Article title: A System Dynamics Approach for Assessing the Impact of Cyber Attacks on Critical Infrastructures

Béla Genge, István Kiss, Piroska Haller
Petru Maior University of Tg. Mureș, Department of Informatics, Tg. Mureș, Romania
Email: bela.genge@ing.upm.ro, istvan.kiss@stud.upm.ro, phaller@upm.ro

Abstract – The massive proliferation of information and communications technologies (hardware and software) into the heart of modern critical infrastructures has given birth to a unique technological ecosystem. Despite the many advantages brought about by modern information and communications technologies, the shift from isolated environments to "systems-of-systems" integrated with massive information and communications infrastructures (e.g., the Internet) exposes critical infrastructures to signifcant cyber threats. Therefore, it is imperative to develop approaches for identifying and ranking assets in complex, large-scale and heterogeneous critical infrastructures. To address these challenges, this paper proposes a novel methodology for assessing the impact of cyber attacks on critical infrastructures. The methodology is inspired by research in system dynamics and sensitivity analysis. The proposed behavioral analysis methodology computes the covariances of the observed variables before and after the execution of a specific intervention involving the control variables. Metrics are proposed for quantifying the significance of control variables and measuring the impact propagation of cyber attacks. Experiments conducted on the IEEE 14-bus and IEEE 300-bus electric grid models, and on the well-known Tennessee Eastman chemical process demonstrate the efficiency, scalability and cross-sector applicability of the proposed methodology in several attack scenarios. The advantages of the methodology over graph-theoretic and electrical centrality metric approaches are demonstrated using several test cases. Finally, a novel, stealthy cyberphysical attack is demonstrated against a simulated power grid; this attack can be used to analyze the precision of modern anomaly detection systems..

A system dynamics approach for assessing the impact of cyber attacks on critical infrastructures

Béla Genge,[1] István Kiss, Piroska Haller

Department of Informatics, Petru Maior University of Tirgu Mures, N. Iorga Street, No. 1, Tirgu Mures, Mures, 540088 Romania

Abstract

The massive proliferation of information and communications technologies (hardware and software) into the heart of modern critical infrastructures has given birth to a unique technological ecosystem. Despite the many advantages brought about by modern information and communications technologies, the shift from isolated environments to "systems-of-systems" integrated with massive information and communications infrastructures (e.g., the Internet) exposes critical infrastructures to significant cyber threats. Therefore, it is imperative to develop approaches for identifying and ranking assets in complex, large-scale and heterogeneous critical infrastructures. To address these challenges, this paper proposes a novel methodology for assessing the impact of cyber attacks on critical infrastructures. The methodology is inspired by research in system dynamics and sensitivity analysis. The proposed behavioral analysis methodology computes the covariances of the observed variables before and after the execution of a specific intervention involving the control variables. Metrics are proposed for quantifying the significance of control variables and measuring the impact propagation of cyber attacks.

Experiments conducted on the IEEE 14-bus and IEEE 300-bus electric grid models, and on the well-known Tennessee Eastman chemical process demonstrate the efficiency, scalability and cross-sector applicability of the proposed methodology in several attack scenarios. The advantages of the methodology over graph-theoretic and electrical centrality metric approaches are demonstrated using several test cases. Finally, a novel, stealthy cyber-physical attack is demonstrated against a simulated power grid; this attack

[1] Corresponding author: Béla Genge (bela.genge@gmail.com)

can be used to analyze the precision of modern anomaly detection systems.

Keywords
Critical Infrastructures; Cyber Attacks; Impact Assessment; System Dynamics; Sensitivity Analysis; Smart Grid

Submitted: January 12, 2015
Revision: April 3, 2015
Accepted: April XX, 2015

1. Introduction

Modern critical infrastructure assets such as power plants, water supply systems and electric power grids are moving from isolated environments to "systems-of-systems" integrated with massive information and communications technology infrastructures such as the Internet. In fact, the proliferation of information and communications hardware and software into the heart of critical infrastructures has given birth to a unique technological ecosystem. Modern critical infrastructures encompass a variety of objects ranging from sensors and actuators, RFID systems, industrial equipment and video surveillance cameras to generic personal computers and networking devices [11]. These sophisticated critical infrastructures deliver advanced services and features, enhance controllability, reliability and safety, and facilitate the implementation of novel infrastructure paradigms such as the smart grid.

As a result of technological advancements, critical infrastructures are not only subject to traditional information and communications technology attacks, but also to a new breed of cyber-physical attacks. These complex and sophisticated attacks involve the exploitation of the cyber and physical characteristics of a critical infrastructure asset in order to cause significant damage. An exemplar is Stuxnet, which is widely believed to be the first piece of malware that was specifically designed to impact the operations of an industrial installation [4, 15]. In fact, Stuxnet's ability to rewrite control logic raises the specter of a new class of threats that exploit software vulnerabilities to cause considerable damage to physical processes. Symantec [28] recently revealed the existence of Dragonfly, a new targeted attack against critical infrastructures. The attack, which was originally thought to be directed at the energy sector, is now believed to be targeted against the pharmaceutical industry and may well be one of the most effective cyber espionage weapons to date. Dragonfly's ability to penetrate the core of industrial systems and steal proprietary data raises serious concerns about the capabilities of future malware.

Motivated by these threats, several researchers (see, e.g., [8, 14, 13]) have proposed approaches for enhancing the security of current and future critical infrastructures. In fact, a significant body of research has focused on understanding critical infrastructure vulnerabilities and quantifying the security risks of critical infrastructure implementations [5, 27]. One important outcome is the assessment of the impact of cyber attacks on the normal functioning of physical processes. In such situations, it is imperative not only to demonstrate and evaluate the disruptive impacts of cyber attacks [12], but also to quantify the impacts and ultimately provide rankings of the importance of specific cyber assets.

To address the existing gaps, this paper proposes a novel methodology for assessing the impacts of cyber attacks on critical infrastructures. The proposed Cyber Attack Impact Assessment (CAIA) methodology is inspired by research in system dynamics [9]. The methodology compares the behavior of complex physical processes in the presence and absence of accidental or deliberate interventions in order to evaluate the significance of cyber assets. The interventions may be required to respond to cyber attacks as well as faults and random events. The methodology is applicable to large-scale, hierarchical and heterogeneous installations, but most importantly, it may be used to evaluate the impacts of disturbances (e.g., caused by cyber attacks) in a variety of production systems. Since the methodology relies on input data (i.e., one-way communications from a physical process to an assessment engine) and leverages measurements that are already available at specific network points (e.g., data historians), it may be applied in critical production environments where the impacts of control commands can be assessed and quantified in order to identify critical control variables. This is a significant aspect of the proposed methodology, which differentiates it from existing methodologies. Furthermore, the results of the methodology can be used by various network planning and risk assessment techniques to verify the assurance of mechanisms that protect critical system components.

The proposed methodology is experimentally validated from several perspectives. First, the effects of various parameters on the results are assessed using the IEEE 14-bus electric grid model [23]. Next, the scalability of the methodology is demonstrated using attack scenarios implemented in the context of the IEEE 300-bus electric grid model [18]. Following this, the cross-sector applicability of the methodology is evaluated using the well-known Tennessee Eastman chemical process system [7]. The advantages of the methodology over graph-theoretic methodologies and electrical centrality measures [1, 29] are also demonstrated using several test cases. Finally, a novel stealthy attack that leverages the methodology results is evaluated using the IEEE 14-bus electric grid model. This attack is well suited to testing anomaly detection systems because it targets multiple low-impact cyber assets to cause major infrastructure disruptions.

2. Related work

Several researchers have attempted to evaluate and quantify the impacts of cyber attacks on physical processes in critical infrastructures. Kundur et al. [20] proposed a graph-based model to evaluate the influence of control loops on physical processes; this approach was used to assess the impacts of cyber attacks on electric power generation. Sgouras et al. [25] evaluated the impact of cyber attacks on a simulated smart metering infrastructure; the denial-of-service attacks against smart meters and utility servers caused severe communications interruptions. Sridhar and Govindarasu [26] evaluated the impacts of cyber attacks on wide-area frequency control applications in power systems; their research showed that cyber attacks can significantly impact system stability by causing severe drops in system frequency.

Bilis et al. [1] proposed a systematic approach that uses five metrics derived from complex network theory to assess the impacts of cyber attacks on electric power systems. The metrics were used to rank nodes in a graph-based representation of an electric grid. Despite its applicability to large electric power grids, the work of Bilis and colleagues has two key drawbacks: (i) while a structural assessment may provide interesting results for synthetic power grids, it cannot reproduce the behavior of a real physical process; and (ii) the approach is domain-dependent (i.e., focused on electric power systems) and its application to other domains (e.g., chemical plants) may not be feasible.

In order to address problems with structural assessments, Wang et al. [29] proposed electrical centrality metrics to identify the critical nodes in electric power systems. The approach leverages the electrical admittance matrix to infer the electrical properties of the underlying physical processes. Despite its proven applicability to power systems, the approach has two main shortcomings: (i) it does not account for the complexity of electric grids, where the behavior of the physical processes are governed by control loops implemented in the cyber realm; and (ii) as in the case of the graph-theoretic approach of Bilis et al. [1], the electrical centrality measures are specifically designed for electric power systems and are not applicable to other critical infrastructures.

Krotofil et al. [19] and Genge and Siaterlis [13] have studied the impacts of attacks on physical processes after they reach their emergency shutdown limits. Despite its applicability in the chemical sector, these approaches cannot deal with attacks that trigger perturbations that do not necessarily lead to system shutdown.

The research efforts discussed above focus on the impacts of cyber attacks on the normal functioning of physical processes from several perspectives. However, they are all directed at specific scenarios, physical processes and domain-dependent equipment. In contrast, the CAIA methodology proposed in this paper is applicable to a variety of critical infrastructures and can quantify the impacts of cyber attacks using simulation-based results as well as data originating from production systems. Moreover, the methodology can be applied in a hierarchical and multiphase fashion, which makes it highly scalable and appropriate for large-scale critical infrastructures. Indeed, unlike most existing approaches, the CAIA methodology accounts for the complexity of the physical and cyber dimensions of critical infrastructures and can be used in the presence of control loops that govern critical infrastructure behavior.

3. Cyber attack impact assessment methodology

This section presents the Cyber Attack Impact Assessment (CAIA) methodology. It begins with an overview of the architecture of a modern critical infrastructure, and proceeds to present the design considerations and details of the CAIA methodology.

