DECENTRALIZED CYBER-PHYSICAL SECURITY APPLICATIONS FOR THE FUTURE GRID ENERGY MANAGEMENT SYSTEM

A Dissertation
Presented to
The Academic Faculty

by

Leilei Xiong

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Electrical and Computer Engineering

Georgia Institute of Technology
May 2019

DECENTRALIZED CYBER-PHYSICAL SECURITY APPLICATIONS FOR THE FUTURE GRID ENERGY MANAGEMENT SYSTEM

Approved by:

Dr. Santiago Grijalva, Advisor
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Lukas Graber
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Raheem Beyah
School of Electrical and Computer Engineering
Georgia Institute of Technology

Dr. Polo Chau
School of Computational Science and Engineering
Georgia Institute of Technology

Dr. Maryam Saeedifard
School of Electrical and Computer Engineering
Georgia Institute of Technology

Date Approved: December 14, 2018

TABLE OF CONTENTS

LIST OF TABLES

Page

LIST OF FIGURES

Page

SUMMARY

The objective of this work is to enhance the state estimator software application in the energy management system. The state estimator is responsible for processing raw power system measurements received from substation equipment in the field and filtering out the errors to determine the most likely state of the power system in terms of bus voltage magnitudes and angles. Because the state estimator solution serves as the input to other critical downstream applications in the energy management system, such as contingency analysis and optimal power flow, it is key that a unique solution can be found, and that the solution is accurate.

The changing nature of the power grid is introducing new challenges for the state estimator. The first challenge is the need to process more data in less time. New sensors such as phasor measurement units that have a finer sampling granularity are creating more data for the state estimator to process. Market deregulation has created the need to monitor larger interconnections. Also, faster system dynamics due to the integration of non-dispatchable renewables such as wind and solar are creating a need for faster state estimation, so that power system operators have a clearer understanding of the time-varying system behavior. With all of these new driving forces, it becomes increasingly more important to move away from the conventional central state estimator architecture towards a decentralized approach that is more scalable. However, it is unclear from the existing research literature just how much decentralization is ideal. To address this question, we investigate the impact of the number of sub-problems on the performance of the state estimator.

The second challenge is the need for more robust power system cybersecurity. Over the past decade, the power grid has become increasingly intertwined with information and communication technology, transforming it into a cyber-physical system ("smart grid"). While this interconnectivity comes with many benefits, it has also rendered the grid more vulnerable to cyber-attacks. Currently the main lines of defense against grid cyber-attacks are traditional IT cybersecurity measures, such as firewalls, air gaps, and antivirus software. While these measures are necessary, they are insufficient when used alone since they cannot detect cyber-physical attacks on the measurement data. The purpose of the state estimator is to filter out measurement errors, so it is a natural candidate for defending against data-related cyber-attacks. In this work, we present a cyber-physical security assessment co-simulator that uses an enhanced state estimator to identify whether a control command is malicious. We also introduce a novel approach for *N*-1 RTU cyber-physical security assessment that ranks RTUs by how critical the impact of their loss is on the state estimator.

1 INTRODUCTION

1.1 Dissertation Overview

The objective of this work is to enhance the state estimator software application in the energy management system to deal with two major challenges facing the power grid of the future. The first challenge is the need to process more measurement data in less time. Several major forces are driving this need. First, new sensors such as phasor measurement units that have a finer sampling granularity are creating more data for the state estimator to process. Next, market deregulation has created the need to monitor larger interconnections. Finally, faster system dynamics due to the integration of non-dispatchable renewables such as wind and solar are creating a need for faster state estimation, so that power system operators can have a clearer understanding of the time-varying system behavior. Conventional state estimators used by regional transmission organizations and utilities today have a central architecture, which does not scale well as the number of measurements increases. Hence, the existing research literature has proposed moving towards a decentralized approach that is more scalable. However, it is unclear from the literature just how much decentralization is ideal. To address this question, we investigate the impact of the number of sub-problems on the performance of the state estimator.

The second challenge is the need for more robust power system cybersecurity. Over the past decade, the power grid has become increasingly intertwined with information and communication technology, transforming it into a cyber-physical system known as the "smart grid." While this interconnectivity comes with many benefits, it also renders the grid more vulnerable to cyber-attacks. Currently the main lines of defense against grid cyber-attacks are traditional IT cybersecurity measures, such as firewalls, air gaps, and antivirus software. While these measures are necessary, they are insufficient

when used alone since they cannot detect cyber-physical attacks, such as attacks on the measurement data. The purpose of the state estimator is to filter out measurement errors, so it is a natural candidate for defending against cyber-physical attacks. In this work, we present a cyber-physical security assessment co-simulator that uses an enhanced state estimator to identify whether a control command is malicious. We also introduce a novel approach for *N*-1 RTU cyber-physical security assessment that ranks RTUs by how critical the impact of their loss is on the state estimator.

This dissertation is organized as follows. The remainder of Chapter 1 introduces how power systems work today, the architecture of the energy management system, and discusses in greater detail the challenges and needs for the state estimator. Chapter 2 introduces the three different functions of the state estimator (state estimation, network observability analysis, and bad data detection) and the mathematical formulations for each. Chapter 3 presents the novel work investigating the impact of the number of sub-problems on the performance of the state estimator. Chapter 4 reviews the state of power system cybersecurity today. Chapter 5 presents the novel work on cyber-physical security assessment. Chapter 6 concludes the dissertation and lists the contributions of this work.

1.2 Traditional Power Systems

The traditional power system has three different functions: generation, transmission, and distribution. Historically, large-scale power plants (typically fueled by non-renewable resources such as nuclear power, coal, and natural gas) were used to generate electricity, because they were more efficient and more cost-effective than smaller plants. These generators are located far away from load centers, so the power they produce is increased to a very high voltage by step-up transformers at a local substation before it is sent across a network of long transmission lines in order to reduce the overall resistive losses in the system. Once power reaches a substation near the end user, step-down transformers at that substation reduce the voltage to a safer and more

useable level. Finally radial distribution lines carry the lower-voltage power to electrical loads (residential, commercial, or industrial consumers of electricity). In order for the entire interconnected system to maintain stable operation, the total amount of generation must equal the total demand of the loads plus system losses at all times. A typical grid contains hundreds of generators, hundreds of transformers, thousands of circuit breakers that protect major equipment, as well as thousands of high voltage transmission lines and lower voltage distribution lines. Figure 1 is a simple illustration of the traditional power system architecture.

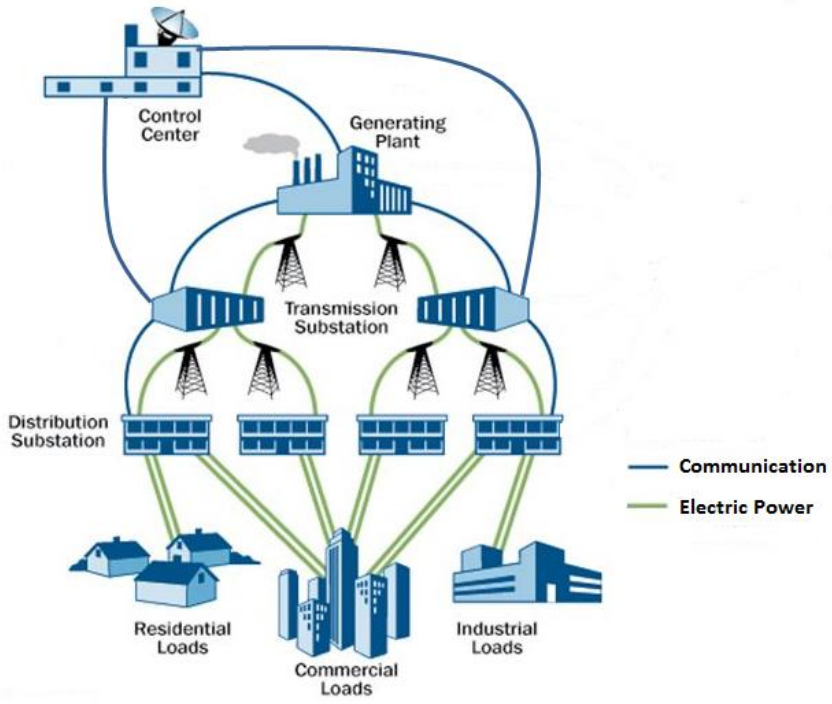

Figure 1. Traditional power system architecture

In the continental United States, three alternating-current (AC) power grids have evolved to mostly operate independently of one another with only limited transfer of electricity between them: the Eastern Interconnection, the Western Interconnection, and the Texas Interconnection. The Eastern and Western Interconnections are also tied to the Canadian power grid. These interconnections are made up of over 3,000 electric utilities, most of whom do not perform all three functions of generation, transmission, and

distribution. In regulated markets, utilities typically own generation, transmission, and distribution assets and thus perform all three functions. In deregulated markets, utilities tend to engage only in distribution while the transmission network is owned by companies and organizations that are obligated to provide indiscriminate access to various power producers.

In deregulated regions in the United States, the transmission network is controlled by five independent system operators (ISOs) and four regional transmission organizations (RTOs), as shown in Figure 3. These five ISOs are California ISO, Electric Reliability Council of Texas (ERCOT), ISO New England, Midcontinent ISO, and New York ISO. The four RTOs are ISO New England (also an ISO), Midcontinent ISO (also an ISO), PJM Interconnection, and Southwest Power Pool. ISOs operate a regional grid, administer the region's wholesale electricity market, and provide reliability planning for the region's bulk grid. RTOs perform the same functions as ISOs, but they have greater responsibility for the transmission network, such as coordinating, controlling, and monitoring the power system in their region. In regions where there is no ISO/RTO, utilities serve these same functions [1].

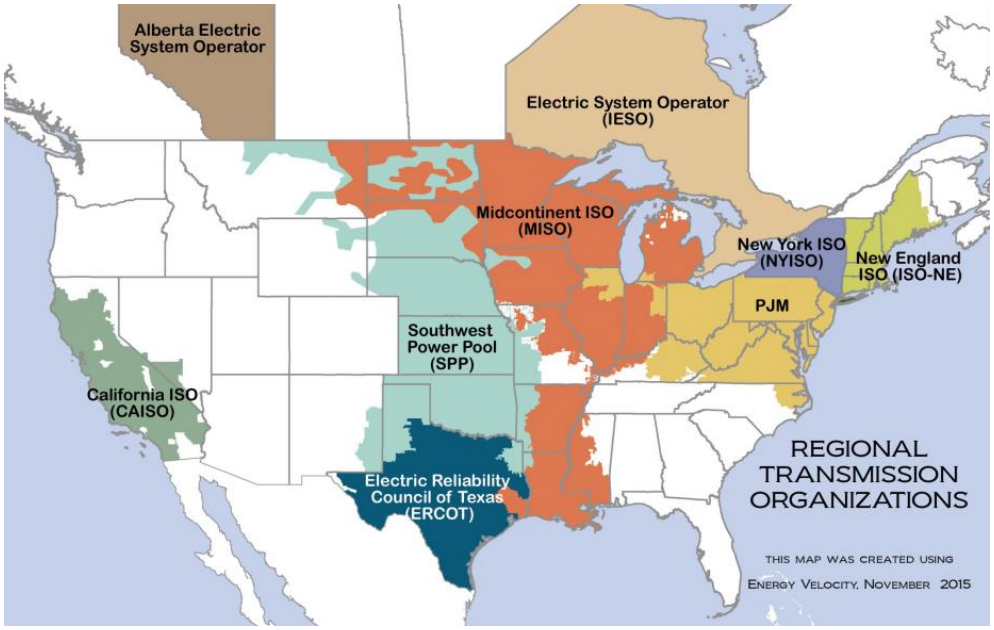

Figure 2. ISOs and RTOs in North America [2]

The actual monitoring and operation of the power grids is performed by entities known as balancing authorities, whose area of control is known as a balancing area or control area. Each balancing authority ensures in real time that the power balance in their area between supply and demand is maintained at all times, including system losses and transfers of electricity between control areas. ISOs and RTOs also serve as balancing authorities. The rest are electric utilities that have taken on the responsibility. In the continental United States, there are 66 balancing authorities in total, as shown in Figure 3. The Eastern Interconnection consists of 31 balancing authorities in the United States and 5 in Canada. The Western Interconnection consists of 34 balancing authorities in the United States, 2 in Canada, and 1 in Mexico. The Texas Interconnection, which covers most of Texas, consists of only a single balancing authority (ERCOT) [3].

