IM'09. IFIP/IEEE International Symposium on*. IEEE, 2009, pp. 33–40.

[25] *Xml: The Complete Reference*. Osborne, 2001. [Online]. Available: <http://books.google.com/books?id=1XErctH8PhoC>

[26] D. M. Nicol, W. H. Sanders, S. Singh, and M. Seri, "Usable global network access policy for process control systems," *IEEE Security and Privacy*, vol. 6, pp. 30–36, 2008.

[27] J. Glover, M. Sarma, T. Overbye, and T. Overbye, *Power system analysis and design*. Thomson, 2008.

[28] R. T. S. T. F. of the Application of Probability Methods Subcommittee, "IEEE reliability test system," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-98, no. 6, pp. 2047–2054, Nov. 1979.

[29] T. Dean, L. Kaelbling, J. Kirman, and A. Nicholson, "Planning under time constraints in stochastic domains," *Artificial Intelligence*, vol. 76, pp. 35–74, July 1995.

[30] E. LeMay, W. Unkenholz, D. Parks, C. Muehrcke, K. Keefe, and W. H. Sanders, "Adversary-driven state-based system security evaluation," in *Proceedings of the 6th International Workshop on Security Measurements and Metrics*, ser. MetriSec '10. New York, NY, USA: ACM, 2010, pp. 5:1–5:9. [Online]. Available: <http://doi.acm.org/10.1145/1853919.1853926>

[31] L. Wang, T. Islam, T. Long, A. Singhal, and S. Jajodia, "An attack graph-based probabilistic security metric," in *Data and Applications Security XXII*, ser. Lecture Notes in Computer Science, V. Atluri, Ed. Springer Berlin / Heidelberg, 2008, vol. 5094, pp. 283–296, 10.1007/978-3-540-70567-3_22. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-70567-3_22

[32] E. LeMay, M. D. Ford, K. Keefe, W. H. Sanders, and C. Muehrcke, "Model-based security metrics using adversary view security evaluation (advise)," in *QEST*. IEEE Computer Society, 2011, pp. 191–200.

[33] J. Yu, Y. Ramana Reddy, S. Selliah, S. Reddy, V. Bharadwaj, and S. Kankanahalli, "TRINETR: An architecture for collaborative intrusion detection and knowledge-based alert evaluation," *Advanced Engineering Informatics*, vol. 19, no. 2, pp. 93–101, 2005.

[34] P. Porras, M. Fong, and A. Valdes, "A mission-impact-based approach to INFOSEC alarm correlation," in *Proceedings of the Symposium on Recent Advances in Intrusion Detection*. Springer-Verlag, 2002, pp. 95–114.

[35] K. Alsubhi, E. Al-Shaer, and R. Boutaba, "Alert prioritization in intrusion detection systems," in *IEEE Network Operations and Management Symposium*, 2008, pp. 33–40.

[36] W. Lee and X. Qin, "Statistical causality analysis of infosec alert data," *Managing Cyber Threats*, pp. 101–127, 2005.

[37] B. Stott, O. Alsac, and A. F.L., "Analytical and computational improvements in performance index ranking algorithms for networks," *International Journal of Electrical Power and Energy Systems*, vol. 7, no. 3, pp. 154–160, 1985.

[38] G. C. Ejebe and B. F. Wollenberg, "Automatic contingency selection," *IEEE Transactions on Power Apparatus and Systems*, vol. 65, no. 1, pp. 859–109, 1979.

[39] T. Mikolinnas and B. Wollenberg, "An advanced contingency selection algorithm," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-100, no. 2, pp. 608–617, feb. 1981.

[40] G. Irisarri and A. Sasson, "An automatic contingency selection method for on-line security analysis," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-100, no. 4, pp. 1838–1844, april 1981.

[41] T. Guler and G. Gross, "Detection of island formation and identification of causal factors under multiple line outages," *Power Systems, IEEE Transactions on*, vol. 22, no. 2, pp. 505–513, may 2007.

[42] T. Halpin, R. Fischl, and R. Fink, "Analysis of automatic contingency selection algorithms," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-103, no. 5, pp. 938–945, may 1984.

[43] NERC. (2005) System performance following loss of two or more bulk electric system elements (category c). [Online]. Available: www.nerc.com/files/TPL-003-0.pdf

[44] C. Davis and T. Overbye, "Multiple element contingency screening," *Power Systems, IEEE Transactions on*, vol. 26, no. 3, pp. 1294–1301, aug. 2011.

Saman Zonouz is an Assistant Professor in the Electrical and Computer Engineering Department at the University of Miami. He received his Ph.D. in Computer Science from the University of Illinois at Urbana-Champaign in 2011. He has worked on intrusion response and recovery, information flow-based security metrics for power-grid critical infrastructures, online digital forensics analysis and monitor-less recoverable applications. His research interests include: computer security and survivable systems, control/game theory, intrusion response and recovery systems, automated intrusion forensics analysis, and information flow analysis-based security metrics.



Charles M. Davis received the B.S. degree in electrical engineering from Louisiana Tech University in 2002 and the M.S. and Ph.D. degrees from the Electrical and Computer Engineering Department at the University of Illinois at Urbana-Champaign. He is currently working at PowerWorld Corporation. His research interests include linear sensitivities, power system analysis, power system visualization, and power system operational reliability.

Katherine R. Davis received the B.S. degree in electrical engineering from the University of Texas at Austin in 2007 and the M.S. and Ph.D. degrees from the University of Illinois Urbana-Champaign in 2009 and 2011. She is a software engineer and senior consultant at PowerWorld Corporation and an Adjunct Assistant Professor in Electrical and Computer Engineering Department at the University of Illinois Urbana-Champaign. Her research interests include data-enhanced power system modeling and analysis and making the grid more robust with respect to bad data.

Robin Berthier is a research scientist at the University of Illinois at Urbana-Champaign, working with Prof. William H. Sanders. Robin graduated from the Reliability Engineering Department at the University of Maryland in 2009. His doctoral dissertation with Prof. Michel Cukier, and focused on the issue of honeypot sensors deployed on large networks. He introduced a new architecture to increase the scalability of high-interaction honeypots, and combined network datasets of different granularities to offer unique attack forensics capabilities. His current research interests include advanced intrusion detection systems and the security of critical infrastructures.



Rakesh B. Bobba is a Research Assistant Professor in the College of Engineering at the University of Illinois, Urbana-Champaign with appointments in Information Trust Institute and Electrical and Computer Engineering Department. His research interests are in the security of distributed and networked systems with a current focus on cyber-physical systems including critical infrastructures such as the power grid and cloud computing. He received M.S. and Ph.D. degrees in Electrical and Computer Engineering from the University of Maryland at College Park in 2007 and 2009, respectively.



William H. Sanders is a Donald Biggar Willett Professor of Engineering, the Interim Head of the Department of Electrical and Computer Engineering, and the Director of the Coordinated Science Laboratory (www.csl.illinois.edu) at the University of Illinois at Urbana-Champaign. He is a professor in the Department of Electrical and Computer Engineering and Affiliate Professor in the Department of Computer Science. He is a Fellow of the IEEE and the ACM, a past Chair of the IEEE Technical Committee on Fault-Tolerant Computing, and past Vice-Chair of the IFIP Working Group 10.4 on Dependable Computing. He was the founding Director of the Information Trust Institute (www.iti.illinois.edu) at Illinois.