3.1. Critical infrastructure architecture

A modern critical infrastructure has a hierarchical structure comprising two layers: (i) physical layer, which encompasses sensors, actuators and hardware devices that interact with the physical processes; (ii) and the cyber layer, which encompasses the information and communications technology

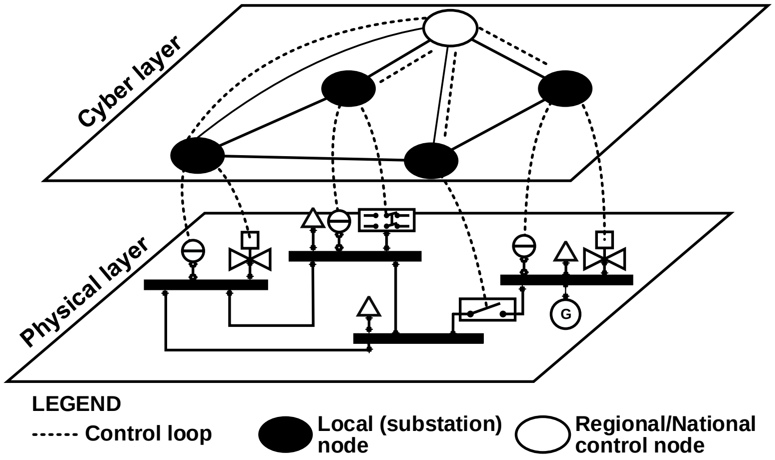

Figure 1: Critical infrastructure architecture.

hardware and software needed to monitor the physical processes and implement complex control loops. Figure 1 shows a critical infrastructure with its physical and cyber layers. From an operational point of view, hardware controllers (i.e., programmable logical controllers (PLCs)), receive data from sensors, elaborate a local actuation strategy and send commands to actuators. These hardware controllers also send the data received from sensors to supervisory control and data acquisition (SCADA) servers and execute the commands that they receive from the SCADA servers.

Critical infrastructure assets vary in size from a few sensors and generic personal computers to thousands of control objects, RFID devices, industrial equipment and video surveillance cameras organized in a hierarchical structure and distributed across a large geographical area. This technological ecosystem incorporates control loops at various locations to manage operations and ensure the correct functioning of the underlying physical processes. They may interact directly with sensor and actuator nodes or remotely with other controllers in a hierarchical and distributed control system architecture. Essentially, control loops constitute the core of a critical infrastructure and, therefore, need to be protected from threats to their cyber and physical dimensions.

3.2. Design considerations

The design assumes large-scale physical processes with several control loops spread across regional, national or international boundaries. Control

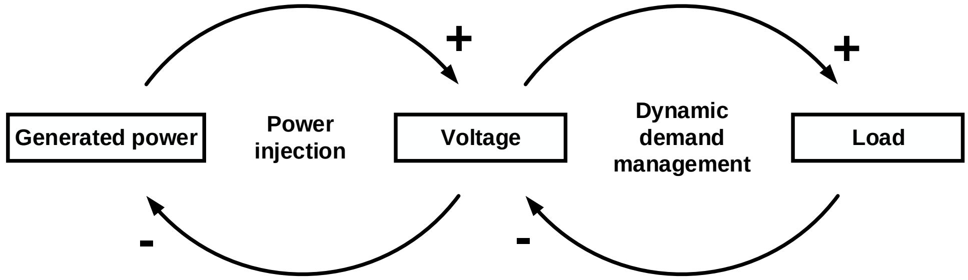

Figure 2: Causal loop for an electric power grid.

loops rely on observed variables and effect changes to the physical process state via control variables. A real installations may incorporate thousands of observed and control variables. The wide variety of cyber attacks and their potential implementations [3] raise significant concerns about the problem size and the ability to produce feasible cyber attack impact assessments.

In an attempt to address these challenges, the CAIA methodology adopts a procedure inspired by system dynamics [9]. System dynamics can explain how the dependencies between various cyber/physical assets and the process behavior change over time. The approach was originally proposed by Forrester [10] in the mid 1950s as a formal modeling methodology for understanding the behavior of complex systems over time. Since then, the approach has been applied in a variety of domains, including industrial processes, socioeconomic systems, policy analysis and design.

The main building blocks are causal loop diagrams that map system causalities and stock-flow diagrams that support qualitative system analyses. Causal loop diagrams are transformed into stock-flow diagrams in order to enable quantitative as well as causal understanding of system behavior. Figure 2 shows an example causal loop diagram for a simplified scenario involving an electric power grid. The positive causal link between the generated power and substation voltage levels denotes that changes in the generated power result in voltage level changes in the same direction (increase or decrease). On the other hand, the negative causal link between voltage and generated power denotes that voltage increases (i.e., above certain limits (over-voltages)) may trigger a decrease in the generated power (i.e., a change in the opposite direction). However, as shown in the figure, voltage levels are also influenced by user demand and may trigger the execution of dynamic load shedding algorithms to ensure grid stability. The most significant aspect of this example, however, is the causal relationship between load and generated power. Despite the absence of a direct causal link between the two, an indirect causal link can be identified via the substation voltage levels. For example, an increase in user demand leads to voltage level decreases that, in turn, could trigger an increase in the generated power.

Forrester's system dynamics approach was shown to be applicable in identifying control loop dominance [9] (i.e., control loops that dominate system behavior). The approach involves seven steps during which each variable of interest in a control loop is deactivated by setting a constant control variable value. Following this, the process model is run over a specific time interval to determine if it exhibits the same or different behavior pattern.

Forrester's approach was refined by Huang et al. [17] with regard to the sensitivity analysis of control loops. Sensitivity analysis measures how sensitive a model is to changes in its control parameter values. Huang et al. proposed a quantitative analysis of loop dominance, a significant improvement over Ford's work [9] that mainly provides binary (yes/no) answers. More specifically, Huang et al. measure sensitivity by computing the relative variance between model behavior in the deactivated control loop case and the reference model behavior. This approach ensures the quantification of the relative contribution of a control loop to the behavior of a certain variable of interest.

The CAIA methodology proposed in this work builds on the behavioral analysis of physical processes proposed by Ford [9] and the sensitivity assessment approach of Huang et al. [17]. However, the methodology is further refined and adapted in order to embrace the complexity and ubiquitous nature of cyber attacks on critical infrastructures. This leads to three important observations:

- Research has revealed that cyber attacks on critical infrastructures leverage a variety of attack vectors that may not necessarily include traditional hardware-based control loops [2]. For instance, human-machine interfaces (HMIs), servers and human operators are also integral critical infrastructure components that may close control loops at the local, regional or national levels by means of various information and communications hardware and software. Therefore, CAIA-based analysis is not limited to traditional control variables, but to any variable exposed to the cyber realm that may be compromised and used to influence the behavior of a physical process.

- CAIA does not simply deactivate control loops (i.e., by repeating the last value as in [17]); instead, it supports a variety of control parameter changes. This enables the quantification of the impact of diverse cyber attacks on the behavior of a physical process.

- CAIA is applied to individual control parameters and it does not provide a recipe for aggregating all possible attack scenarios and all available control parameter combinations. Such an attempt is unlikely to succeed given the complexity of cyber attacks and the large number of attack vectors available to malicious actors (as described in [2]). Instead, CAIA effectively quantifies the impact of specific types of attacks on the behavior of a physical process. These results are useful for ranking cyber assets based on their sensitivity to specific cyber attacks, identifying the most vulnerable cyber assets and designing control networks.

Essentially, CAIA computes the covariances of observed variables before and after the execution of a specific intervention. This provides a metric for quantifying the significance of each control variable on the correct functioning of a critical infrastructure. Compared with other approaches, CAIA is well suited to scenarios where the physical process model is not available as well as production systems for which control and measurement variable data is available. The latter is particularly useful in post-event analysis where the emphasis is on establishing the relative impacts of specific faults and of deliberate cyber attacks on the normal functioning of a system.

Figure 3 illustrates the application of the CAIA methodology to an infrastructure. CAIA uses the measured values of the observed variables to evaluate the impacts of deliberate interventions, possibly even cyber attacks, on the control variables. In a real-world scenario, CAIA may leverage system event logs to infer the start of the assessment period if the interventions are reported in the logs. For each such event, CAIA records the values of the observed variables and applies the technique described in this section to compute the impact ranking of a cyber asset (i.e., control variable). Note that, although the results presented in the following sections are based on simulations of physical processes and cyber attacks, CAIA can be deployed in real installations where, in general, all the required parameter values are available in the form of historical and real-time data.

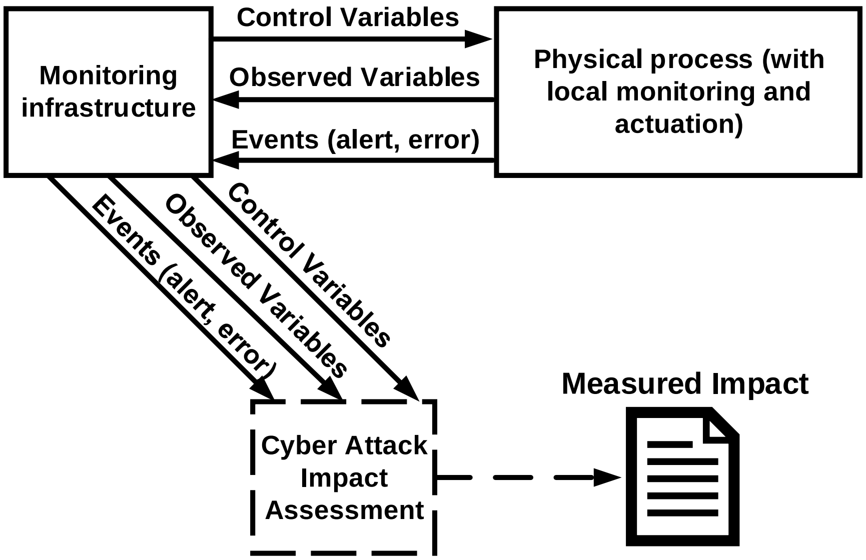

Figure 3: CAIA methodology.

3.3. Formal description

The CAIA methodology relies on measurements made at a set of discrete time instants $T = \{1, 2, \ldots, t, \ldots, m\}$. $J = \{1, 2, \ldots, j, \ldots, n\}$ is the set of observed variables and $I = \{1, 2, \ldots, i, \ldots, k\}$ is the set of control variables. Furthermore, Y^i $(i \in I)$ is a two-dimensional $m \times n$ matrix containing m measurements of n observed variables corresponding to an intervention applied via the control variable u_i. Y^i_{tj} and Y^0_{tj} denote the t^{th} measurements of the j^{th} observed variable for a scenario involving a specific intervention and for a scenario without intervention, respectively.

CAIA compares the values of the observed variables before and after the execution of a specific intervention. The mean value of the j^{th} observed variable for interventions on the i^{th} control variable is defined by:

$$\bar{Y}^i_j = \frac{1}{m} \sum_{t \in T} Y^i_{tj} \quad \forall i \in I, \forall j \in J \tag{1}$$

The intervention-free mean value for the j^{th} observed variable is defined by:

$$\bar{Y}^0_j = \frac{1}{m} \sum_{t \in T} Y^0_{tj} \quad \forall j \in J \tag{2}$$

Next, the cross covariance between the j^{th} intervention-free observed variable and the j^{th} observed variable with intervention on the i^{th} control variable is computed as:

$$C_{ij} = \frac{1}{m} \sum_{t \in T} \frac{(Y_{tj}^i - \bar{Y}_j^i)(Y_{tj}^0 - \bar{Y}_j^0)}{\bar{Y}_j^0 \bar{Y}_j^i} \quad \forall i \in I, \forall j \in J \tag{3}$$

The C_{ij} values ($\forall i \in I$, $\forall j \in J$) comprise the matrix of cross covariances, also called the impact matrix. This is an intermediate, yet detailed, result of the impact assessment procedure. As indicated by Equation (3), the rows of the impact matrix correspond to interventions on control variables while the columns correspond to observed variables. This two-dimensional structure expresses the impacts of cyber attacks on the local and global scales. Thus, the impact matrix can provide useful insights into the impacts of attack propagation on neighboring assets and the propagation of disturbances to remote assets.