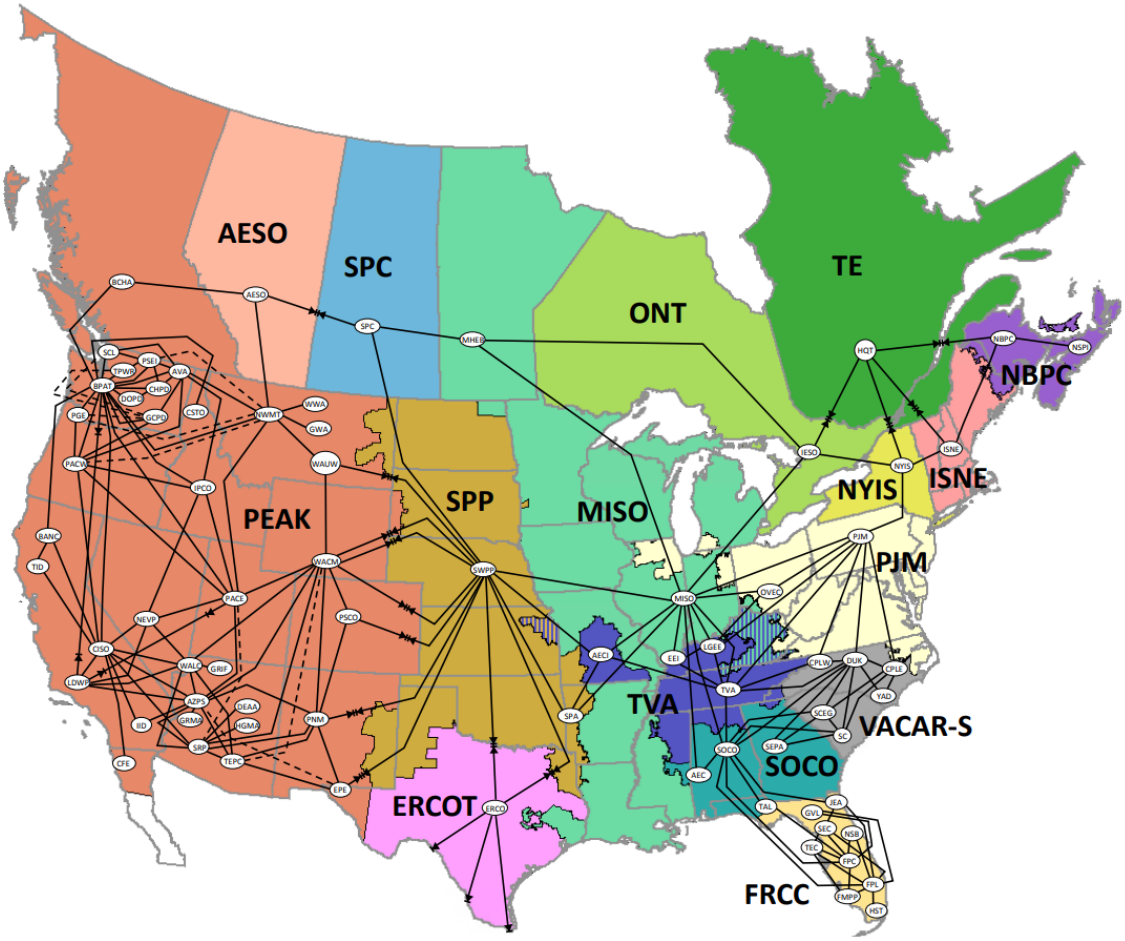

Figure 3. Balancing authorities in North America [4]

1.3 The Energy Control Center and the Energy Management System

Every balancing authority has at least two energy control centers (one primary and one backup), where power system operators monitor, control, and coordinate the physical and economic operation of the power system in that control area. Information about monitored equipment in the field (such as analog measurements and status indicators) is periodically communicated from remote terminal units (RTUs) located at various substations back to the supervisory control and data acquisition (SCADA) master station at the energy control center. RTUs are typically scanned at a rate of once every 2 seconds. This information is then used by the energy management system (EMS), a sophisticated suite of tools used by power system operators to monitor and control the operation of the transmission and bulk generation system.

A typical EMS is composed of many complex computer applications, as shown in Figure 4. Traditionally these applications can be divided into several overarching categories: SCADA/EMS monitoring, economic functions, operations planning, and security assessment. Not shown in Figure 4 is the presentation of EMS results to the operator in the form of a large-scale visualization. These tools combined provide the operator with an accurate picture of the current state of the system as well as offer guidelines on control actions that ensure its secure and economic operation. Currently the information management and decision making process at each control center is handled in a centralized manner, and the computations are typically performed serially.

The focus of the work in this dissertation is on the state estimator, which falls under the SCADA/EMS monitoring category. The state estimator includes three main functions: state estimation (determining the most likely state of the system), network observability analysis (determining if a unique state estimation solution can be found), and bad data detection (determining if measurements are consistent with one another based on the physics of the system).

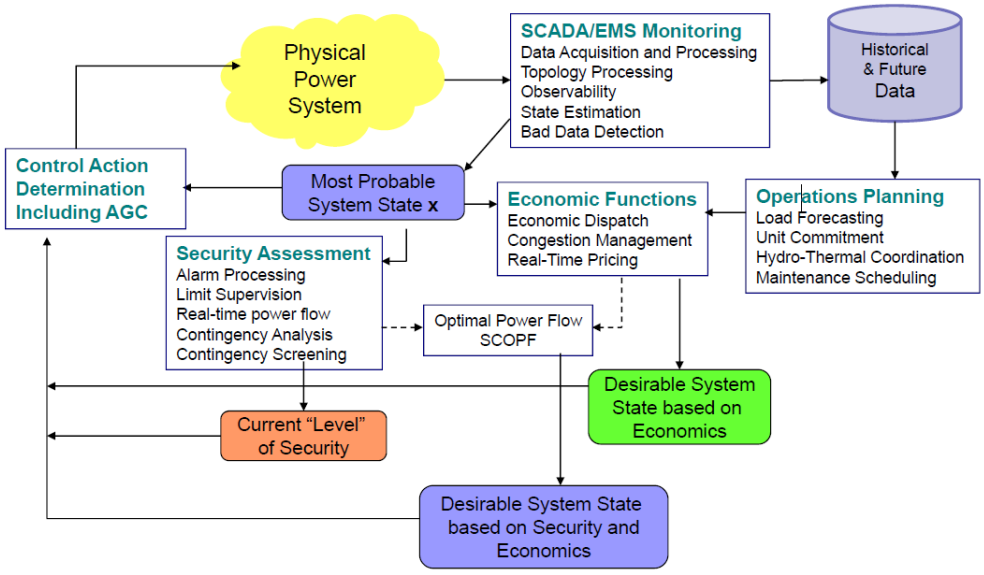

Figure 4. Architecture of an energy management system [5]

1.4 New Challenges Facing the Power Grid

The power grid is currently facing a number of new challenges (illustrated in Figure 5). The first challenge is more data than ever before. New technologies such as phasor measurement units (PMUs) are straining the grid's communication and data processing infrastructure with their higher measurement rate (typically 30 to 60 samples per second) as compared to the SCADA measurement rate (typically one sample every 2 seconds). Also, utilities need to share more information and monitor the system over larger geographic areas due to new market regulations and pricing competition. Additionally, the integration of renewables and energy storage, interest in updating the distribution system to facilitate demand response, and two-way power flow are driving the need for more accurate models and better supervision [6]. All of these driving forces are creating more information than ever before that need to be processed and analyzed by the EMS in a short period of time. To address this issue, there is interest in moving away from the traditional centralized architecture seen in power systems to a more distributed one. One alternative is centralized data acquisition but distributed computation across

multiple processors in the same computing cluster. Another option is completely separate data acquisition and computation. A distributed architecture offers several advantages, such as semi-autonomous operation for each region and faster computation due to smaller problem sizes.

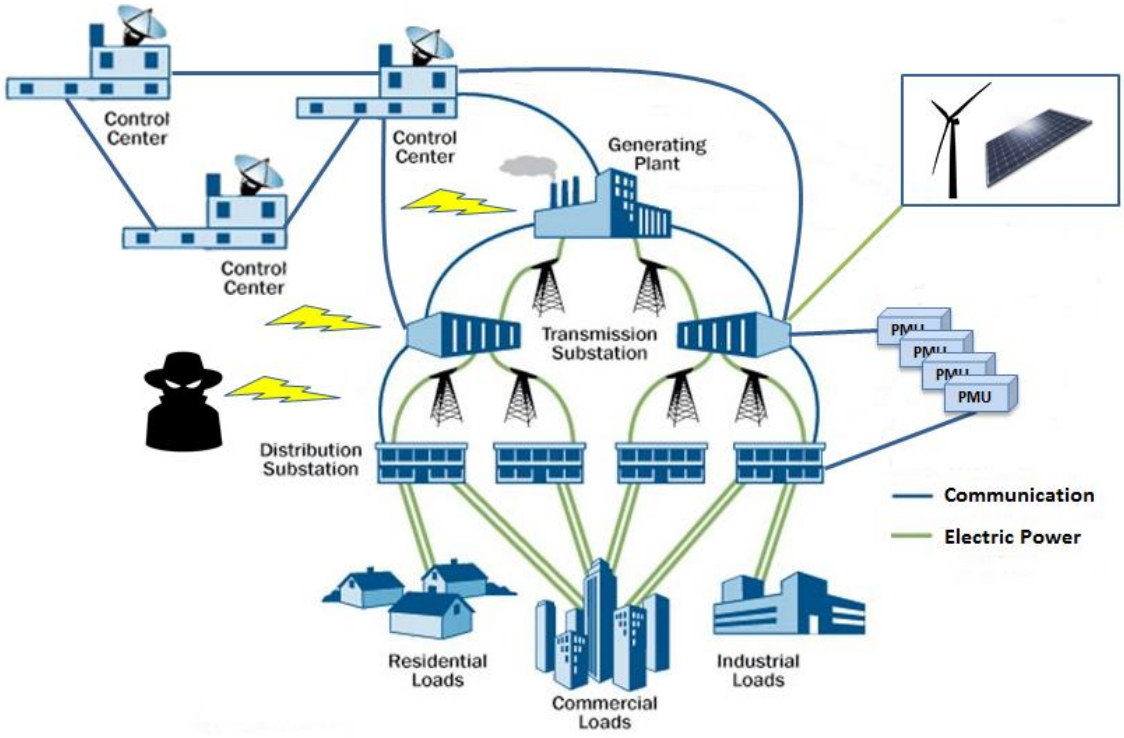

Figure 5. New challenges facing the power grid

The second challenge is the threat of cyber-attacks. As the smart grid incorporates more information and communication technology, more avenues of attack become open to malicious cyber actors, rendering the grid more vulnerable to cyber-attacks. While traditional cybersecurity measures such as firewalls, air gaps, and antivirus software play an important role in protecting the IT infrastructure of the grid, they are insufficient when used alone since the power grid is a cyber-physical system that has complex interdependencies between the physical electrical network and the communication network that supports it. While researchers have been aware of the potential impact of cyber-physical attacks on the grid since the mid-2000s, the idea had remained mostly

academic in nature. It was not until December 2015 that the first known successful power system cyber-attack occurred in Ukraine. Hackers were able to compromise information systems of three distribution companies and disrupted power to approximately 225,000 customers [7]. This event illustrated the potential severity and scale of impact that cyber-attacks could have on the grid.

1.4.1 Need for Enhanced State Estimation

Power system state estimation (SE) is a critical function in the EMS that determines the most likely state of the system (bus voltage magnitudes and angles) from raw unsynchronized field measurements, which are collected by substation RTUs and then sent to the EMS front end computer at the control center. These measurements include real and reactive power injections, transmission line flows, and bus voltage magnitudes. Because measurements are subject to instrumentation and communication errors, it is the role of the state estimator to detect and filter these errors. Power system operators rely on the solution of the state estimator to provide them with an accurate picture of the system's real-time behavior. The SE solution serves as the input to other downstream EMS applications, such as contingency analysis and optimal power flow. Currently SE is serially computed for a single control area.

The computational speed of the SE problem depends on the number of measurements and the number of states that need to be estimated. As the number of sensors in the power system increases, the size of the SE problem will also continue to grow. In order to solve this problem in the same amount of time as current operations or even less time, the global SE problem must be decomposed into a series of smaller problems. This approach can take the form of distributed state estimation (DSE), where the SE problem for a single ISO/RTO is rewritten as a series of smaller problems that are then solved in parallel. For DSE, the data acquisition is centralized, but the computation is distributed across multiple processors, albeit most likely using a cluster in the same

physical location. Another form is multi-area state estimation (MASE), where multiple control areas coordinate solving their joint SE problem. For MASE, both the data acquisition and computation are completely separate.

This work proposes to explore two questions that have been largely undiscussed in the literature for DSE and MASE: 1) how the global SE problem could be automatically decomposed into smaller problems, and 2) the impact of the number of sub-problems on the computational speed. In the case of DSE, the boundaries of the sub-problems have frequently been determined through trial and error, and the number of problems is assumed to be given [8]. In MASE, the boundaries for each sub-problem are defined by the physical boundaries of each control area, and the number of problems is defined by the number of control areas involved in the joint SE. By examining the impact of the number of sub-problems on the computational speed, we can begin to understand the benefits and limitations of decentralized SE.

1.4.2 Need for Cyber-Physical Security Assessment

Growing concerns over national security have brought intense scrutiny upon the vulnerability of power systems to cyber-attacks. Thus an increasingly critical direction of research has been to evaluate the potential impact of cyber-attacks on the electric grid. Traditionally power system security assessment only takes into account physical contingencies, such as the outage of a transmission line or generator. In the event of a cyber-attack, communication infrastructure could also be targeted. In general the grid was designed with implicit trust in its communication channels and system components, so it contains numerous cyber vulnerabilities. Recently there has been a push within the power industry to put cybersecurity measures in place and close some of these vulnerabilities. However, while traditional methods of cybersecurity protection are important and should continue to be implemented, they will not always be successful in preventing attackers

from penetrating the system. Thus it is important to understand the impact that various cyber-attacks could have on the grid.