The impact matrix also provides useful information in the case of observed variables that are equally significant. However, in some domains, such as the chemical sector, minor variations of certain products may be more relevant than major variations of other products. As a result, the impact assessment is further refined using a weighted impact matrix:

$$C_{ij}^w = \omega_j C_{ij} \quad \forall i \in I, \forall j \in J \tag{4}$$

where ω_j is the weight associated with the j^{th} observed variable. Obviously, $\omega_j = 1$ ($\forall j \in J$) when the observed variables are equally significant.

While the weighted impact matrix may provide adequate details required for impact assessment, specific scenarios (e.g., risk assessments) may require numerical estimates of the significance of each control variable and observed variable. To address this, two approaches are proposed for quantifying the impacts of cyber attacks on different assets and ranking the assets.

The first approach relies on the impact matrix to provide an impact ranking of cyber attacks on the i^{th} control variable:

$$R_i^c = \frac{\sum\limits_{j \in J} C_{ij}^w}{\sum\limits_{l \in I} \sum\limits_{j \in J} C_{lj}^w} \quad \forall i \in I \tag{5}$$

where R_i^c quantifies the sensitivity of the control variable to cyber attacks by aggregating the impact propagation for all observed variables.

The second approach is geared for other scenarios that require the identification of observed variables that are sensitive to cyber attacks that target different control variables. In such cases, the following equation, which quantifies the impact ranking of the j^{th} observed variable from the perspective of cyber attacks targeting control variables, is used:

$$R_j^o = \frac{\sum\limits_{i \in I} C_{ij}^w}{\sum\limits_{l \in I} \sum\limits_{j \in J} C_{lj}^w} \quad \forall j \in J \tag{6}$$

3.4. Assessment interval

As mentioned above, an attack impact assessment is triggered by system-level events. The end of the assessment interval may be specified using a time window whose length may be known *a priori* or may be determined dynamically.

In the case of a physical process whose the system response time γ is known *a priori*, the window length is fixed and is, therefore, predefined. However, γ alone may not cover the entire interval required for the process to reach steady state. Therefore, a second term is added to γ to accommodate the deviations of the observed variable values from their mean observed values. The total length of the time window, denoted by τ, is given by:

$$\tau = \gamma + \max(t - t_e || Y_{tj}^i - \bar{Y}_j^0| \leq \delta_j) \quad t, t_e \in T,\, j \in J \tag{7}$$

where t is the current measurement time for the j^{th} observed variable, t_e is the time at which the start event was recorded and δ_j is the maximum tolerable mean deviation of the j^{th} observed variable. CAIA uses Equation (7) to record the behavior of the physical process at least for time γ and additionally for the maximum time needed for the deviations of all the observed variables to fall below a tolerable value δ_j.

Figure 4 shows an example of the time window computation. In this case, the system response time γ is approximately 1.8 seconds. However, the assessment stops only after the deviation decreases below a specified threshold.

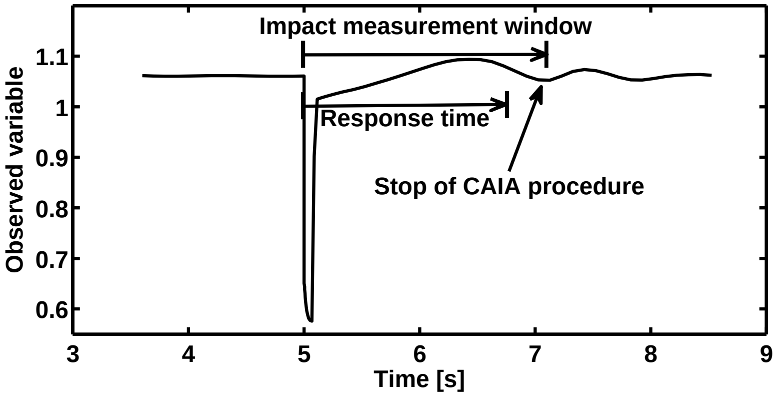

Figure 4: Computation of the assessment interval time window.

4. Experimental results

The CAIA methodology is evaluated from several perspectives. First, the effects of various parameter changes on the CAIA results are evaluated using the IEEE 14-bus electric grid model [23]. Following this, the scalability of CAIA is demonstrated by conducting experiments with the IEEE 300-bus electric grid model [18]. Next, the applicability of CAIA to other domains is showcased using a complex chemical process system [7]. Comparisons are also made between CAIA and graph-theoretic and electric centrality metric approaches [1, 29]. Finally, a novel stealthy cyber attack scenario that leverages the CAIA results is evaluated using the IEEE 14-bus model.

A CAIA prototype was implemented in MATLAB. The IEEE 14-bus and IEEE 300-bus electric grid models were simulated with the MATLAB PSAT toolbox [22]. The Tennessee Eastman chemical process [7] was simulated using MATLAB Simulink.

4.1. Parameter evaluation

The evaluation of the effects of the CAIA parameter values on the assessment results used the IEEE 14-bus model [23]. The model includes five generators (total generated power of 825 MVA), eleven loads, three transformers, fourteen buses and twenty branches. The IEEE 14-bus model was enhanced with control loops specific to real-world power systems such as power system stabilizers, automatic voltage regulators and turbine governors. Furthermore, secondary voltage regulators, including cluster controllers and central area controllers, were integrated in the IEEE 14-bus model to provide a realistic

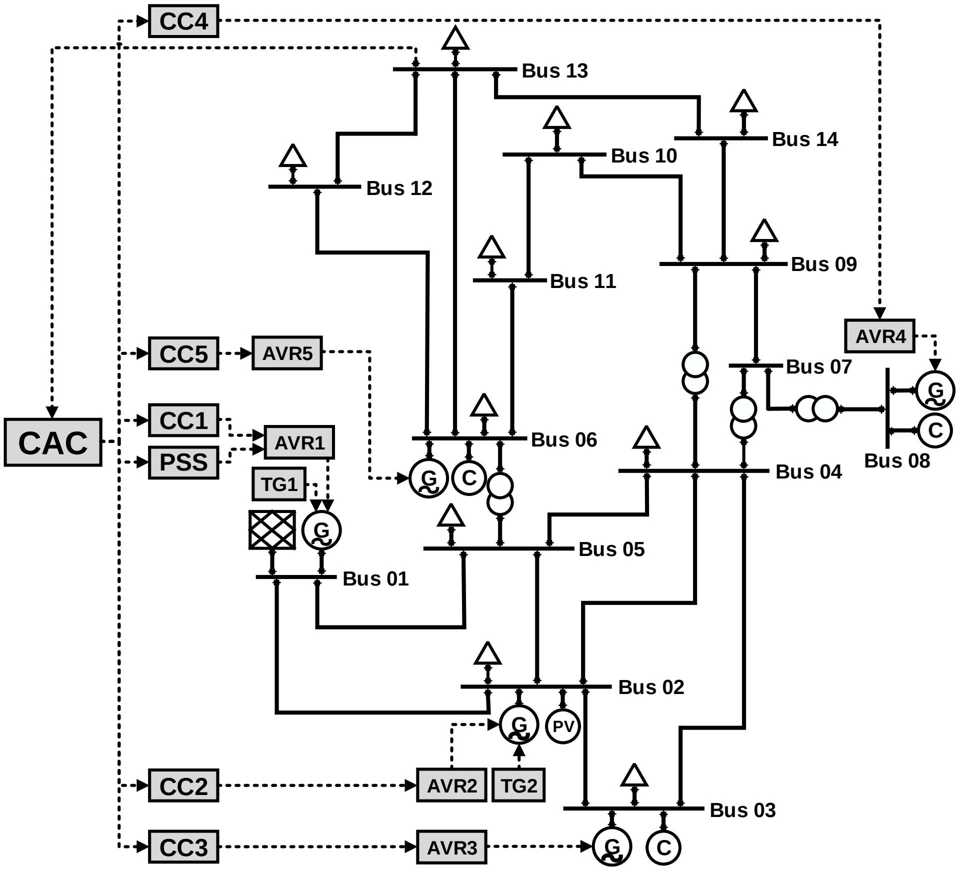

Figure 5: IEEE 14-bus model and its associated controllers.

control-loop-rich environment. Figure 5 shows a graphical representation of the IEEE 14-bus model. Note that the control loops are represented using dashed lines.

4.1.1. Stop condition parameters

The evaluation of the effects of various time window values on the stop condition leveraged the impact rankings of control variables R_i^c. The evaluation was performed by implementing a bus fault attack on each individual bus line. As shown in Figures 6(a) and 6(b), the measured impact on each control variable stabilizes after 0.15 seconds. Thus, the CAIA evaluation assumed a time window of 0.15 seconds for the IEEE 14-bus model as well as for the models used in the other experiments.

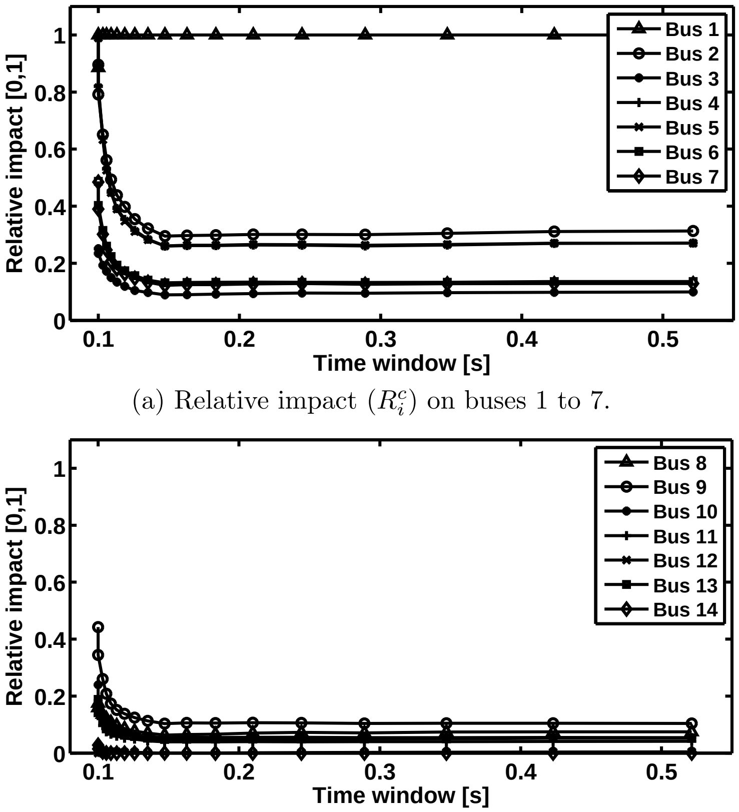

(a) Relative impact (R_i^c) on buses 1 to 7.

(b) Relative impact (R_i^c) on buses 8 to 14.

Figure 6: Effect of time window interval size on the impact values for the IEEE 14-bus model.

The results indicate major differences in the impact rankings of different control variables (bus-level circuit breakers). Control variable 1 has the highest ranking while the remaining circuit breaker control variables have significantly lower impact rankings. This is because substation 1 has a high power generator equipped with active control devices, which greatly influences the evolution of voltage levels throughout the electric grid. From the perspective of traditional risk assessment techniques [5, 8], this is an important finding because such a classification of control variables can help focus protection efforts on critical assets.

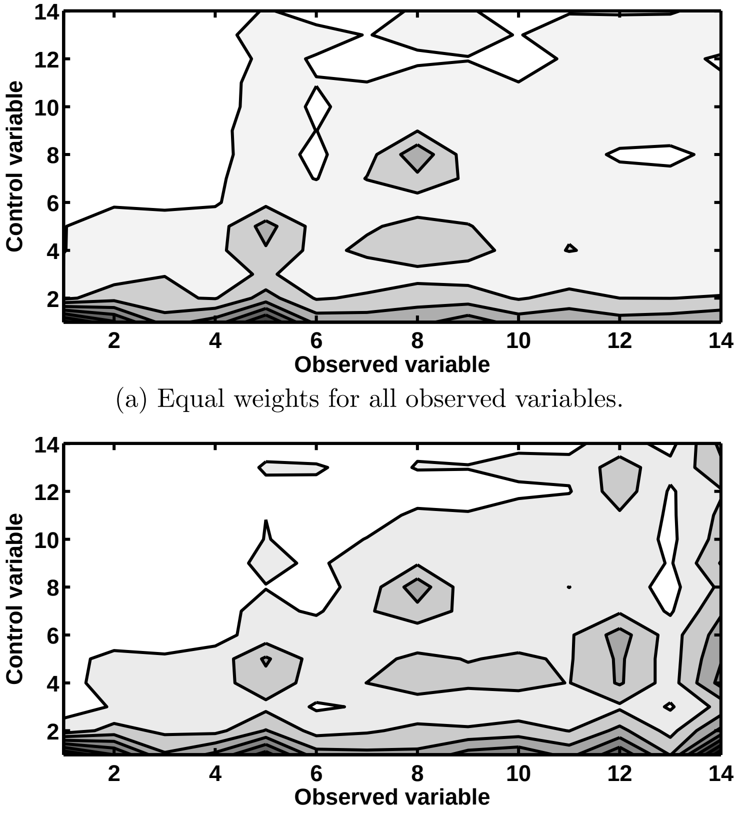

(a) Equal weights for all observed variables.

(b) Increased weights for observed variables (bus line voltage levels) 10, 12 and 14.

Figure 7: Effect of observed variable weights on the impact matrix for the IEEE 14-bus model.

4.1.2. Weighted impact matrix parameters

The heterogeneous characteristic of critical infrastructure components may require the use of different weights when computing the CAIA impact matrices. For example, in the case of electric grids, the observed variables in a specific substation (e.g., voltage levels) may be configured with larger weights to ensure that their significance is properly represented in the impact matrix. A cyber attack that causes substation-level bus faults in the IEEE 14-bus model is used to illustrate the effects of weighted observed variables on the impact rankings. Note that the observed variables in this case correspond to the substation voltage levels.

First, CAIA was used to derive the impact matrix with identical weights

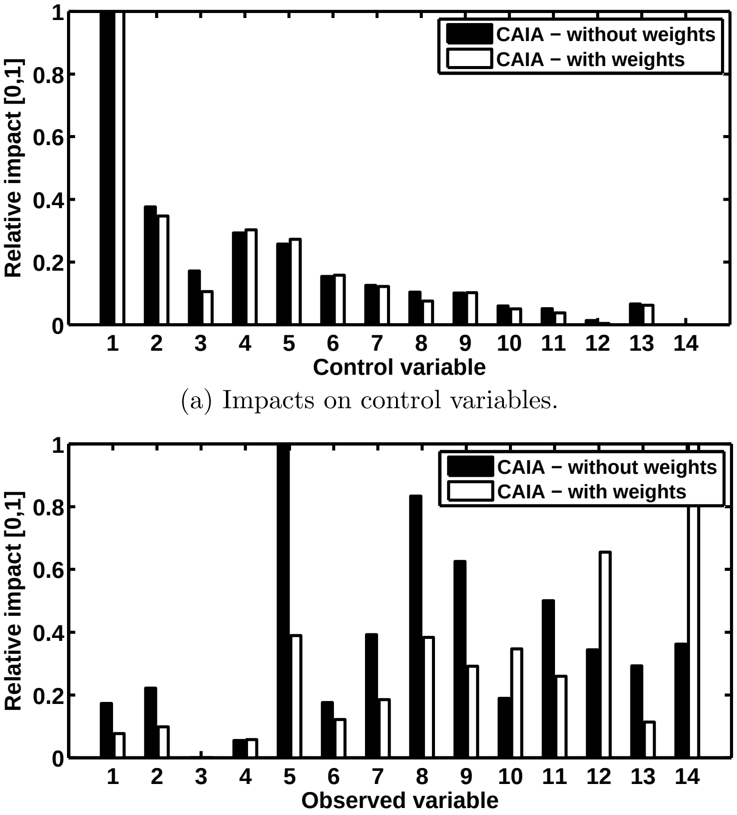

(a) Impacts on control variables.

(b) Impacts on observed variables.

Figure 8: Effect of observed variable weights on the impact ranking for the IEEE 14-bus model.

for the observed variables. As already indicated by previous results and further confirmed by the impact matrix in Figure 7(a), bus 1 is the most significant because of the attached generator and controllers. These results are also underscored by the impact rankings of the control variables (R_i^c) shown in Figure 8(a). However, the impact rankings of the observed variables (R_j^o) in Figure 8(b) are different because the voltage levels on buses 5, 8 and 9 are highly susceptible to cyber attacks that target control variables.

Next, the voltage levels (i.e., observed variables) for substations 10, 12 and 14 were assigned weights of 1.5, 2.0 and 2.3, respectively, while the remaining observed variables had the same initial weights of 1.0. The results in Figure 7(b) show an increase in the significance of attacks that target these observed variables. However, since the weights were assigned to variations in the values of the observed variables, the impact rankings for the control variables (R_i^c) do not exhibit significant changes (see Figure 8(a)). As expected, the increased weights assigned to observed variables 10, 12 and 14 yield significant increases in their impact rankings (see Figure 8(b)). These results demonstrate the importance of tuning CAIA parameters. Indeed, such adjustments better reflect the unique aspects of an analyzed infrastructure, increasing the accuracy of the results.

4.1.3. Attack parameters

As confirmed by previous research [19], different attacks can lead to profoundly different outcomes and impact results. Therefore, it is imperative to further clarify the applicability of the CAIA methodology with respect to different attack parameters. To accomplish this, the CAIA output is evaluated from two perspectives. First, the effect of different attack durations on the CAIA impact rankings is evaluated. Next, the effect of different control variable changes (i.e., attack magnitudes) on the CAIA results is evaluated Both the evaluations use a bias attack to intentionally change the input voltages of automatic voltage regulators by a certain percentage (automatic voltage regulators control the functioning of generators). The inputs include a set-point given by a central controller and a measured voltage, which may be measured locally or received from a remote phasor measurement unit. The attack scenario assumes that an attacker can modify the values sent by phasor measurement units and inject the altered measurement packets as automatic voltage regulator inputs.

More formally, let $T_A = \{t_s, \ldots, t_e\}$ be the attack duration, where t_s is the attack start time and t_e is the attack end time such that $T_A \subseteq T$. Furthermore, let u_i^t denote the value of the i^{th} control variable at time $t \in T$. Then, the bias attack on u_i^t is defined by:

$$\widetilde{u}_i^t = \begin{cases} u_i^t, & \text{for } t \notin T_A \\ u_i^t + \dfrac{\alpha}{100} u_i^t & \text{for } t \in T_A \end{cases} \tag{8}$$

where $\widetilde{u}_i^t$ is the attacker's modified value of the i^{th} control variable at time t and α is the percentage of bias intervention on u_i^t.

Figure 9(a) shows the effects of cyber attack duration on the CAIA results. Attacks with durations of less than one second may lead to different impact

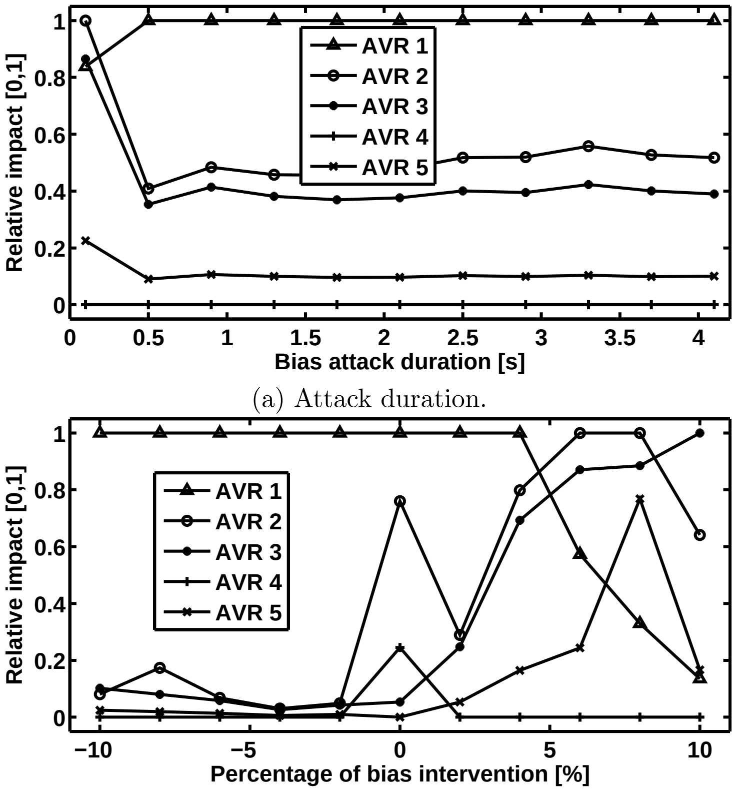

(a) Attack duration.

(b) Attack magnitude.

Figure 9: Effect of attack parameters on the impact rankings for the IEEE 14-bus model.

values. However, when the attack duration is increased, the uncompromised automatic voltage regulators can stabilize the process and the impact rankings are unchanged. Of course, simultaneously attacking several automatic voltage regulator inputs would result in radically different impact rankings. Nevertheless, CAIA provides a systematic approach for individually assessing the significance of the control and observed variables, which can then be used to identify groups of critical assets for which special protection mechanisms could be implemented.

Figure 9(b) shows the effects of the bias attack magnitude on the CAIA results. Note that the impact of the attack is highly dependent on the attack parameters and control loop configuration. In particular, the CAIA results are mostly influenced by the automatic voltage regulator configurations. Specifically, AVR1 is highly susceptible to cyber attacks and that its compromise may have severe repercussions on power grid stability. This susceptibility is due to the fact that AVR1 is configured to have a 200 p.u. gain value, which is ten times larger than the gain values of the other automatic voltage regulators.

4.2. Scalability evaluation

This section evaluates CAIA scalability using the large-scale IEEE 300-bus electric grid model [18], which incorporates 300 substations, 69 generators and 411 branches. Two attacks are implemented. The first attack triggers bus faults on the substation lines by changing the control variables that open and close bus line circuit breakers. The second attack causes severe load loss (10% of each attached load) by manipulating the control variables associated with the dynamic load management hardware. In both cases, the observed variables correspond to the substation voltage levels.