One of the main purposes of the state estimator is to detect and correct erroneous measurements before they are propagated to downstream EMS applications, so it can be viewed as a natural first line of defense against cyber-attacks on the bulk grid. Erroneous measurements usually occur from instrumentation and/or communication failure, but threat actors could also maliciously manipulate measurements to intentionally mislead power system operators during a data-related cyber-attack. (Refer to Section 4.4.1 where the Stuxnet attackers intentionally fooled monitoring systems and human operators into believing that the centrifuges were operating normally when they were actually spinning themselves apart.) The intent could be either to fool human operators into directly making an incorrect control decision for the grid (for example, if operators believed that the grid was in a different state than it actually was) or to mask the situation when the grid is in an emergency state.

Another important purpose of the state estimator is to identify when the system becomes unobservable (refer to Section 2.3). Even under normal circumstances, the system can lose observability when enough critical measurements are lost due to instrumentation and/or communication failure. With sufficient knowledge of the system, cyber-attackers could seek to exploit this vulnerability by launching denial-of-service attacks on RTUs that transmit critical measurements. This could cause power system operators to lose situational awareness, and they would become either unable to detect that an attack is happening or unable to respond to further attacks.

Much of the existing literature regarding cyber-attacks on the grid has focused on studying stealth deception attacks, such as false data injection where arbitrary errors can be introduced into the estimated state by manipulating the power system measurements. In this work, we investigate two other types of attacks: 1) bad command injection attacks, and 2) denial-of-service attacks on remote terminal units. In bad command injection, an

unwanted command is sent to a substation RTU. In a denial-of-service attack on an RTU, communication between the attacked RTU and the control center is blocked, so measurements from that RTU are not available to the state estimator. By studying these two types of attacks, we can better understand and ultimately mitigate them.

2 POWER SYSTEM STATE ESTIMATION

Chapter 2 provides background information on power system state estimation. It contains five different sections. Section 2.1 introduces the state estimator as well as its role in power system control and operation. Section 2.2 discusses the main state estimation methods that are currently used in energy management systems today, including the formulations of each algorithm as well as their strengths and weaknesses. Section 2.3 presents existing methods for observability analysis. Section 2.4 presents existing methods for bad data detection.

2.1 Introduction

Power system state estimation is a critical function in the energy management system that determines the most likely state of the power system (bus voltages represented in phasor form as magnitudes and phase angles) from raw unsynchronized field measurements, which are collected by substation RTUs and then sent to the EMS front end computer at the control center. These SCADA measurements can include real and reactive power injections, transmission line flows, and bus voltage magnitudes. Because measurements are subject to instrumentation and communication errors, it is the role of the state estimator to detect and filter these errors. Power system operators rely on the solution of the state estimator to provide them with an accurate picture of the system's real-time behavior.

The output of the state estimator, the most likely system state, is the input to many other important time-sensitive EMS applications, such as security assessment functions like contingency analysis and economic functions like security-constrained optimal power flow (as seen in Figure 6). Currently the state estimator problem is solved once every half a minute to once every few minutes depending on regional regulations and practices. Table 1 shows the exact solution rate for each ISO/RTO in North America.

California ISO's state estimator and New York ISO's state estimator run automatically the most frequently at once every 30 seconds [9], [10] while ERCOT's state estimator runs the least frequently at once every 5 minutes [11], [12]. (The state estimator can run upon manual request in addition to automatically running.) The solution rate is directly related to how often other EMS applications such as security and market functions are automatically run. By contrast, SCADA measurements are collected at a rate of approximately one sample every 2-4 seconds. These measurements are typically not timestamped or synchronized. A key assumption behind the design of the state estimator is that all SCADA measurements are collected at the same point in time or close enough to the same point in time as to not make a difference in practical implementation.

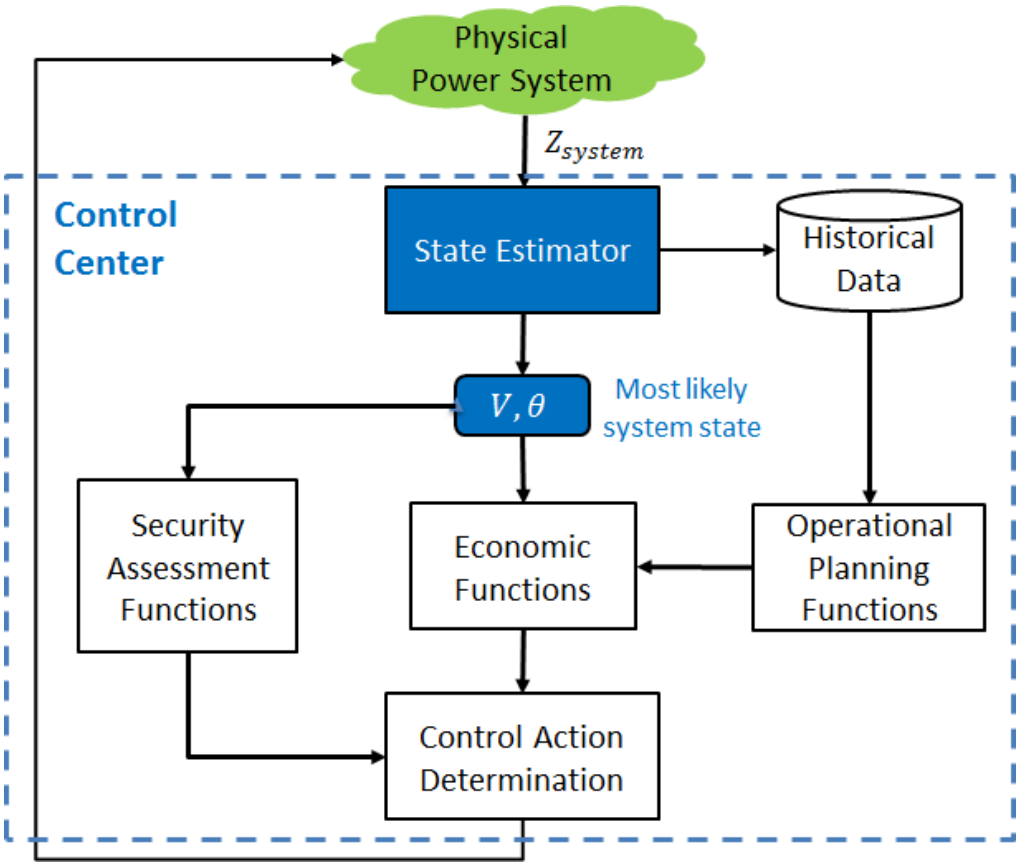

Figure 6. The importance of state estimation in the EMS

Table 1. SE Solution Rate in North America

ISO/RTO Name	SE Solution Rate [s]
California ISO [9]	30
ERCOT [11], [12]	300
ISO-NE [13]	180
Midwest ISO [14]	60
New York ISO [10]	30
Southwest Power Pool [15]	90

In addition to state estimation, the state estimator also includes three other functions: topology error processing, observability analysis, and bad data detection/processing. The topology error processor uses the state estimator inputs to determine the most likely physical configuration of the power system (which is outside the scope of this work). The observability analysis function determines whether a unique solution can be found during state estimation. The bad data detection function flags measurements and removes them if they are found to be inconsistent with other system measurements. The relationships between all four functions are illustrated in Figure 7. Observability analysis and bad data detection are discussed in greater detail in Sections 2.3 and 2.4 respectively.

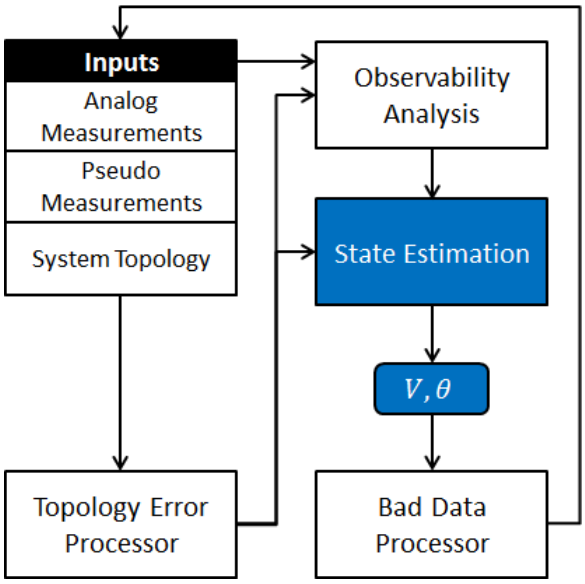

Figure 7. State estimator process

2.2 Solution Methods

There exist many solution methods for power system state estimation. The most well-known and widely used is the weighted least squares (WLS) state estimator, which was first introduced by F. C. Schweppe in 1970 [16], [17]. Assuming measurement errors can be modeled as independent random variables with Gaussian distributions that have zero mean and known variance, a minimization problem can be formulated in terms of the residuals of each measurement. There are three different versions of the WLS state estimator: AC, decoupled, and DC. The main computational burden for AC SE is the calculation and decomposition of the gain matrix. If the sensitivity of the real power equations to changes in the bus voltage magnitude and the sensitivity of the reactive power equations to changes in the bus angles are low, then the decoupled formulation can be used instead, in which the off-diagonal blocks of the measurement Jacobian are ignored in calculating the gain matrix. The main advantages that it has over the AC formulation are 1) it uses less memory, and 2) it is computationally faster due to the use of smaller and constant gain sub-matrices [18]. Further simplification leads to the DC formulation, where all bus voltage magnitudes are assumed to be 1.0 per unit.

The main disadvantage of the WLS state estimator is that when the Normal Equations are ill-conditioned, the WLS state estimator may become numerically unstable and fail to reach convergence. Ill-conditioning can arise from having a large number of injection measurements, very large weighting factors, and the presence of both short and long lines at the same bus. To combat the weakness of the Normal Equations, most methods avoid using an ill-conditioned gain matrix [19]. Also, the WLS state estimator does not deal well with outliers that might be bad data, leading to a need for robust state estimators (discussed briefly in Section 2.2.7) [20]. Despite these issues, the conventional WLS state estimator is still the most widely used algorithm in EMS systems, especially the decoupled and DC formulations. However, AC has become the preferred formulation

in power systems literature due to its higher level of accuracy and the cheap availability of modern computing power.

2.2.1 AC Weighted Least Squares

Let z represent a set of power system measurements. Then $z = h(x) + e$, where x is the estimated state vector (bus voltages and angles), h is the vector of functions relating the state variables to the error-free measurements, and e is a vector of measurement errors, which are assumed to have a Gaussian distribution with mean 0 and variance σ^2.

The WLS estimator minimizes the objective function

$$J(x) = [z - h(x)]^T R^{-1}[z - h(x)], \tag{1}$$

where R is a diagonal matrix of the measurement error variances. To obtain the minimum x, we take the partial derivative of the objective function

$$g\left(x^{(k)}\right) = \frac{\partial J(x)}{\partial x} = -H\left(x^{(k)}\right)^T R^{-1}\left(z - h\left(x^{(k)}\right)\right) \tag{2}$$

and set it equal to 0. Here $x^{(k)}$ is the state vector at iteration k. H is the measurement Jacobian equal to $\frac{\partial h(x)}{\partial x}$.

By applying the Gauss-Newton method [19], we obtain the Normal Equations

$$\left[G\left(x^{(k)}\right)\right]\Delta x^{(k+1)} = -g\left(x^{(k)}\right), \tag{3}$$

where the gain matrix G is

$$G\left(x^{(k)}\right) = \frac{\partial g(x^{(k)})}{\partial x} = H\left(x^{(k)}\right)^T R^{-1} H\left(x^{(k)}\right). \tag{4}$$

Then x is solved for iteratively until a convergence tolerance ε is reached. The expressions for h and H are presented in Section 2.2.2.

The exact algorithm is described in detail below:

Step 1. Set iteration index k equal to 0.

Step 2. Initialize the state vector $x^{(k)}$, typically as a flat start (voltage magnitudes are set to 1, and voltage angles are set to 0).

Step 3. Calculate gain matrix $G\left(x^{(k)}\right)$.

Step 4. Calculate $-g\left(x^{(k)}\right)$.

Step 5. Decompose $G\left(x^{(k)}\right)$ using Cholesky factorization and solve for $\Delta x^{(k+1)}$.

Step 6. Check if $\max\left|\Delta x^{(k)}\right| \leq \varepsilon$ (a user-defined stopping condition)? If yes, stop. If no, update $x^{(k+1)} = x^{(k)} + \Delta x^{(k)}$ and $k = k + 1$, then go back to Step 3.

2.2.2 Polar and Rectangular Coordinates

Similar to the power flow algorithm, AC state estimation can be formulated in either polar or rectangular coordinates. Polar coordinates are more intuitive in power systems analysis since bus voltages are already typically represented in phasor form. However, rectangular coordinates can enable simpler computation.