Both the attacks yield impact matrices with sparse distributions of impact values. Figure 10(a) shows the impact of the bus fault attack on each control variable and the propagation of perturbations to neighboring and remote substations. As expected, the (sparse) diagonal indicates that an attack on a bus control variable significantly impacts the associated substation.

However, disturbances also propagate to remote substations. In particular, the observed variables starting from bus 270 are highly sensitive to the majority of attacks that target different bus line control variables. The main reason is that, compared with the vast majority of other substations where the voltage levels range between 66 kV and 345 kV, these are low-voltage substations with voltage levels ranging from 0.6 kV to 20 kV. Consequently, attacks on bus control variables associated with these substations have low impacts on the high-voltage substations.

The computation of the relative impact values based on the impact matrix provide numerical estimates of the attack impact. As discussed above, the relative impact may be computed from the perspective of the control variables (R_i^c) or from the perspective of the observed variables (R_j^o). Figure 11(a) presents the results from the control variable perspective. Due to their high-voltage profiles, the substations up to substation 270 dominate the impact ranking. However, the impact ranking is profoundly different when viewed from the observed variable perspective (Figure 11(b)). Specifically, the impact ranking is dominated by observed variables starting from

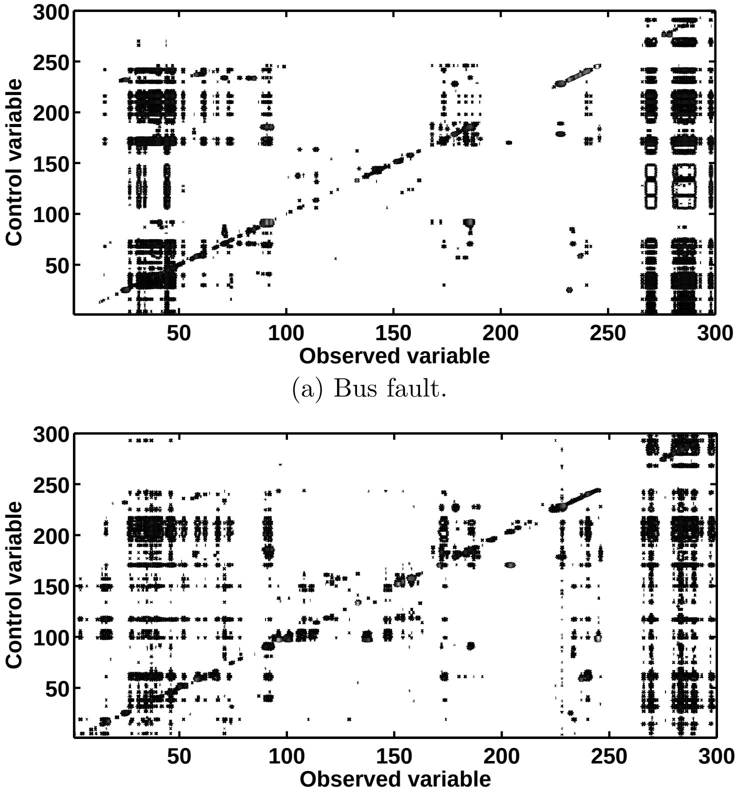

(a) Bus fault.

(b) Loss of load.

Figure 10: Impact matrices for the IEEE 300-bus electricity grid model for two attack scenarios.

substation 270 (i.e., low-voltage substations), which are highly sensitive to disturbances originating from remote high-voltage substations.

Similar results are obtained for the second attack as shown in Figure 10(b). However, the relative impact values have a different distribution. Figure 12(a) shows that the impacts on the control variables are highly dependent on the magnitudes of the decoupled loads. According to the model, substations 97, 99 and 100 have significant real and reactive power demands: substation 97 has 14.1 MW of real demand and 650 MVAr of reactive demand, substation 99 has 777 MW of real demand and 215 MVAr of reactive demand while substation 100 has 535 MW of real demand and 55 MVAr of reactive demand. Upon applying the 10% load loss due to the attack, the disturbances caused to these substations lead to severe voltage fluctuations with the peak impact

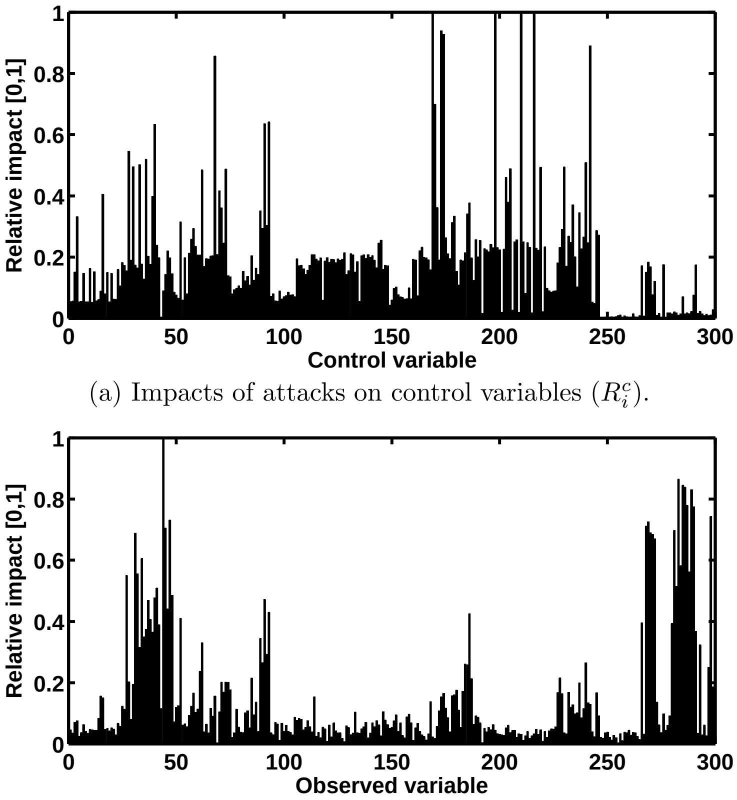

(a) Impacts of attacks on control variables (R_i^c).

(b) Impacts of attacks on observed variables (R_j^o).

Figure 11: Relative impact values in the bus fault attack scenario for the IEEE 300-bus model.

values shown in Figure 12(a).

Figure 12(b), which presents the impacts on the observed variables (R_i^o), shows that the low-voltage substations (from 270 to 300) are also sensitive to load loss. The results are best explained by the intrinsic power system components and connections for which disturbance propagation is governed by the electrical properties of the physical process. Other researchers (e.g., [29] have also confirmed this observation and the fact that, in the particular case of power systems, electrical properties may be inferred from the network admittance matrix. However, as demonstrated later in this paper, while these assumptions may hold for electrical power flow analysis, significant limitations are imposed on dynamic analyses and scenarios where complex control loops govern power grid behavior.

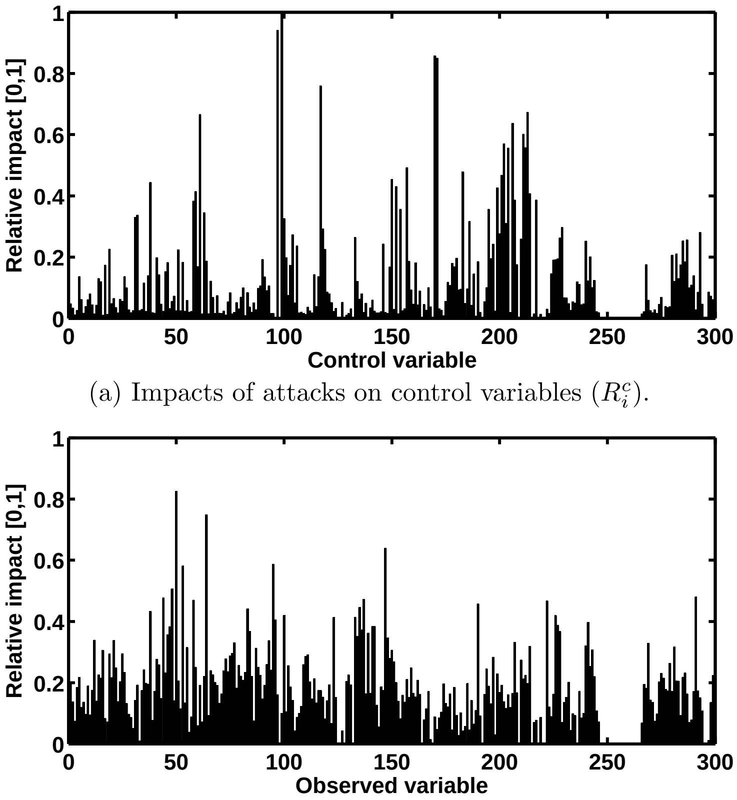

(a) Impacts of attacks on control variables (R_i^c).

(b) Impacts of attacks on observed variables (R_j^o).

Figure 12: Relative impact values in the loss of load attack scenario for the IEEE 300-bus model.

4.3. Cross-sector applicability

This section demonstrates CAIA's cross-sector applicability using the complex Tennessee Eastman chemical plant model [7]. Figure 13 shows the Tennessee Eastman chemical plant model, which has five main units: reactor, condenser, stripper, vapor/liquid separator and compressor. Each unit requires an automated control loop (dashed line) to ensure stable operation. The model incorporates twelve control (input) variables (denoted by MV), 41 observed (output) variables (denoted by Y) and 50 internal states. Several control approaches have been proposed for the Tennessee Eastman chemical plant model; the evaluation presented in this paper uses the decentralized control system of Ricker [24]. The evaluation involved executing an integrity attack that modifies the control variable values sent to the actuators that

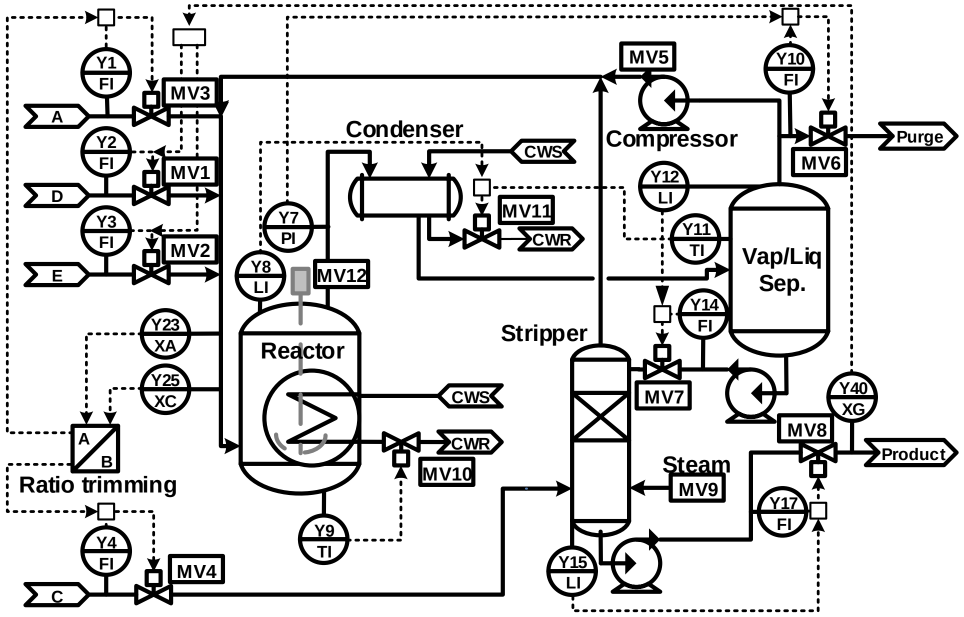

Figure 13: Tennessee Eastman chemical plant model.

manipulate the physical process.