Assume there exists a power system with N buses and M sensor measurements. In polar coordinates, the voltage at bus a is represented in its phasor form as $V_a = |V_a|\angle\theta_a$. In rectangular coordinates, the voltage at bus a is represented in Cartesian form as $V_a = e_a + jf_a$. This expression can be rewritten as:

$$e_a = V_a\cos(\theta_a) \tag{5}$$

$$f_a = V_a\sin(\theta_a), \tag{6}$$

$$|V_a|^2 = e_a^2 + f_a^2. \tag{7}$$

AC WLS SE in Polar Coordinates

In polar coordinates, the state vector $x^{(k)}$ for a system with N buses and M measurements contains $2N - 1$ elements with N voltage magnitudes and $N - 1$ voltage angles, since the angle of the slack (reference) bus is assumed to be 0. Hence the state vector would be:

$$x^{(k)} = [\theta_2 \quad \theta_3 \quad \cdots \quad \theta_N \quad V_1 \quad V_2 \quad V_3 \quad \ldots \quad V_N]^T. \tag{8}$$

In $h\left(x^{(k)}\right) \in \mathbb{R}^{2N-1}$, the expressions for the real and reactive power injections at bus m are:

$$P_a = \sum_{k=1}^{N} V_a V_k (G_{ak} \cos(\theta_{ak}) + B_{ak} \sin(\theta_{ak})) \tag{9}$$

$$Q_a = \sum_{k=1}^{N} V_a V_k (G_{ak} \sin(\theta_{ak}) - B_{ak} \cos(\theta_{ak})) \tag{10}$$

where G_{ab}, B_{ab}, G_{aa}, and B_{aa} represent the respective entries of the bus admittance matrix $Y_{bus} = G + jB$. Following the two-port π-model of a transmission line (as seen in Figure 8), the expressions for the real and reactive power flow from bus a to bus b are:

$$P_{ab} = V_a^2 g_{ab} + V_a V_b (-g_{ab} \cos(\theta_{ab}) - b_{ab} \sin(\theta_{ab})) \tag{11}$$

$$Q_{ab} = -V_a^2 \left(\frac{B_{sh}}{2} + b_{ab}\right) + V_a V_b (-g_{ab} \sin(\theta_{ab}) + b_{ab} \cos(\theta_{ab})) \tag{12}$$

where $g_{ab} + jb_{ab}$ is the line impedance and B_{sh} is the line shunt capacitance.

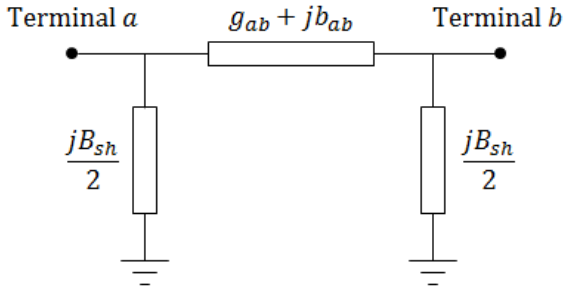

Figure 8. Two-port transmission line model

The measurement Jacobian $H\left(x^{(k)}\right) \in \mathbb{R}^{M \times (2N-1)}$ has the following structure:

$$H\left(x^{(k)}\right) = \begin{bmatrix} \frac{\delta P_{inj}}{\delta \theta} & \frac{\delta P_{inj}}{\delta V} \\ \frac{\delta P_{flow}}{\delta \theta} & \frac{\delta P_{flow}}{\delta V} \\ \frac{\delta Q_{inj}}{\delta \theta} & \frac{\delta Q_{inj}}{\delta V} \\ \frac{\delta Q_{flow}}{\delta \theta} & \frac{\delta Q_{flow}}{\delta V} \\ 0 & \frac{\delta V_{mag}}{\delta V} \end{bmatrix} \tag{13}$$

In $H\left(x^{(k)}\right)$, the entries related to real power injections at bus a are:

$$\frac{\delta P_a}{\delta \theta_a} = \sum_{k=1}^{N} V_a V_k(-G_{ak}\sin(\theta_{ak}) + B_{ak}\cos(\theta_{ak}) - V_a^2 B_{aa} \tag{14}$$

$$\frac{\delta P_a}{\delta \theta_b} = V_a V_b(G_{ab}\sin(\theta_{ab}) - B_{ab}\cos(\theta_{ab})) \tag{15}$$

$$\frac{\delta P_a}{\delta V_a} = \sum_{k=1}^{N} V_k(G_{ak}\cos(\theta_{ak}) + B_{ak}\sin(\theta_{ak}) + V_a G_{aa} \tag{16}$$

$$\frac{\delta P_a}{\delta V_b} = V_a(G_{ab}\cos(\theta_{ab}) + B_{ab}\sin(\theta_{ab})) \tag{17}$$

The entries related to reactive power injections at bus a are:

$$\frac{\delta Q_a}{\delta \theta_a} = \sum_{k=1}^{N} V_a V_k(G_{ak}\cos(\theta_{ak}) + B_{ak}\sin(\theta_{ak}) - V_a^2 G_{aa} \tag{18}$$

$$\frac{\delta Q_a}{\delta \theta_b} = V_a V_b(-G_{ab}\cos(\theta_{ab}) - B_{ab}\sin(\theta_{ab})) \tag{19}$$

$$\frac{\delta Q_a}{\delta V_a} = \sum_{k=1}^{N} V_k(G_{ak}\sin(\theta_{ak}) - B_{ak}\cos(\theta_{ak}) - V_a B_{aa} \tag{20}$$

$$\frac{\delta Q_a}{\delta V_b} = V_a(G_{ab}\sin(\theta_{ab}) - B_{ab}\cos(\theta_{ab})) \tag{21}$$

The entries related to real power flows between bus a and bus b are:

$$\frac{\delta P_{ab}}{\delta \theta_a} = V_a V_b\left(g_{ab}\sin(\theta_{ab}) - b_{ab}\cos(\theta_{ab})\right) \tag{22}$$

$$\frac{\delta P_{ab}}{\delta \theta_b} = -V_a V_b\left(g_{ab}\sin(\theta_{ab}) - b_{ab}\cos(\theta_{ab})\right) \tag{23}$$

$$\frac{\delta P_{ab}}{\delta V_a} = 2V_a g_{ab} - V_b\left(g_{ab}\cos(\theta_{ab}) + b_{ab}\sin(\theta_{ab})\right) \tag{24}$$

$$\frac{\delta P_{ab}}{\delta V_b} = -V_a\left(g_{ab}\cos(\theta_{ab}) + b_{ab}\sin(\theta_{ab})\right) \tag{25}$$

The entries related to reactive power flows between bus a and bus b are:

$$\frac{\delta Q_{ab}}{\delta \theta_a} = -V_a V_b(g_{ab}\cos(\theta_{ab}) + b_{ab}\sin(\theta_{ab})) \tag{26}$$

$$\frac{\delta Q_{ab}}{\delta \theta_b} = V_a V_b(g_{ab}\cos(\theta_{ab}) + b_{ab}\sin(\theta_{ab})) \tag{27}$$

$$\frac{\delta Q_{ab}}{\delta V_a} = -2V_a\left(\frac{B_{sh}}{2} + b_{ab}\right) - V_b(g_{ab}\sin(\theta_{ab}) - b_{ab}\cos(\theta_{ab})) \tag{28}$$

$$\frac{\delta Q_{ab}}{\delta V_b} = -V_a(g_{ab}\sin(\theta_{ab}) - b_{ab}\cos(\theta_{ab})) \tag{29}$$

The entries related to voltage magnitude at bus a are:

$$\frac{\delta V_a}{\delta \theta_a} = 0 \tag{30}$$

$$\frac{\delta V_a}{\delta \theta_b} = 0 \tag{31}$$

$$\frac{\delta V_a}{\delta V_a} = 1 \tag{32}$$

$$\frac{\delta V_a}{\delta V_b} = 0 \tag{33}$$

AC WLS SE in Rectangular Coordinates

In rectangular coordinates, the state vector $x^{(k)}$ for a system with N buses and M measurements contains $2N-1$ elements with N real voltage components e and $N-1$ imaginary voltage components f, since the angle of the slack (reference) bus is assumed to be 0. Hence the state vector would be:

$$x^{(k)} = [e_1 \quad e_2 \quad e_3 \quad \cdots \quad e_N \quad f_2 \quad f_3 \quad \ldots \quad f_N]^T. \tag{34}$$

In $h\left(x^{(k)}\right) \in \mathbb{R}^{2N-1}$, the expressions for the real and reactive power injections at bus a are [21]:

$$P_a = e_a \sum_{k=1}^{N}(G_{ak}e_k - B_{ak}f_k) + f_a \sum_{k=1}^{N}(G_{ak}f_k + B_{ak}e_k) \tag{35}$$

$$Q_a = e_a \sum_{k=1}^{N}(-G_{ak}f_k - B_{ak}e_k) + f_a \sum_{k=1}^{N}(G_{ak}e_k - B_{ak}f_k) \tag{36}$$

G_{ab}, B_{ab}, G_{aa}, and B_{aa} represent the respective real and imaginary entries of the bus admittance matrix $Y_{bus} \in \mathbb{R}^{N\times N}$, where $Y_{bus} = G + jB$.

The expressions for the real and reactive power flow from bus a to bus b are:

$$P_{ab} = (e_a^2 + f_a^2)g_{ab} - g_{ab}e_ae_b + b_{ab}e_af_b - g_{ab}f_af_b - b_{ab}f_ae_b \tag{37}$$

$$Q_{ab} = -(e_a^2 + f_a^2)\left(\frac{B_{sh}}{2} + b_{ab}\right) + g_{ab}e_af_b + b_{ab}e_ae_b - g_{ab}f_ae_b + b_{ab}f_af_b \tag{38}$$

where $g_{ab} + jb_{ab}$ is the line impedance and B_{sh} is the line shunt capacitance.

The expression for the voltage magnitude squared at bus a is, as discussed earlier:

$$|V_a|^2 = e_a^2 + f_a^2 \tag{39}$$

The measurement Jacobian $H\left(x^{(k)}\right) \in \mathbb{R}^{M\times(2N-1)}$ has the following structure:

$$H\left(x^{(k)}\right) = \begin{bmatrix} \frac{\delta P_{inj}}{\delta e} & \frac{\delta P_{inj}}{\delta f} \\ \frac{\delta P_{flow}}{\delta e} & \frac{\delta P_{flow}}{\delta f} \\ \frac{\delta Q_{inj}}{\delta e} & \frac{\delta Q_{inj}}{\delta f} \\ \frac{\delta Q_{flow}}{\delta e} & \frac{\delta Q_{flow}}{\delta f} \\ \frac{\delta V_{mag}^2}{\delta e} & \frac{\delta V_{mag}^2}{\delta f} \end{bmatrix} \tag{40}$$

In $H\left(x^{(k)}\right)$, the entries related to real power injections at bus a are:

$$\frac{\delta P_a}{\delta e_a} = \sum_{k=1}^{N}\left(G_{ak}e_k - B_{ak}f_k\right) + G_{aa}e_a + B_{aa}f_a \tag{41}$$

$$\frac{\delta P_a}{\delta e_b} = G_{ab}e_a + B_{ab}f_a \tag{42}$$

$$\frac{\delta P_a}{\delta f_a} = \sum_{k=1}^{N}\left(G_{ak}f_k + B_{ak}e_k\right) - B_{aa}e_a + G_{aa}f_a \tag{43}$$

$$\frac{\delta P_a}{\delta f_b} = G_{ab}f_a - B_{ab}e_a \tag{44}$$

The entries related to reactive power injections at bus a are:

$$\frac{\delta Q_a}{\delta e_a} = \sum_{k=1}^{N}\left(-G_{ak}f_k - B_{ak}e_k\right) - B_{aa}e_a + G_{aa}f_a \tag{45}$$

$$\frac{\delta Q_a}{\delta e_b} = G_{ab}f_a - B_{ab}e_a \tag{46}$$

$$\frac{\delta Q_a}{\delta f_a} = \sum_{k=1}^{N}\left(G_{ak}e_k - B_{ak}f_k\right) - G_{aa}e_a - B_{aa}f_a \tag{47}$$

$$\frac{\delta Q_a}{\delta f_b} = -G_{ab}e_a - B_{ab}f_a \tag{48}$$

The entries related to real power flows between bus a and bus b are:

$$\frac{\delta P_{ab}}{\delta e_a} = 2g_{ab}e_a - g_{ab}e_b + b_{ab}f_b \tag{49}$$

$$\frac{\delta P_{ab}}{\delta e_b} = -g_{ab}e_a - b_{ab}f_a \tag{50}$$

$$\frac{\delta P_{ab}}{\delta f_a} = 2g_{ab}f_a - g_{ab}f_b - b_{ab}e_b \tag{51}$$

$$\frac{\delta P_{ab}}{\delta f_b} = -g_{ab}f_a + b_{ab}e_a \tag{52}$$

The entries related to reactive power flows between bus a and bus b are:

$$\frac{\delta Q_{ab}}{\delta e_a} = -2\left(\frac{B_{sh}}{2} + b_{ab}\right)e_a + g_{ab}f_b + b_{ab}e_b \tag{53}$$

$$\frac{\delta Q_{ab}}{\delta e_b} = -g_{ab}f_a + b_{ab}e_a \tag{54}$$

$$\frac{\delta Q_{ab}}{\delta f_a} = -2\left(\frac{B_{sh}}{2} + b_{ab}\right)f_a - g_{ab}e_b + b_{ab}f_b \tag{55}$$

$$\frac{\delta Q_{ab}}{\delta f_b} = g_{ab}e_a + b_{ab}f_a \tag{56}$$

The entries related to voltage magnitude at bus a are:

$$\frac{\delta V_a^2}{\delta e_a} = 2e_a \tag{57}$$

$$\frac{\delta V_a^2}{\delta e_b} = 0 \tag{58}$$

$$\frac{\delta V_a^2}{\delta f_a} = 2f_a \tag{59}$$

$$\frac{\delta V_a^2}{\delta f_b} = 0 \tag{60}$$

2.2.3 Computational Challenges

The main computational burden associated with AC SE is the calculation and decomposition of the gain matrix $G = H^T R^{-1} H$. Because the measurement Jacobian H is sparse and R is diagonal, $G \in \mathbb{R}^{(2N-1)\times(2N-1)}$ is structurally/numerically symmetric, sparse (although less sparse than H), and non-negative definite in general (positive definite for a fully observable network). For computational and memory efficiency, the gain matrix is built and stored as a sparse matrix. It can be rewritten as

$$G = \sum_{k=1}^{M} H_k^T R_{kk}^{-1} H_k \tag{61}$$

where H_k is a k-th row of the measurement Jacobian (very sparse) and R_{kk} is the corresponding element of the measurement covariance matrix.

By applying Cholesky decomposition to G, we obtain the non-unique triangular factors L and L^T, where $G = LL^T$. The sparsity of L may vary, depending on the way the decomposition is performed and based on the ordering method. Note that for systems that are not fully observable, the Cholesky decomposition may not exist, and hence no state estimation solution can be found (see Section 2.3).

Assuming a Cholesky decomposition exists, it is combined with the Normal Equations to obtain

$$\left[L\left(x^{(k)}\right)L\left(x^{(k)}\right)^T\right]\Delta x^{(k+1)} = -g\left(x^{(k)}\right). \tag{62}$$

Then forward substitution and back substitution steps are used to solve for entries of $\Delta x^{(k+1)}$.

2.2.4 Decoupled Weighted Least Squares

One way to reduce the computational burden of calculating and decomposing G is to use several simplifying assumptions to arrive at an approximate gain matrix with constant sub-matrices. Because the sub-matrices are constant and smaller, this formulation requires less overall memory usage, and Cholesky decomposition only needs to be performed during the first iteration. Similar to the power flow algorithm, those assumptions are:

- The sensitivity of the real power equations with respect to changes in bus voltage magnitudes is low.
- The sensitivity of the reactive power equations with respect to changes in bus voltage angles is low.

In other words, we can assume that the measurement Jacobian H (in polar form) can be approximated as

$$H\left(x^{(k)}\right) = \begin{bmatrix} \frac{\delta P_{inj}}{\delta\theta} & \frac{\delta P_{inj}}{\delta V} \\ \frac{\delta P_{flow}}{\delta\theta} & \frac{\delta P_{flow}}{\delta V} \\ \frac{\delta Q_{inj}}{\delta\theta} & \frac{\delta Q_{inj}}{\delta V} \\ \frac{\delta Q_{flow}}{\delta\theta} & \frac{\delta Q_{flow}}{\delta V} \\ 0 & \frac{\delta V_{mag}}{\delta V} \end{bmatrix} \approx \begin{bmatrix} \frac{\delta P_{inj}}{\delta\theta} & 0 \\ \frac{\delta P_{flow}}{\delta\theta} & 0 \\ 0 & \frac{\delta Q_{inj}}{\delta V} \\ 0 & \frac{\delta Q_{flow}}{\delta V} \\ 0 & \frac{\delta V_{mag}}{\delta V} \end{bmatrix} = \begin{bmatrix} H_{AA} & H_{AR} \\ H_{RA} & H_{RR} \end{bmatrix} \tag{63}$$

Then the new approximate gain matrix is

$$G \approx \begin{bmatrix} G_{AA} & 0 \\ 0 & G_{RR} \end{bmatrix} \tag{64}$$

where $= diag([R_A \quad R_R])$, $G_{AA} = {H_{AA}}^T R_A^{-1} H_{AA}$, and $G_{RR} = {H_{RR}}^T R_R^{-1} H_{RR}$. Then solve

$$G_{AA}\Delta\theta = H_{AA}^T R_A^{-1} \Delta z_A / V \tag{65}$$

$$G_{RR}\Delta V = H_{RR}^T R_R^{-1} \Delta z_R / V \tag{66}$$

where $\Delta z_A = z_A - h_A(\hat{x})$ and $\Delta z_R = z_R - h_R(\hat{x})$.

2.2.5 DC Weighted Least Squares

Further assumptions can be made to completely simplify the system models. Again, similar to power flow, those assumptions are:

- All bus voltage magnitudes are known to be 1.0 per unit.
- Shunt elements and branch resistances are ignored.
- Only real power measurements are considered.

In this simplest case, the relationship between the real power measurements and the bus angles is linear:

$$z_A = H_{AA} x_A + e_A \tag{67}$$

Then the DC WLS estimator is:

$$(H_{AA}^T R_A^{-1} H_{AA})\Delta x_A = H_{AA}^T R_A^{-1} \Delta z_A \quad (68)$$

2.2.6 Key Assumptions and Limitations

As mentioned in the previous sections, there are multiple assumptions behind the design of the WLS state estimator, and some necessary conditions need to be met in order for a unique and accurate solution to be found. These assumptions and necessary conditions are summarized below:

- Measurement errors are assumed to be independent and have a Gaussian distribution with a mean of 0 and a known variance of σ^2. These variances are chosen based on assumptions regarding metering accuracy and/or historical data for the measurement errors.
- The gain matrix G needs to be relatively well-conditioned in the Normal Equations, or else the WLS state estimator may become numerically unstable and fail to reach convergence. Ill-conditioning can arise from having a large number of injection measurements, very large weighting factors, and the presence of both short and long lines at the same bus. When the gain matrix is ill-conditioned, alternate methods could be used instead of Cholesky factorization to mitigate the computational issues.
- The state estimation solution cannot be found when the Cholesky decomposition does not exist (e.g. for many systems that are not fully observable).
- There must be at least one voltage measurement for each observable island.

2.2.7 Robust State Estimation

Although the focus of this dissertation is on the WLS state estimator due to its prevalence in the power industry, it is worthwhile to briefly review the problem formulation for robust state estimation since it has some advantages compared to the

WLS formulation. First we must define statistical robustness. A state estimator is robust if the estimated system state is insensitive to major deviations in a limited number of redundant measurements. There is frequently a trade-off between robustness and computational speed. As discussed in Section 2.2.6, one major assumption behind the design of the WLS state estimator is that measurement errors are independent, and each error has a Gaussian distribution with a mean of 0 and a known variance. While this assumption may be generally true under normal operating conditions, there are occasionally situations where gross errors can occur due to telemetry noise/failure or even potentially cyber-attacks (which will be discussed in greater detail in Chapter 4).

Before reviewing specific robust state estimation methods, first the definition of an outlier needs to be introduced. In statistics, an outlier is an observation point that is distant from other observations due to either natural variability in the measurement or measurement error. In power systems, outliers could be either measurements that do not contain any actual error but due to the structure of its corresponding equation appear unusual, or they could be incorrectly recorded measurements that differ from their actual value. Naturally occurring outliers are difficult to identify and will strongly bias the state estimate when they contain errors. An outlier measurement that has an undue amount of influence on the state estimate and where the corresponding row of the measurement Jacobian H lies away from the rest of the factors is referred to as a leverage point.

There are many different objective functions and solution methods for robust state estimation, but the general problem formulation for an M-estimator is as follows:

$$\begin{aligned} &\min \sum_{i=1}^{M} \rho(r_i) \\ &s.t.\ \ z = h(x) + r \end{aligned} \tag{69}$$

where $\rho(r_i)$ is a selected function of the measurement residual r_i, $z \in \mathbb{R}^M$ is the set of measurements, $x \in \mathbb{R}^{2N-1}$ is the state vector, and $h(x) \in \mathbb{R}^{2N-1}$ is the measurement

function that corresponds to perfect error-free measurements. The function ρ should have at least the following properties:

- $\rho(r) = 0$ when $r = 0$
- $\rho(r) > 0 \; \forall \; r$
- $\rho(r) = \rho(-r)$
- $\rho(r)$ is monotonically increasing in both $+r$ and $-r$ directions

One example of a function that satisfies all of these conditions is the Least Absolute Value (LAV) M-estimator $\rho(r_i) = |r_i|$.

In general, M-estimation problems can be solved using algorithms such as Newton's method and iteratively reweighted least squares. The LAV estimator, which can be formulated as a linear programming problem, can take advantage of optimization algorithms such as the simplex and interior point methods.

2.3 Network Observability Analysis

The objective of observability analysis is to determine whether a unique estimate can be found for the power system state based on the network topology as well as the type and location of available measurements. This analysis is first performed offline during the initial installation of a state estimator to verify whether the existing installed sensors are sufficient to fully observe the system. It is also performed online prior to running the state estimator to ensure that an estimate can be found from the measurements received at each time step. Sometimes communication or metering errors/failures can cause a power system to become unobservable, in which case the system will be divided into a series of isolated observable islands, each with its own independent phase angle reference (i.e. slack bus).

Observability analysis methods can be divided into two different classes: numerical and topological. Both classes of methods were developed using the same general assumptions:

- Paired P and Q measurements (i.e. each real injection/flow measurement has a corresponding reactive injection/flow measurement)
- Decoupled measurement model (see Section 2.2.4, cannot be used if assumptions do not hold or if current magnitude measurements are included)

2.3.1 Numerical Methods

There are multiple methods for numerical observability analysis. The method described below is based on the nodal variable formulation. Referring back to the decoupled state estimation model from Section 2.2.4, $\theta - P$ observability and $V - Q$ observability can be tested separately if the P and Q measurements are paired. Unlike in $\theta - P$ analysis, there needs to be at least one voltage magnitude measurement in $V - Q$ analysis for each observable island.

The observability of a system depends solely on the topology of the system and the type/location of measurements. It does not depend on branch parameters or the system operating state, so simplifying assumptions can be made for the system parameters. For example, all branches can be assumed to have an impedance of $j1.0$ per-unit, and all bus voltages can be assumed to be 1.0 per-unit as seen in the DC state estimation model (as seen in Section 2.2.5). Then the DC power flows can be written as

$$P_b = A\theta \tag{70}$$

where $P_b \in \mathbb{R}^M$ is the vector of branch power flows, $A \in \mathbb{R}^{M \times N}$ is the branch-bus incidence matrix, and $\theta \in \mathbb{R}^N$ is the vector of bus voltage phase angles. The branch-bus incidence matrix is defined as

$$A(i,j) = \begin{cases} 1 & \text{if bus } j \text{ is the sending end of branch } i \\ -1 & \text{if bus } j \text{ is the receiving end of branch } i \\ 0 & \text{else} \end{cases} \tag{71}$$

If there is an estimate $\hat{\theta}$ such that

$$H_{AA}\hat{\theta} = 0 \tag{72}$$

yet the corresponding branch flow is nonzero, i.e.

$$P_b = A\hat{\theta} \neq 0 \tag{73}$$

then this estimate $\hat{\theta}$ is an unobservable state, and the branches b that have nonzero flows are unobservable branches.

The algorithm for identifying observable islands is as follows:

Step 1. Remove all branches that have no incident measurements.

Step 2. Form $G_{AA} = H_{AA}^T R_A^{-1} H_{AA}$.

Step 3. Factorize G_{AA} using Cholesky factorization. When a zero pivot is encountered, replace that entry with 1.0 and the corresponding entry of the right hand side is assigned an arbitrary value.

Step 4. Identify and remove all unobservable branches and all injections that are incident to these unobservable branches.

Step 5. If no more unobservable branches can be found, identify the observable islands that are divided by the unobservable branches and stop. Else proceed back to Step 2.

2.3.2 Topological Methods

There are also multiple methods for topological observability analysis. (However, there is considerably less literature on the subject compared to numerical observability analysis.) Topological methods differ from numerical methods in that no floating point computations are necessary. Instead they are based entirely on logical operations and thus only need information about network topology and measurement type/location. Similar to the numerical methods, it is assumed that measurements are available in real and reactive pairs, and thus the real part of the decoupled measurement model can be used for observability analysis.

The seminal paper on topological observability analysis [22] was published in 1980. In it, the authors developed a new graph theoretical algorithm after recognizing that

the rank of the H matrix depends only on the system topology. The important theoretical conclusion from their paper is that H is full rank (i.e. the associated network X is observable) if and only if there exists a spanning tree T of full rank (a spanning tree is a tree that is incident to every bus of network X). A tree T is full rank if every branch of T can be assigned in a unique fashion. The algorithm they proposed is divided into two parts, one that processes the line flow measurements and one that processes the injection measurements. The algorithm works as follows:

Part 1. Construct a forest F^l of branches with line flow measurements. Remove branches that form loops with F^l (cotree branches). Remove branches that are incident to no measurement. Identify and label as redundant injection measurements that are incident only to branches of F^l or to cotree branches. Identify the boundary injection measurements. Perform an initial measurement assignment of the branches of F^l (i.e. assign each branch to the flow measurement represented by that branch), which will be known as a.