In the first case, equal weights were assigned to the observed variables. The impact matrix in Figure 14(a) shows that significant impacts occur on the variables that participate in control loops (Y1–Y4, Y7–Y12, Y14, Y15, Y17, Y23, Y25 and Y40) whereas only minor impacts occur on the remaining observed variables.

In the second case, the weights used in [7] to calculate the total production costs were employed. Specifically, weights in the range 0.0318 to 30.44 were assigned to thirteen observed variables; the remaining variables were assigned weights of one. Figure 14(b) shows that assigning different weights to the observed variables leads to significant increases in the impacts on the observed variables for attacks on the majority of control variables. This behavior is shown by the vertical lines in Figure 14(b), which indicate profound changes in the significance of the observed variables with increased weights and their corresponding influence on the CAIA results. Using different assignments of weights customizes the CAIA methodology to different infrastructures, yielding more accurate impact matrices and more meaningful results.

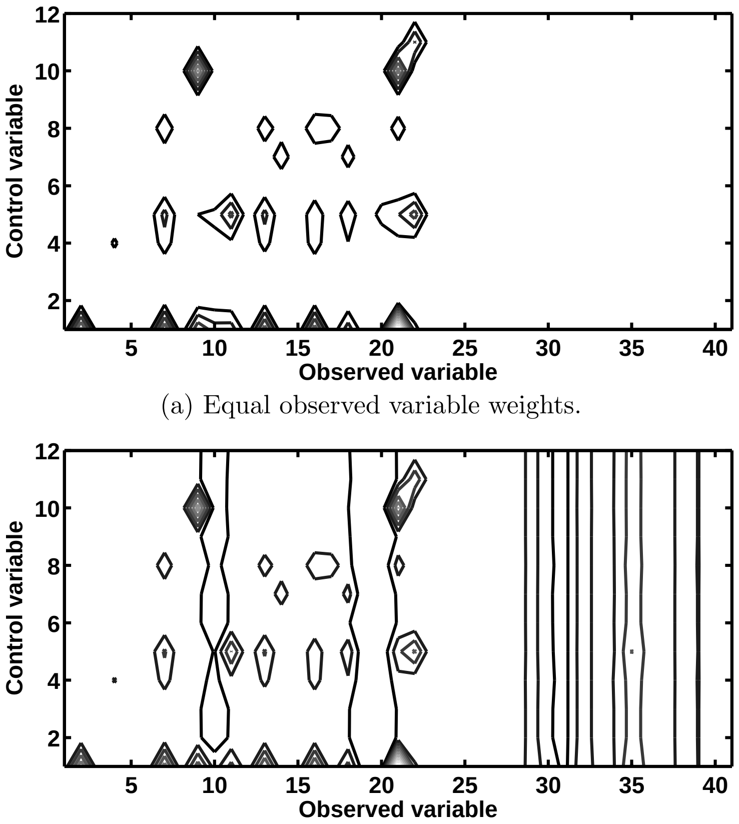

(a) Equal observed variable weights.

(b) Different observed variable weights.

Figure 14: CAIA impact matrix for the Tennessee Eastman chemical process model.

4.4. Comparison with other approaches

Numerous approaches have been proposed for assessing the vulnerabilities in electric power grids. The most common approaches focus on structural analysis based on centrality measures from graph theory. Bilis et al. [1] have developed a heuristic methodology that relies on five centrality metrics from graph theory to identify the most critical nodes in an electric power grid. The five metrics, degree centrality, eccentricity, betweenness centrality, centroid centrality and radiality, quantify the significance of nodes from a purely structural (i.e., topological) perspective. However, Hines et al. [16] and other researchers [6, 29] have demonstrated that, while structural assessments may be used to design synthetic power grids [30], they are not well suited to reproduce the behavior of real electric power grids. Structural abstractions

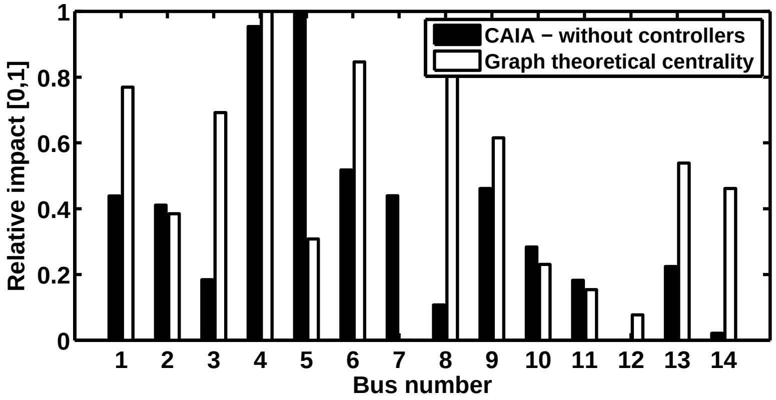

Figure 15: Comparison of the CAIA and graph-theoretic methodology results.

convey only the topological characteristics of electric grids while grid behavior is governed by physics (e.g., Kirchhoff's laws). Therefore, it is important to develop methodologies that capture the underlying physics of electric grids such as that of Wang et al. [29], which is based on electrical centrality metrics.

In the following, the CAIA results are compared with those obtained using the graph-theoretic and electrical centrality metric approaches. The evaluation is conducted using the IEEE 14-bus model. Since the graph-theoretic and electrical centrality metric approaches focus on assessments of the significance of nodes (i.e., bus lines), the CAIA evaluation incorporates a cyber attack that targets bus lines by triggering individual bus faults.

First, the CAIA results are compared with those obtained using the graph-theoretic methodology of Bilis et al. [1]. The heuristic algorithm and the five centrality measures described in [1] were implemented. Since the approach of Bilis et al. does not consider control loops, the CAIA methodology was applied to the controller-free IEEE 14-bus model.

Figure 15 compares the results of the two methodologies. Since the approach of Bilis et al. [1] is exclusively based on node connectivity, nodes that are more connected have greater impacts. However, the CAIA results and those presented in the previous sections demonstrate that node significance is highly dependent on node role and on the attached devices (e.g., stability controllers) that have important roles in maintaining system stability.

Next, the CAIA results are compared with the results of the electrical eigenvector centrality metric methodology proposed by Wang et al. [29]. Unlike the graph-theoretic methodology of Bilis et al. [1], the electrical eigen-

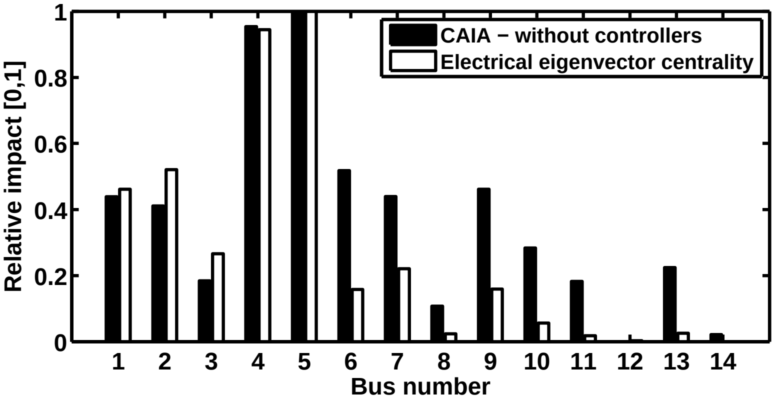

Figure 16: Comparison of the CAIA and electrical centrality metric methodology results.

vector centrality metric methodology provides a measure of the significance of a particular node in a network based on an adjacency matrix. In fact, the eigenvector centrality metric methodology builds on the electric admittance matrix, which embodies the electrical properties of the physical process. The electrical eigenvector centrality metric methodology was implemented as described in [29] and the results were compared with the CAIA results for the controller-free IEEE 14-bus model. Figure 16 shows that significant similarities exist in the rankings obtained using the two methodologies. The differences are due to the fact that CAIA accounts for the output of the entire electric grid model in a dynamic simulation while the electrical eigenvector centrality metric methodology relies exclusively on the admittance matrix, which does not consider the time dimension.

At this point, it is important to highlight the key advantages delivered by the CAIA methodology. Many elements are omitted when a methodology focuses purely on the structural properties or on the physical properties of a critical infrastructure. This is significant because the physical dimension of a critical infrastructure is tightly linked to the cyber domain via control loops, which can profoundly affect the behavior of the physical processes. Therefore, CAIA is capable of assessing not only controller-free physical processes, but also the more complex and ultimately more realistic critical infrastructure itself, in which the underlying processes are governed by a hierarchical structure of automated and human control loops. Figure 17 illustrates the unique capabilities of CAIA by comparing the results with and without control loops for the IEEE 14-bus model in the presence of a bus fault attack on

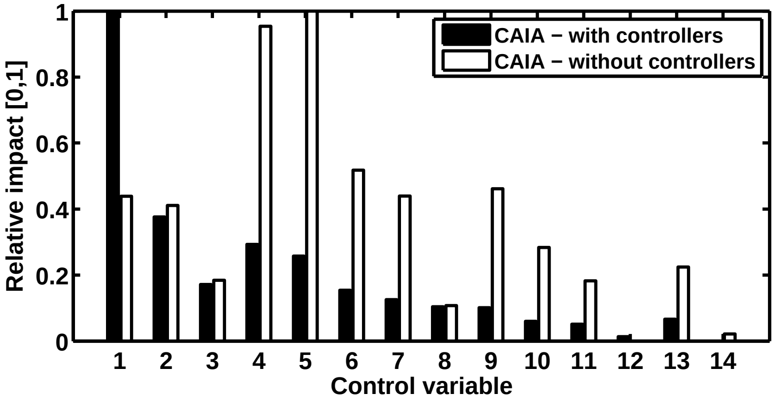

Figure 17: Comparison of CAIA results with and without control loops for the IEEE 14-bus model.

bus line control variables.

In the controller-free scenario, buses 4 and 5 have the highest relative impact values, which rank them as the most significant in the IEEE 14-bus model (Figure 18(a) where the measured impacts are represented using gray circles). This is explained by the positions of the buses in the electric grid and by the fact that they are responsible for transferring the power produced by the generators connected to buses 1, 2 and 3 to the upper (i.e., "Nordic") region. Since the power produced by the generators connected to buses 1, 2 and 3 is 94% (775 MVA) of the total generated power, their disconnection (via buses 4 and 5) leads to a voltage collapse in the upper region.