Part 2. Take the boundary injection measurements from Part 1. Process them in stages one at a time until T is a critical tree or until all boundary injections have been considered. The inputs to each stage are the boundary injection under consideration x_* and the forest (with measurement assignment) output from previous stages. The output of each stage is an updated forest (with measurement assignment) that has fewer components or contains more unmeasured nodes than the input forest.

> **Step 2.1.** Let F_* be the component of the forest of flow-measured branches that contains x_* and T_* be the input forest that contains x_*. The tree T_* is directed away from x_* as follows. Every branch b in T_* that is incident to x_* has x_* labeled as the negative $(-)$ node of b, and the other node is labeled as positive $(+)$.

Step 2.2. After T_* has been directed away from x_*, a sequence of pairwise disjoint sets of branches B_i^l is constructed iteratively. B_0^l is set equal to F_*. B_1^l consists of all branches incident to x_* that are not in B_0^l.

Step 2.3. The steps for deriving B_{i+1}^l are shown in Figure 9. Select the first branch b in B_1^l . Set x as the positive node of b. Check to see if b is in the edges of the input forest T_*^l. If so, check if the initial measurement assignment of b is to node x. If not, check to see if x is in the nodes of the input forest T_*^O. If either x is in the nodes of the input forest T_*^O or the initial measurement assignment of b is not to node x, set the negative node of b to be the T-predecessor branch of node x. If the initial measurement assignment of b^- is to node x, then add every branch that is incident to node x that is not in a previous B_i^l to B_{i+1}^l. If not, add only b^- to B_{i+1}^l. Continue processing each branch until all branches are processed or if a sink is hit.

Step 2.4. If a sink (an unmeasured node not already in T_* or a node belonging to another component of T) is hit by a branch b, then a backtracking process is initiated (illustrated in Figure 10). The procedure recursively selects a sequence of branches, one from each B_i^l, modifies the tree, and updates the assignment. There are five simple subroutines: 'add to tree,' 'any,' 'delete,' 'erase,' and 'search.' 'Add to tree' adds b to T^l and assigns x to b, i.e. $a(b) = x$. 'Any' searches through B_i^l for any branch that is incident to x. 'Delete' removes b from T^l. 'Erase' removes $b^- = -_T(x)$ from T^l and any other T-predecessor branches. 'Search' looks for a branch in B_i^l that is also in T_*^l and is assigned to x.

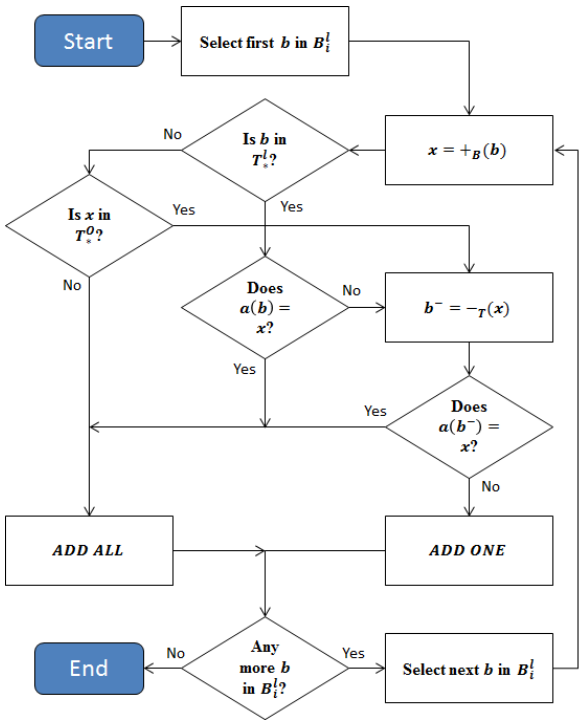

Figure 9. Topological OA flow chart for building the next B_i^l [22]

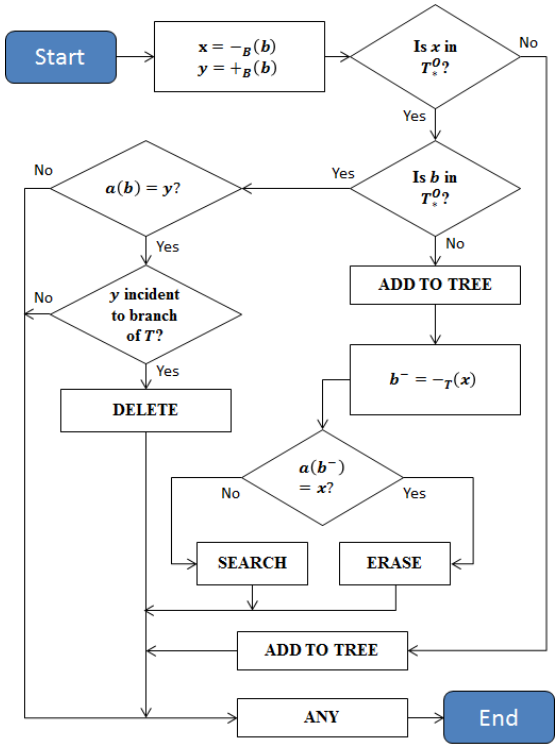

Figure 10. Topological OA flow chart for processing the next branch [22]

After this seminal paper was published, the authors published a number of other papers related to topological observability analysis and state estimation. Most notably, they developed an algorithm that determines the maximal observable subnetwork, i.e. the largest region where bus voltages may be estimated when confronted with a deficiency in measurements [23]. They also developed an algorithm to adding pseudo measurements of estimated bus loads to improve network observability [24]. In addition, they worked on determining the detectability of bad measurements based on the network topology and location measurements (discussed in Section 2.4) [25].

A second influential paper on topological observability analysis [26] was published in 1982. The authors directly search for an observable spanning tree in the measurement graph, using an algorithm for matroid intersections. They reframed the topological observability analysis problem as a graph with N vertices and L edges, where each edge needs to be colored with (i.e. assigned) one of $K + 1$ colors. The objective is to find a spanning tree such that for each color, at most one edge of the spanning tree has that color. Similar to [22], the flow measurements are processed first, followed by the injection measurements. The exact algorithm can be found in [26], Figure 2.

Since those two papers, other papers on the subject have emerged. In 1991, a new algorithm that used the concept of augmenting sequences, including a direct comparison between the proposed algorithm and a numerical observability analysis algorithm, was presented, showing that their topological algorithm was considerably faster [27]. At the same time, another paper that involved a fast algorithm for finding the minimum spanning tree of an augmented graph was presented [28]. The first author of [28] also explored using genetic algorithms to solve the combinatorial problem of topological observability [29]. Between 2005 and 2006, more papers that used genetic algorithms and artificial neural networks were presented. This body of available research literature on topological observability analysis is summarized in Table 2.

Table 2. Topological Observability Analysis Literature Review

Year	Method	Author	Ref
1980	Heuristic	Krumpholz et al.	[22]
1982	Matroid intersection	Quintana et al.	[26]
1991	Augmenting sequences	Nucera and Gilles	[27]
1991	Augmented graph	Mori and Tsuzuki	[28]
1992	Genetic algorithms	Mori and Tanaka	[29]
2005	Artificial neural networks	Jain et al.	[30]
2006	Heuristic	Jain et al.	[31]
2006	Genetic algorithms	Vazquez-Rodriguez et al.	[32]

2.4 Bad Data Detection

A key objective of the state estimator is to detect and remove erroneous measurements, assuming there is sufficient redundancy in the measurement set so that random errors can be filtered out. Measurement errors occur due to a variety of reasons:

- Sensor noise, bias, drift, failure, or wiring mistakes
- Communication noise, error, or failure
- Intentional manipulation of sensor measurements during a cyber-attack

Measurements that are obviously wrong or physically impossible (such as negative voltage magnitudes, extremely large measurements, etc.) can be easily detected through simple logic checks. However, more subtle inconsistencies can be difficult to detect without advanced techniques.

In WLS state estimation, bad data detection and removal are performed after the estimation process. Bad data can manifest as either a single erroneous measurement or multiple erroneous measurements. Multiple bad data can be further categorized into non-interacting bad data, interacting but non-conforming bad data, and interacting and conforming bad data. Interacting measurements are measurements that are strongly correlated whose errors significantly affect one another's estimated values. Conforming measurements are measurements that appear to be consistent with one another.

The degree to which measurements are interacting can be measured by the residual covariance matrix Ω. If $\Omega(i,j) \geq \varepsilon$ (user-defined threshold based on the network and measurement topology), then measurement i and measurement j are strongly interacting. Otherwise, they are considered to be weakly interacting or non-interacting. The matrix is calculated as

$$\Omega = SR \tag{74}$$

Based on the DC state estimation model (see Section 2.2.5), the WLS estimator is

$$\Delta\hat{x}_A = (H_{AA}^T R_A^{-1} H_{AA})^{-1} H_{AA}^T R_A^{-1} \Delta z_A \tag{75}$$

where $\Delta z_A = H_{AA}\Delta x + e$. The estimated value of Δz_A is $\Delta\hat{z}_A = H_{AA}\Delta\hat{x}_A = K\Delta z$. The hat matrix K is equal to

$$K = H_{AA} G_{AA}^{-1} H_{AA}^T R_A^{-1} \tag{76}$$

K can be used to provide an idea of the local measurement redundancy around each sensor. A large diagonal entry relative to the off-diagonal entries in K indicates poor local redundancy. The residual sensitivity matrix S measures the sensitivity of measurement residuals to measurement errors and is equal to

$$S = I - K \tag{77}$$

where I is the identity matrix.

In general, measurements are classified into two different groups: critical and redundant. The removal of a critical measurement from the measurement set would cause the system to become unobservable. The measurement residual of a critical measurement is always zero. Sometimes critical measurements can occur in pairs or k-tuples, in which case the removal of all of these measurements will cause the system to become unobservable. A redundant measurement is a non-critical measurement, and it has a nonzero measurement residual. Bad data can only be detected if removing the erroneous measurement does not render the system unobservable. Hence, errors that occur in critical

measurements cannot be detected. A single measurement containing bad data can be identified if and only if it is not critical and it does not belong to a critical pair.

2.4.1 The Largest Normalized Residual Test

The most commonly used bad data detection test in EMS systems today is the largest normalized residual test. The normalized residual of measurement i is calculated as follows:

$$r_i^N = \frac{|r_i|}{\sqrt{\Omega_{ii}}} \tag{78}$$

Assuming there is a single bad measurement that is not a critical measurement or a member of a critical pair, the largest normalized residual will correspond to the erroneous measurement. This will also hold true for multiple bad data as long as the measurements are non-interacting. The algorithm is as follows:

Step 1. Solve WLS state estimation and obtain the measurement residual vector $r_i = z_i - h_i(\hat{x})$.

Step 2. Compute the normalized residuals.

Step 3. Find the index k that corresponds to the largest r_i^N.

Step 4. If $r_k^N > c$ (user-defined threshold), remove the k-th measurement which is suspected to be bad and proceed to Step 1. Else stop, no bad measurements are suspected.

2.4.2 Key Assumptions and Limitations

As discussed in the previous sections, existing bad data detection techniques such as the largest normalized residual test will definitely detect and remove a bad measurement if:

- The measurement is not critical (i.e. removing it will not make the system unobservable).

- Removing the measurement does not create any critical measurements.

If there is multiple bad data, then the largest normalized residual test can detect bad measurements one at a time if the measurements are non-interacting. If there are multiple interacting bad measurements:

- The largest normalized residual test may still correctly identify bad data if the measurements are non-conforming (i.e. the errors are not consistent with each other).
- The largest normalized residual test may fail to identify conforming measurements.

3 DECOMPOSITION-BASED STATE ESTIMATION

In this chapter, we present a new faster state estimation method, which involves automatically dividing the central SE problem for a large system into a series of smaller problems as well as empirically exploring the effect of the number of sub-problems on the computational speed [33]. Currently the state estimator performs its data processing and computation for a single control area in a central location (the control center). These computations are carried out serially once every 30 seconds to few minutes, depending on the balancing authority (see Table 1). As the number of devices used to monitor the power grid increases and the need to monitor very large interconnections grows, centralized state estimation becomes too slow to be feasible. Our proposed solution to these challenges is decomposition-based state estimation.

The organization of this chapter is as follows. Section 3.1 motivates the need for decomposition-based state estimation and provides the scope of the research. Section 3.2 reviews the relevant research literature on speeding up the performance of the state estimator. Section 3.3 introduces the general consensus ADMM problem as well as discusses its merits and disadvantages as a distributed optimization method. Section 3.4 describes the importance of referencing the local slack buses to a single global slack bus. Section 3.5 discusses the methodology behind the fast decomposition-based SE algorithm. Section 3.6 introduces a problem formulation that would achieve greater solution accuracy.

3.1 Introduction

Centralized state estimation becomes infeasible with an increase in the number of monitoring devices and the need driven by market deregulation for utilities to monitor very large interconnections. First, as more measurements are added into the system, more speed is needed to process them in real time. Secondly, faster processing would enable

real-time wide-area SE [6]. Because the speed of SE depends on the size of the problem (the number of measurements and the number of states to be estimated), the ability to decompose the global SE problem into smaller problems is essential. Once the problem is decomposed, then parallel or distributed computing can be applied to each problem, as seen in distributed state estimation (DSE) where multiple nodes work together to jointly estimate the state of an area [34]. A related problem to DSE is multi-area state estimation (MASE), where several control areas work together to jointly estimate their combined state [35], [36].

This chapter explores two questions that have been largely undiscussed in the literature for DSE and MASE: 1) how the global SE problem could be automatically decomposed into smaller problems, and 2) the impact of the number of sub-problems on the computational speed. In the case of DSE, the boundaries of the sub-problems have frequently been determined through trial and error, and the number of problems is assumed to be given [8]. In MASE, the boundaries for each sub-problem are defined by the physical boundaries of each control area, and the number of problems is defined by the number of control areas involved in the joint SE. The results from this work suggest that methodical decomposition of the global SE problem beyond the natural boundaries of a control area can significantly increase the speed of SE.

The main contributions of this work are as follows. First, we introduce automatic graph partitioning as a systematic method for virtually dividing a power system into an arbitrary number of sub-areas. Then we propose a fast decomposition-based SE approach, an enhancement to traditional WLS SE, which uses the partitioning results to automatically decompose the global SE problem into smaller sub-problems and solve them in unison using the alternating direction method of multipliers (ADMM). Lastly we

explore empirically the impact of the number of sub-problems on the computational speed of the global SE problem for a serial implementation and extrapolate the speedup seen for the considered IEEE test cases to larger power systems. To the best of our knowledge, this is the first time the impact of partition granularity on the speed of the SE problem has been quantitatively evaluated. The proposed approach enables an investigation of how far typical SE problems should be decomposed for the best speed gains.

3.2 Relevant Work

In the literature, there are two related branches of research that focus on speeding up the computational performance of the state estimator. One is DSE, where the burden of SE for a large system is spread over multiple processors or threads (frequently in the same geographic location). The other is MASE, where the burden of SE for a very large interconnection is distributed over multiple processors for each control area (in separate geographic locations). A good survey of different DSE and MASE schemes from prior to 2010 is provided in [35]. The authors classified different schemes by the level of area overlap, the computing architecture, the level of coordination, measurement synchronization, process synchronization, and solution methodology.

For the level of area overlap, schemes could either have completely non-overlapping areas, boundary-bus overlapping areas, virtual bus overlapping areas, tie-line overlapping areas, or extended overlapping areas (see Figure 11). For the architecture, there are hierarchical and decentralized schemes. In a hierarchical architecture, a central coordinating computer communicates with each local state estimator, which is run on local processors either in separate or the same geographic location. In a decentralized architecture, each local state estimator communicates with its neighbors without the help of a central computer. For the level of coordination, each area can submit its results to the

central processor after full convergence of its local SE, each area can submit its results after every local iteration, or some combination of the two approaches could apply. For process synchronization, each local processor can be run synchronously for a hierarchical architecture where the coordination takes place at the central SE level, synchronously or asynchronously if the coordination takes place at the iteration level, or asynchronously for a decentralized architecture. Finally, regarding solution methodology, the majority of surveyed methods rely on the WLS SE formulation while decentralized schemes tend to use Lagrangian relaxation-based algorithms or formulate the WLS equations as an optimization problem.

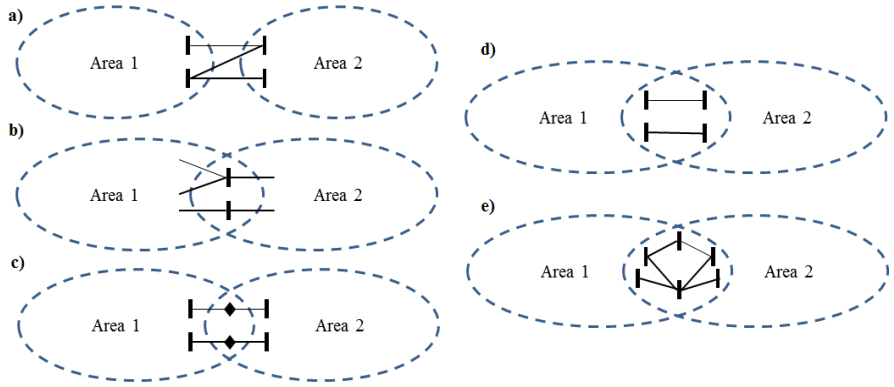

Figure 11. Different levels of area overlap for MASE schemes [35]: a) non-overlapping areas; b) boundary-bus overlapping areas; c) virtual bus overlapping areas; d) tie-line overlapping areas; e) extended overlapping areas.

Since 2010, there have been some other relevant works on DSE and MASE. Meliopoulos et al. [37] introduced a fully distributed three-phase linear state estimator at the substation level, which uses a detailed substation model and all locally available measurements including those from PMUs. Xie et al. [38], [39] described a fully distributed SE algorithm that does not require local observability of all control areas and where the communication topology can differ from the physical topology. Kekatos and Giannakis [36] presented a fully distributed robust SE algorithm that uses ADMM, which

also does not require local observability. Both algorithms use the DC state estimation model, which as discussed in Section 2.2 we know to be not as accurate as the AC model.

Table 3. DSE and MASE Literature Review

Reference, author	Area overl. [a]	Solut. meth. [b]	Estim. state [c]	Meas. type [d]	Coord. scheme [e]
[37] Meliopoulos	NO	NE-H	-	P	It
[38], [39] Xie	NO	NE-H	Opt	C	It
[36] Kekatos and Giannakis	MO	NE-O	Opt	C	It
[40] Karimipour and Dinavahi	NO	NE-H	Unknown	P	It

[a] NO: non-overlap; MO: minimally overlap; FO: fully overlap.
[b] NE: Normal equations; R: relaxation; O: optimization; H: heuristic.
[c] Opt: optimal; Sub: suboptimal.
[d] C: conventional only; P: considers PMU.
[e] SE: SE level; It: iteration level.

3.3 The Alternating Direction Method of Multipliers

ADMM is a simple but powerful distributed convex optimization algorithm that has many applications in large-scale distributed computing and optimization. It blends the benefits of two predecessor algorithms, dual decomposition and augmented Lagrangian methods for convex optimization, and is either equivalent or closely related to many other similar optimization algorithms. It displays the decomposability of dual ascent and the superior convergence properties of the method of multipliers. It was first introduced in the 1970s, and by the mid-1990s, the theory had been mostly established. Although the focus of the algorithm is on distributed computing, the algorithm can also be used serially, where even with no tuning it is competitive with the best known methods for some problems. The main reference for this section on ADMM is [41].

3.3.1 General ADMM Algorithm

In ADMM, the variables x and z are updated in an alternating (sequential) manner, giving the algorithm its name. The ADMM problem formulation is as follows. Consider the problem of minimizing a global objective function that takes the form

$$\begin{aligned} &\min f(x) + g(z) \\ &\text{subject to } Ax + Bz = c \end{aligned} \tag{79}$$

where variables $x \in \mathbb{R}^n$ and $z \in \mathbb{R}^m$, $A \in \mathbb{R}^{p\times n}$, $B \in \mathbb{R}^{p\times m}$, and $c \in \mathbb{R}^p$. The functions f and g are assumed to be closed, proper, and convex. The unaugmented Lagrangian L_0 is assumed to have a saddle point, i.e. $L_0(x^*, z^*, y) \leq L_0(x^*, z^*, y^*) \leq L_0(x, z, y^*)$ $\forall\, x, z, y$. Like in the method of multipliers, the augmented Lagrangian is

$$L_p(x, z, y) = f(x) + g(z) + y^T(Ax + Bz - c) + \frac{\rho}{2}\|Ax + Bz - c\|_2^2. \tag{80}$$

The ADMM iterations are

$$x^{(k+1)} = \underset{x}{\operatorname{argmin}}\, L_p\left(x, z^{(k)}, y^{(k)}\right) \tag{81}$$

$$z^{(k+1)} = \underset{z}{\operatorname{argmin}}\, L_p\left(x^{(k+1)}, z, y^{(k)}\right) \tag{82}$$

$$y^{(k+1)} = y^{(k)} + \rho\left(Ax^{(k+1)} + Bz^{(k+1)} - c\right) \tag{83}$$

where $\rho > 0$ and (k) indicates the k-th interation. The first ADMM equation is the x-minimization step, the second equation is the z-minization step, and the third equation is the dual variable update.

The algorithm can also be written in a scaled form. The scaled dual variable u is

$$u = \left(\frac{1}{\rho}\right)y \tag{84}$$

Using the expression for u, ADMM can be written as

$$x^{(k+1)} = \underset{x}{\operatorname{argmin}}\left(f(x) + \frac{\rho}{2}\left\|Ax + Bz^{(k)} - c + u^{(k)}\right\|_2^2\right) \tag{85}$$

$$z^{(k+1)} = \underset{z}{\operatorname{argmin}}\left(g(z) + \frac{\rho}{2}\left\|Ax^{(k+1)} + Bz - c + u^{(k)}\right\|_2^2\right) \tag{86}$$

$$u^{(k+1)} = u^{(k)} + Ax^{(k+1)} + Bz^{(k+1)} - c \tag{87}$$

By defining the primal residual at iteration $k + 1$ as

$$r^{(k+1)} = Ax^{(k+1)} + Bz^{(k+1)} - c \tag{88}$$

we see that $u^{(k)}$ is related to the running sum of the residuals

$$u^{(k)} = u^{(0)} + \sum_{j=1}^{k} r^{(j)} \tag{89}$$

The dual residual at iteration $k + 1$ is

$$s^{(k+1)} = \rho A^T B(z^{(k+1)} - z^k) \tag{90}$$

ADMM can be considered as a special version of the method of multipliers where a single Gauss-Seidel pass over the variables x and z is used instead of the usual joint minimization step. Separating the minimization over the variables into two steps is what allows for decomposition when f or g are separable.

3.3.2 Convergence Properties and Stopping Criteria

In theory, ADMM displays the following convergence properties:

- Residual convergence
  - The iterates approach feasibility
  - $r^{(k)} \to 0$ as $k \to \infty$
  - $s^{(k)} \to 0$ as $k \to \infty$
- Objective convergence
  - The objective function approaches optimal value
  - $f\left(x^{(k)}\right) + g\left(z^{(k)}\right) \to p^*$ as $k \to \infty$
- Dual variable convergence
  - The dual variable approaches a dual optimal point
  - $y^{(k)} \to y^*$ as $k \to \infty$

In practice, ADMM can be slow to converge to high accuracy. However, it often converges to a modest level of accuracy that is sufficient for many applications within a few tens of iterations. The relatively slower convergence of ADMM distinguishes it from algorithms like Newton's method, where high accuracy can be obtained in only a few iterations.

A reasonable stopping criterion is that both the primal and dual residuals are sufficiently small, i.e.

$$\left\|r^{(k)}\right\|_2 \leq \epsilon^{primary} \text{ and } \left\|s^{(k)}\right\|_2 \leq \epsilon^{dual} \tag{91}$$

where $\epsilon^{primary} > 0$ and $\epsilon^{dual} > 0$. For example, these tolerances could be chosen as follows using an absolute and relative ϵ

$$\epsilon^{primary} = \sqrt{p}\epsilon^{abs} + \epsilon^{rel}\max\{\left\|Ax^{(k)}\right\|_2, \left\|Bz^{(k)}\right\|_2, \|c\|_2\} \tag{92}$$

$$\epsilon^{dual} = \sqrt{n}\epsilon^{abs} + \epsilon^{rel}\left\|A^T y^{(k)}\right\|_2 \tag{93}$$

where $\epsilon^{abs} > 0$ is an absolute tolerance, $\epsilon^{rel} > 0$ is a relative tolerance such as 10^{-3}, p and n are the dimensions of the vectors from Section 3.3.1.

3.3.3 Consensus ADMM Algorithm

The consensus problem is the problem in which multiple agents need to put forth candidate values, communicate with one another, and come to agree on a single consensus value. ADMM-based methods are useful for solving the consensus problem using distributed optimization. The goal is to solve the global problem in such a way that each sub-problem can be handled by its own thread or processor.

The problem formulation for consensus ADMM is as follows. Consider the problem of minimizing a global objective function that is decomposable into *N* parts:

$$\min f(x) = \sum_{i=1}^{N} f_i(x), \tag{94}$$

where $x \in \mathbb{R}^n$ and each smaller objective function f_i is convex. Under these conditions, the original problem can be reformulated in terms of local variables $x_i \in \mathbb{R}^n$ and a common global variable c:

$$\min \sum_{i=1}^{N} f_i(x_i) \text{ s.t. } x_i - c = 0 \tag{95}$$

for $i = 1, \ldots, N$. By introducing these local variables with the constraint that shared local variables should agree, it is possible to split $f_i(x)$ into separate objective functions $f_i(x_i)$.