As stated above, CAIA can also be used to evaluate the impacts of cyber attacks in the presence of control loops. Figure 18(b) shows the CAIA results obtained for the IEEE 14-bus model with a power system stabilizer, automatic voltage regulators, turbine governors, cluster controllers and central area controllers enabled. In this case, the distribution of the bus-level impacts on the power grid is vastly different from when no control loops exist. As illustrated in Figure 17, with the stabilizer and control devices in place, CAIA indicates that buses 1 and 2 are the most critical to the normal functioning of the electric grid. This is because buses 1 and 2 are connected to high power generators equipped with active control devices that greatly influence the evolution of voltage levels throughout the grid (Figure 18(b)). These results can be used to strengthen the security of communications and control loops involving the associated control variables. Furthermore, the CAIA results may be used to drive advanced network planning and risk assessment methodologies such as those proposed by [5, 13].

4.5. Stealthy cyber attack

As already discussed, the cyber attack impact assessment methodology proposed in this work has applications ranging from risk assessment to control network design. However, the methodology also has (less obvious) uses to an attacker. For example, a malicious actor could use CAIA to identify the low-impact control variables. A complex large-scale cyber attack could then be created in which multiple low-impact variables are affected simultaneously to cause severe infrastructure degradation. The most significant aspect of this scenario is that affecting low-impact variables enables an attacker to maintain stealth, hinder detection by process-aware anomaly detection systems and increase the number of compromised cyber assets.

In order to attain his goals, an attacker must have access to measurements taken by sensor devices. Additionally, the attacker must be able to trigger changes to control variable values and record the effects of these changes. Obviously, acquiring such access to critical resources may be difficult. Nevertheless, it is reasonable to assume that the attacker could use laboratory-scale testing facilities as well as simulation models to gain significant knowledge of the targeted infrastructure. After constructing the CAIA impact matrix, the attacker could carefully plan and execute the attack against the real infrastructure.

The IEEE 14-bus model is used to illustrate the effects of a stealthy attack that is designed based on CAIA results. The scenario assumes that the attacker seeks to compromise the maximum number of line breakers (each breaker controls the connectivity between two substations). Figure 19(a) shows the impact matrix computed by the attacker. Next, the attacker has to identify the low-impact breaker variables. This information is shown in Figure 19(b): the attacker has to compromise line breaker 7, followed by breakers 3, 18, 8, 1 and 13, and, finally, breaker 15. Compromising the remaining breakers could lead to more significant impacts that may be visible to operators and/or trigger alarms by anomaly detection systems.

Figure 20 shows the stealthy attack that compromises each breaker and disconnects each substation line in sequence. The results demonstrate that the attacker's actions are stealthy until the eighth circuit breaker is compromised. This is confirmed in the operator's view of the monitored substation voltage levels (Figure 21), where the attacker's actions on the first seven breakers (executed during the zero to six second time interval) lead to minor voltage changes that are within the normal operating limits. However, after the eighth breaker is compromised (after seven seconds), the voltage levels begin to collapse. Such an event generally triggers automated circuit breaking and alarms that enable the operators to isolate and mitigate the effects of the cyber attack. However, by this time, the attacker may already control a large number of power grid assets, enabling him to severely degrade the voltage levels on a large scale.

To further illustrate the application of CAIA to identify low-impact assets when creating stealthy cyber attacks, the effects of two random cyber attacks on the IEEE 14-bus model line breakers are demonstrated. Figure 22 shows that an operator can detect both attacks more rapidly than the stealthy attack discussed above. In the case of the first random attack, the voltage collapse is visible after only two seconds, and after four seconds for the second random attack. These results underscore the fact that malicious entities must plan attacks on critical infrastructure assets very carefully to ensure their successful execution. Other researchers (e.g., [19]) confirm that attackers must be aware of the complexities of the cyber and the physical dimensions of critical infrastructures in order to maximize attack impact.

It is important to emphasize that the successful execution of the stealthy attack described in this section would require significant resources. The attacker could perform most of the CAIA computations in advance based on physical process models and laboratory-scale experiments. However, the attack ultimately has to be launched against a real infrastructure, and for maximal impact, the attacker would need to experiment with the actual target or a close approximation of the target. To defend against complex, stealthy attacks, security personnel could use CAIA to identify the critical cyber assets. These assets could be positioned in critical security zones with special protection mechanisms to reduce the attack risk.

5. Conclusions

The CAIA methodology, which is inspired by systems dynamics, is designed to assess the impacts of cyber attacks on control variables, which are responsible for the correct functioning of physical processes in critical infrastructures. The methodology has been evaluated from several perspectives using a variety of simulated physical processes (IEEE 14-bus, IEEE 300-bus and Tennessee Eastman models). The results demonstrate the effectiveness and broad applicability of CAIA in assessing the impacts of cyber attacks on critical infrastructures. CAIA is superior to graph-theoretic and electrical centrality metric approaches because it supports effective dynamic behavioral analyses of critical infrastructures as well as analyses of critical infrastructures comprising various cyber and physical components, including critical control loops; furthermore, CAIA is applicable in a range of critical infrastructure sectors. Future research will focus on evaluating the applicability of CAIA to production systems and on integrating CAIA results in control network design methodologies.

Acknowledgement

This research was supported by a Marie Curie FP7 Integration Grant under the 7th European Union Framework Programme.

IMPORTANT NOTE TO IJCIP TYPESETTERS: I have checked and edited the references in this paper myself. Please DO NOT MODIFY the references – except to add hyperlinks. Please contact the Journal Manager Ms. Ramya Vasudevan if you have any questions.
Professor Sujeet Shenoi, Editor-in-Chief, IJCIP

References

[1] E. Bilis, W. Kroger and C. Nan, Performance of electric power systems under physical malicious attacks, *IEEE Systems Journal*, vol. 7(4), pp. 854–865, 2013.

[2] J. Brandt, Electric grid facing security threats from all sides, *Smart Grid News* (`www.smartgridnews.com/story/electric-grid-facing-security-threats-all-sides/2014-09-03`), September 3, 2014.

[3] A. Cardenas, S. Amin, Z. Lin, Y. Huang, C. Huang and S. Sastry, Attacks against process control systems: Risk assessment, detection and response, *Proceedings of the Sixth ACM Symposium on Information, Computer and Communications Security*, pp. 355–366, 2011.

[4] T. Chen and S. Abu-Nimeh, Lessons from Stuxnet, *IEEE Computer*, vol. 44(4), pp. 91–93, 2011.

[5] G. Correa-Henao, J. Yusta and R. Lacal-Arantegui, Using interconnected risk maps to assess the threats faced by electricity infrastructures, *International Journal of Critical Infrastructure Protection*, vol. 6(3-4), pp. 197–216, 2013.

[6] E. Cotilla-Sanchez, P. Hines, C. Barrows and S. Blumsack, Comparing the topological and electrical structure of the North American electric power infrastructure, *IEEE Systems Journal*, vol. 6(4), pp. 616–626, 2012.

[7] J. Downs and E. Vogel, A plant-wide industrial process control problem, *Computers and Chemical Engineering*, vol. 17(3), pp. 245–255, 1993.

[8] R. Filippini and A. Silva, A modeling framework for the resilience analysis of networked systems-of-systems based on functional dependencies, *Reliability Engineering and Systems Safety*, vol. 125, pp. 82–91, 2014.

[9] D. Ford, A behavioral approach to feedback loop dominance analysis, *System Dynamics Review*, vol. 15(1), pp. 3–36, 1999.

[10] J. Forrester, Counterintuitive behavior of social systems, *Theory and Decision*, vol. 2(2), pp. 109–140, 1971.

[11] B. Galloway and G. Hancke, Introduction to industrial control networks, *IEEE Communications Surveys and Tutorials*, vol. 15(2), pp. 860–880, 2013.

[12] B. Genge and C. Siaterlis, Analysis of the effects of distributed denial-of-service attacks on MPLS networks, *International Journal of Critical Infrastructure Protection*, vol. 6(2), pp. 87–95, 2013.

[13] B. Genge and C. Siaterlis, Physical process resilience-aware network design for SCADA systems, *Computers and Electrical Engineering*, vol. 40(1), pp. 142–157, 2014.

[14] A. Giani, R. Bent and F. Pan, Phasor measurement unit selection for unobservable electric power data integrity attack detection, *International Journal of Critical Infrastructure Protection*, vol. 7(3), pp. 155–164, 2014.

[15] M. Hagerott, Stuxnet and the vital role of critical infrastructure operators and engineers, *International Journal of Critical Infrastructure Protection*, vol. 7(4), pp. 244–246, 2014.

[16] P. Hines, E. Cotilla-Sanchez and S. Blumsack, Do topological models provide good information about electricity infrastructure vulnerability? *Chaos*, vol. 20(3), article no. 033122, 2010.

[17] J. Huang, E. Howley and J. Duggan, The Ford Method: A sensitivity analysis approach, *Proceedings of the Twenty-Seventh International Conference of the System Dynamics Society*, 2009.

[18] IEEE Test Systems Task Force, 300 Bus Power Flow Test Case, Department of Electrical Engineering, University of Washington, Seattle, Washington (`www.ee.washington.edu/research/pstca/pf300/pg\_tca300bus.htm`), 1993.

[19] M. Krotofil, A. Cardenas, J. Larsen and D. Gollmann, Vulnerabilities of cyber-physical systems to stale data – Determining the optimal time to launch attacks, *International Journal of Critical Infrastructure Protection*, vol. 7(4), pp. 213–232, 2014.

[20] D. Kundur, X. Feng, S. Liu, T. Zourntos and K. Butler-Purry, Towards a framework for cyber attack impact analysis of the electric smart grid, *Proceedings of the First IEEE International Conference on Smart Grid Communications*, pp. 244–249, 2010.

[21] H. MacKenzie, How Dragonfly hackers and RAT malware threaten ICS security, *Belden Industrial Security Blog* (`www.belden.com/blog/industrialsecurity/How-Dragonfly-Hackers-and-RAT-Malware-Threaten-ICS-Security.cfm`), September 15, 2014.

[22] F. Milano, An open source power system analysis toolbox, *IEEE Transactions on Power Systems*, vol. 20(3), pp. 1199–1206, 2005.

[23] F. Milano and M. Anghel, Impact of time delays on power system stability, *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 59(4), pp. 889–900, 2012.

[24] N. Ricker, Decentralized control of the Tennessee Eastman Challenge Process, *Journal of Process Control*, vol. 6(4), pp. 205–221, 1996.

[25] K. Sgouras, A. Birda and D. Labridis, Cyber attack impact on critical smart grid infrastructures, *Proceedings of the Innovative Smart Grid Technologies Conference*, 2014.

[26] S. Sridhar and M. Govindarasu, Model-based attack detection and mitigation for automatic generation control, *IEEE Transactions on Smart Grid*, vol. 5(2), pp. 580–591, 2014.

[27] K. Stouffer, S. Lightman, V. Pillitteri, M. Abrams and A. Hahn, Guide to Industrial Control Systems (ICS) Security, NIST Special Publication 800-82, Revision 2 Final Public Draft, National Institute of Standards and Technology, Gaithersburg, Maryland, 2015.

[28] Symantec, Dragonfly: Cyberespionage Attacks Against Energy Suppliers, version 1.21, Mountain View, California, 2014.

[29] Z. Wang, A. Scaglione and R. Thomas, Electrical centrality measures for electric power grid vulnerability analysis, *Proceedings of the Forty-Ninth IEEE Conference on Decision and Control*, pp. 5792–5797, 2010.

[30] Z. Wang, A. Scaglione and R. Thomas, Generating statistically correct random topologies for testing smart grid communications and control networks, *IEEE Transactions on Smart Grid*, vol. 1(1), pp. 28–39, 2010.

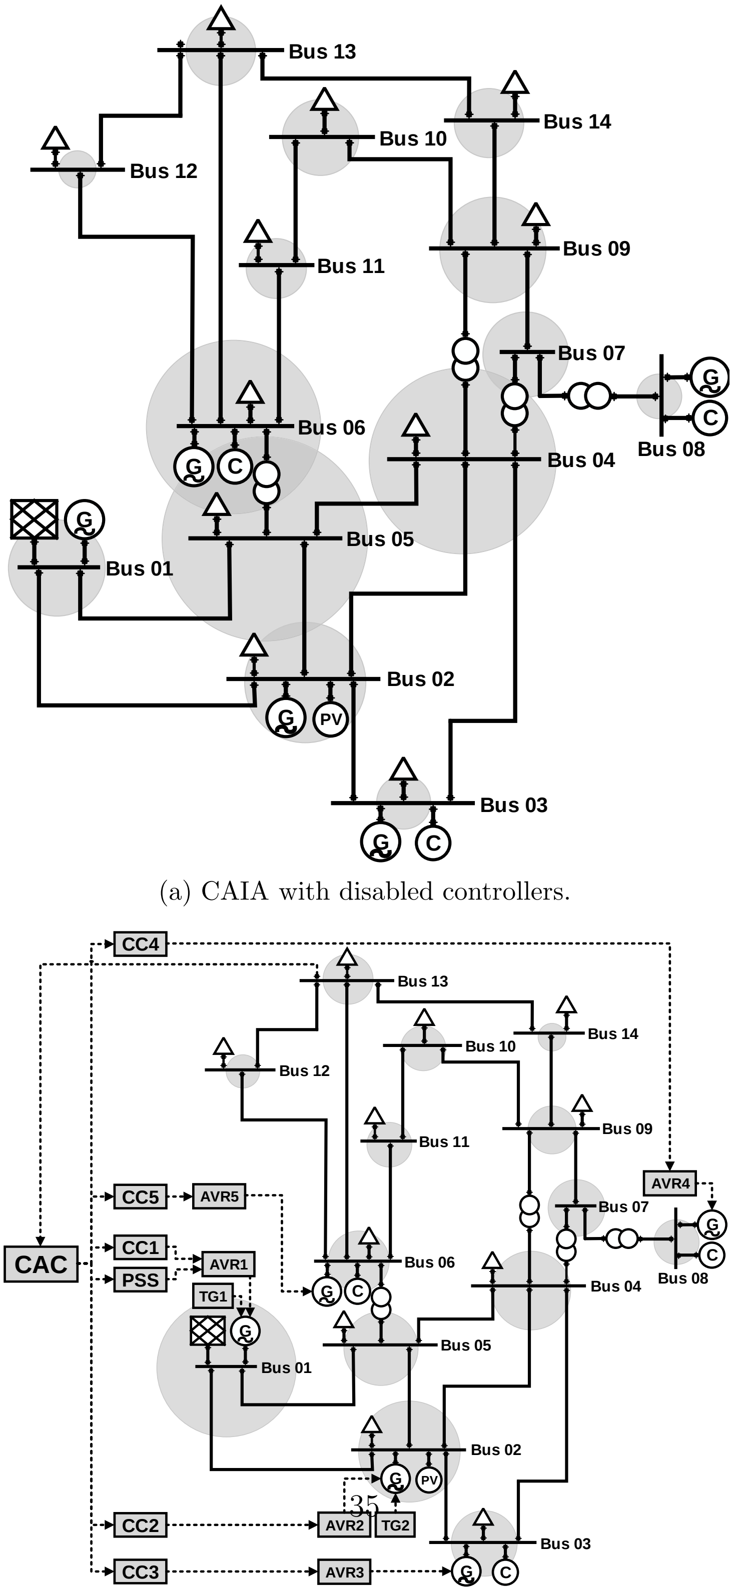

(a) CAIA with disabled controllers.

(b) CAIA with enabled controllers.

Figure 18: Structural view of the CAIA results for the IEEE 14-bus model.

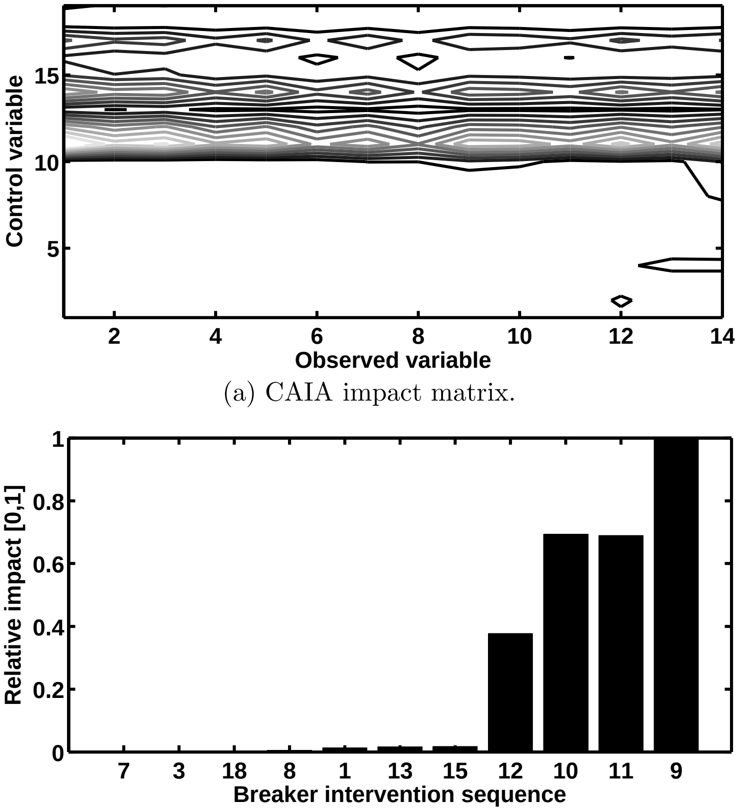

(a) CAIA impact matrix.

(b) Ordered impact ranking of breakers.

Figure 19: Stealthy cyber attack on the IEEE 14-bus model line breakers.

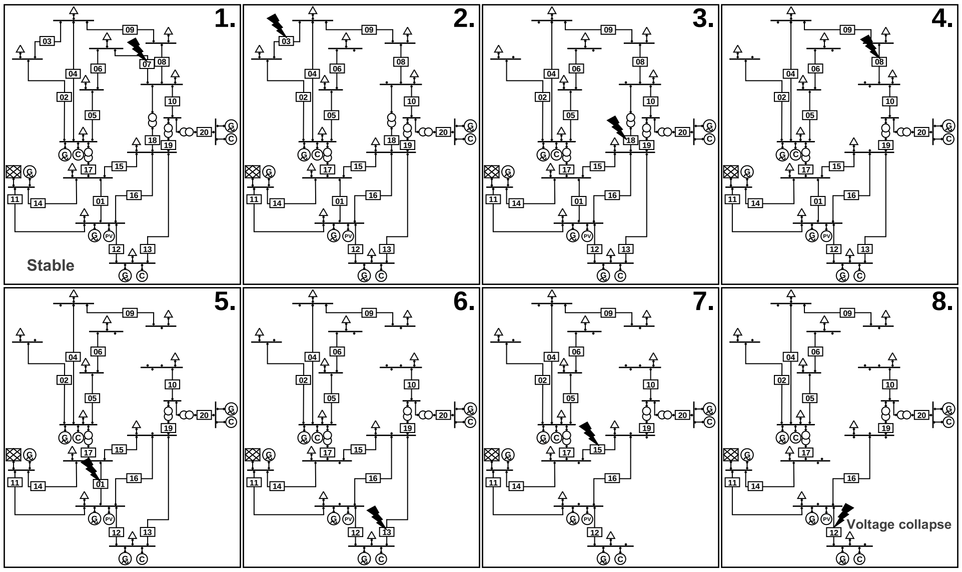

Figure 20: Stealthy cyber attack sequence that disconnects substation lines.

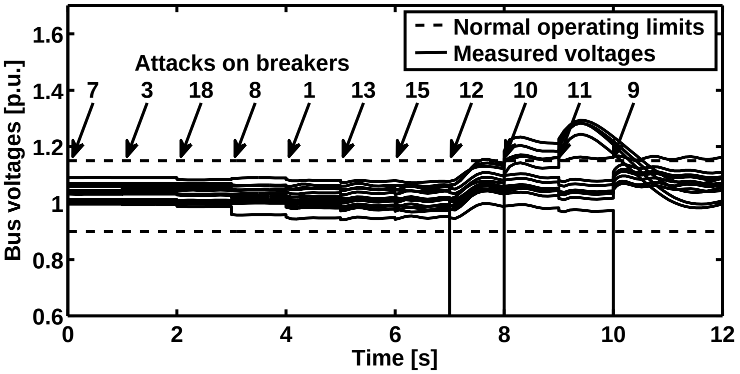

Figure 21: Operator's view of the stealthy attack sequence.

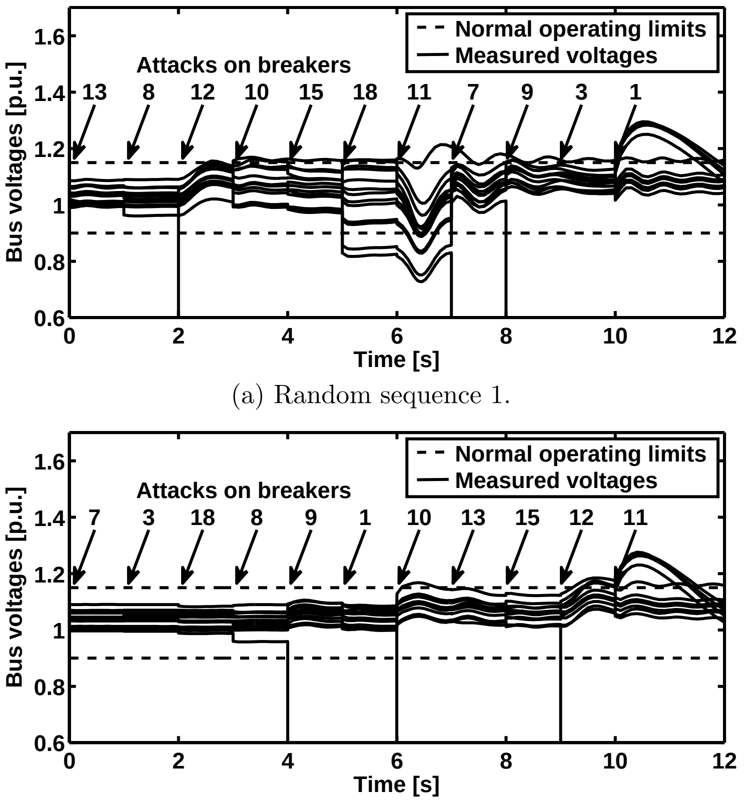

(a) Random sequence 1.

(b) Random sequence 2.

Figure 22: Operator's view of two random cyber attacks.