Then the augmented Lagrangian is:

$$L_p(x_{1:N}, z, y) = \sum_{i=1}^{N} (f_i(x_i) + y_i^T(x_i - c) + \frac{\rho}{2}\|x_i - c\|_2^2), \tag{96}$$

where y is the dual variable and $\rho > 0$ is the penalty factor. By defining c as $\bar{x}$ (the average of the x_i variables), the resulting consensus ADMM algorithm is [41]:

$$x_i^{(k+1)} = \underset{x_i}{\operatorname{argmin}} \left(f_i(x_i) + \left(y_i^{(k)}\right)^T \left(x_i - \bar{x}^{(k)}\right) + \frac{\rho}{2}\left\|x_i - \bar{x}^{(k)}\right\|_2^2 \right) \tag{97}$$

$$y_i^{(k+1)} = y_i^{(k)} + \rho\left(x_i^{(k+1)} - \bar{x}^{(k+1)}\right) \tag{98}$$

For consensus ADMM, the primal and dual residuals are

$$r^{(k)} = (x_1^{(k)} - \bar{x}^{(k)}, \ldots, x_N^{(k)} - \bar{x}^{(k)}) \tag{99}$$

$$s^{(k)} = -\rho\left(\bar{x}^{(k)} - \bar{x}^{(k-1)}, \ldots, \bar{x}^{(k)} - \bar{x}^{(k-1)}\right) \tag{100}$$

The stopping conditions are $\left\|r^{(k)}\right\|_2 \leq \epsilon^{primary}$ and $\left\|s^{(k)}\right\|_2 \leq \epsilon^{dual}$, where

$$\left\|r^{(k)}\right\|_2 = \sqrt{\sum_{i=1}^{N} \left\|x_i^{(k)} - \bar{x}^{(k)}\right\|_2^2} \tag{101}$$

$$\left\|s^{(k)}\right\|_2 = \sqrt{N\rho^2\|\bar{x}^{(k)} - \bar{x}^{(k-1)}\|_2^2} \tag{102}$$

3.3.4 Application to State Estimation

Performing state estimation with nonlinear $h(x)$ functions involves solving nonconvex optimization problems. As discussed in Section 2.2.1, currently AC state estimation relies on the Gauss-Newton method, a modification of Newton's method that

is used to solve nonlinear least squares problems. Typically these models are iteratively linearized via the Gauss-Newton method or using the DC approximation discussed in Section 2.2.5 [36]. Convergence is not guaranteed, not even local convergence as in Newton's method, and the method is sensitive to the initial guess.

When ADMM is applied directly to a nonconvex problem, it may not converge, and even if it does converge, it may not converge to an optimal point. It must be considered as simply another local optimization method. The hope is that it might possibly have better convergence properties than other local optimization methods in terms of faster convergence or convergence to a point with a better objective value. Depending on the initial values of the variables $x^{(0)}$ and $y^{(0)}$ and the parameter ρ, ADMM can converge to different and sometimes non-optimal points.

3.4 Slack Bus Referencing

In centralized state estimation, the calculated states are referenced to the system slack bus (an arbitrarily chosen reference bus). In decentralized state estimation, the solution for each sub-area or sub-problem is referenced to its own slack bus, rather than the global slack bus. Hence, before we integrate the solutions for the different sub-areas or sub-problems, we must first reference them all back to the global slack bus. There are two main ways to achieve this outcome.

The simplest and most accurate method is only possible if phasor measurement units are installed. For multi-area state estimation (solving the joint state estimation problem for several control areas), each sub-area must contain at least one phasor measurement unit. In that scenario, a bus whose phase is measured by a PMU will be selected as the slack bus for that sub-area. Because PMUs use a synchronized GPS clock, measurements gathered from different sub-areas are synchronized to the same time (as opposed to SCADA measurements which are received by the control center asynchronously). Alternately, for distributed state estimation (solving the state estimation

problem for a single control area), the global problem should be partitioned such that each sub-problem contains at least one independent PMU measurement. This in effect would limit the maximum number of possible sub-problems, which would be equal to the number of installed PMUs. The exact methodology for including PMU measurements in the state estimation formulation is discussed in Section 3.6.3.

If phasor measurements are not available, then the second method must be used instead. It is not as accurate as the PMU-based method, but it is still useful because it enables decomposition-based state estimation in the absence of PMU installation. The second method only requires that SCADA measurements be available. Although SCADA measurements are not actually synchronized, we can assume that they are collected at approximately the same time, because the dynamics of the power grid are relatively slow-moving in general. (However, this assumption no longer holds true during a transient power system event, such as a transmission line fault.) This assumption is typical for traditional state estimation, which only includes SCADA measurements. The exact methodology is discussed in Section 3.5.3.

3.5 Fast Decomposition-Based State Estimation

Fast decomposition-based SE uses automatic graph partitioning to decompose the global WLS SE problem into N smaller problems. This reduces the order of the SE problem, which in turn increases the computational speed. There are three main steps in the algorithm. The first is determining how to divide (partition) the global SE problem into N smaller sub-problems. The second is solving the sub-problems in unison with consensus ADMM at each iteration. The last step is reintegrating the sub-problem solutions to obtain the global solution. The overall execution flow is shown in Figure 12.

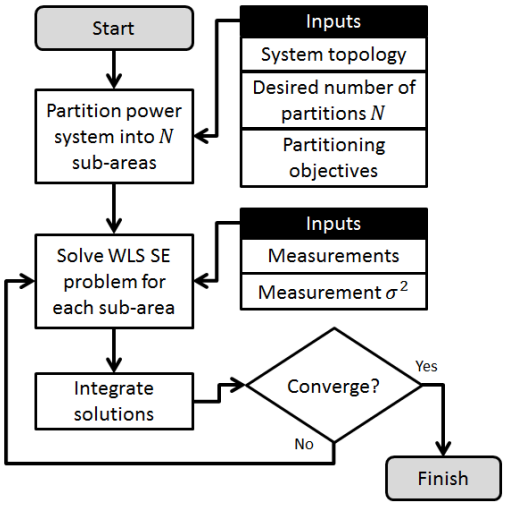

Figure 12. Fast decomposition-based SE algorithm flowchart

3.5.1 Automatic Graph Partitioning

The computational time required for SE increases nonlinearly with the size of the problem, which is determined by the number of estimated states and measurements. Hence, we want to decompose the global SE problem into a series of smaller problems, thus reducing the problem size and improving the overall speed. The most intuitive approach for decomposition is to virtually partition the power system into *N* sub-areas and then assign the appropriate subset of the global measurements to each sub-problem. Because a power system is essentially an undirected graph composed of vertices (buses) and edges (lines), this decomposition can be considered as a graph partitioning problem.

Small partitioning problems can be solved by hand empirically, but large graphs are cumbersome to partition manually, especially when certain partition properties are desirable. Hence, there arose the need for automatic graph partitioning methods. Existing partitioning methods can divide a graph into *N* partitions with a roughly equal number of vertices, where the vertices in each partition are contiguous, while simultaneously reducing the number of edges that traverse the partitions [42].

For this work, there are three desirable properties for how the global SE problem is partitioned. First, the number of buses in each partition should be roughly balanced, so that the number of estimated states in each SE sub-problem is approximately equal. For a serial implementation, this will improve the overall SE solution time since it is the sum of the sub-problem solution times. Unbalanced partitioning would result in disproportionately longer times for larger partitions, which increases the solution time.

Secondly, the number of virtual tie lines (lines that traverse partitions) should be as few as possible. This is due to the way that decomposition-based SE is formulated. In order for partitions to reach consensus about the global system state using ADMM, the state vectors for each sub-problem must include the sub-area's perceptions of the angle and voltage for at least one bus in the neighboring sub-areas. Because adding more virtual tie lines adds more estimated states to the sub-problems which increases the computation time, we want to keep the number of virtual tie lines to a minimum.

The third desirable property is the physical contiguity of the buses in each partition. This is to ensure the observability of each sub-area and serves to reduce the number of virtual tie lines.

Since these properties match the objectives of existing graph partitioning methods, we used a standard graph partitioning toolbox METIS in this work to partition four IEEE test systems. METIS is a set of serial graph partitioning programs based on multi-level recursive bisection and k-way partitioning schemes [43]. Multi-level k-way produces higher quality results than recursive bisection but has issues enforcing contiguity when the desired number of partitions approaches the overall system size. For this work, when the desired number of partitions is less than half the number of buses in the test case, multi-level k-way was selected as the partitioning scheme. Else recursive bisection was selected.

The output of METIS is an assignment list that describes which buses belong in each partition. From those results, virtual tie lines are automatically identified. Then the

measurements for each sub-problem are automatically assigned from the set of global measurements.

The time it takes to partition a power system is negligible compared to the time it takes to perform state estimation. For example, METIS can partition graphs with millions of vertices into hundreds of parts in only a few seconds on current generation workstations and PCs [44]. Most power systems are smaller graphs, on the order of tens of thousands of vertices, so the partitioning time will be even faster.

3.5.2 Problem Formulation

Using the partitioning results, the global SE problem can be decomposed into a series of smaller problems. However, to solve them in unison, the original WLS problem (introduced in Section 2.2) needs to be reformulated to include consensus ADMM. The global objective function that we want to minimize in (1) can be rewritten as a series of smaller problems with local x_i state vectors:

$$f_i(x_i) = [z_i - h_i(x_i)]^T R_i^{-1}[z_i - h(x_i)], \tag{103}$$

where i denotes the local quantity of the WLS SE problem, subject to the constraint that shared x_i states must agree. By combining (103) with the right hand side of (97), we obtain an expression for the new local objective function J_i. Taking the partial derivative of J_i with respect to x_i, we obtain

$$g_i(x_i) = \frac{\partial J_i(x_i)}{\partial x_i} = -2H_i^T(x_i)R_i^{-1}[z_i - h_i(x_i)] + \left(y_i^{(k)}\right)^T + \rho\left(x_i - \bar{x}^{(k)}\right)^T. \tag{104}$$

By setting $g_i(x_i)$ to zero and using the Gauss-Newton method to expand its Taylor series around x_i, then neglecting the higher-order terms, we obtain

$$G_i\left(x_i^{(k)}\right) = 2H_i^T\left(x_i^{(k)}\right)R_i^{-1}H_i\left(x_i^{(k)}\right) + \rho. \tag{105}$$

Then each local state estimation problem is calculated as

$$\left[G\left(x_i^{(k)}\right)\right]\Delta x_i^{(k+1)} = -g\left(x_i^{(k)}\right) \tag{106}$$

$$x_i^{(k+1)} = x_i^{(k)} + \Delta x_i^{(k+1)} \tag{107}$$

3.5.3 Slack Bus Angle Referencing Using SCADA Measurements

In this work, we assume that only SCADA measurements are available to the state estimator. As discussed in Section 3.43.4, the solution of each local state estimation problem is referenced to a local slack bus. In order to integrate the local solutions into the global solution, all states must be referenced to the same global slack bus. The objective is to use SCADA measurements from opposite ends of the tie-lines to estimate the angle difference between pairs of neighboring areas.

A small illustrative example is shown in Figure 13. Assume a given power system is divided into four different areas. Assume the global slack bus angle is $\theta_{1,ref}$ in Area 1. The state estimation problem for Area 1 is solved simultaneously as the state estimation problems for Areas 2, 3, and 4. After solving the state estimation problems, all of the calculated bus angles in Area 1 are relative to $\theta_{1,ref}$, which is set arbitrarily to 0 for simplicity. Similarly the bus angles in Areas 2, 3, and 4 are referenced to $\theta_{2,ref}$, $\theta_{3,ref}$, and $\theta_{4,ref}$ respectively, each of which is also set arbitrarily to 0. The objective is to reference the bus angles in Areas 2, 3, and 4 to the global reference angle $\theta_{1,ref}$. In other words, we want to find $\Delta(\theta_{1,ref} - \theta_{2,ref})$, $\Delta(\theta_{1,ref} - \theta_{3,ref})$, and $\Delta(\theta_{1,ref} - \theta_{4,ref})$.

In our example, Areas 1 and 2 are connected by line A, Areas 2 and 3 are connected by line B, and Areas 2 and 4 are connected by line C. We can estimate $\Delta(\theta_{1,ref} - \theta_{2,ref})$ by assuming that both Area 1 and Area 2 have access to SCADA line flow measurements across line A (i.e. Area 1 and Area 2 overlap by at least one bus). Extend the bus voltage state vector for Area 1 to include the state $\theta_{2,A(Area1)}$, where $\theta_{2,A(Area1)}$ denotes Area 1's estimated value of the bus angle $\theta_{2,A}$. Extend the bus voltage state vector for Area 2 to include the state $\theta_{1,A(Area2)}$, where $\theta_{2,A(Area1)}$ denotes Area 2's estimated value of the bus angle $\theta_{1,A}$. Then $\Delta(\theta_{1,ref} - \theta_{2,ref})$ can be estimated as

$$\Delta(\theta_{1,ref} - \theta_{2,ref}) \cong \frac{1}{2}\left[\left(\theta_{1,A\,(Area1)} - \theta_{1,A(Area2)}\right) + \left(\theta_{2,A(Area1)} - \theta_{2,A(Area2)}\right)\right] \quad (108)$$

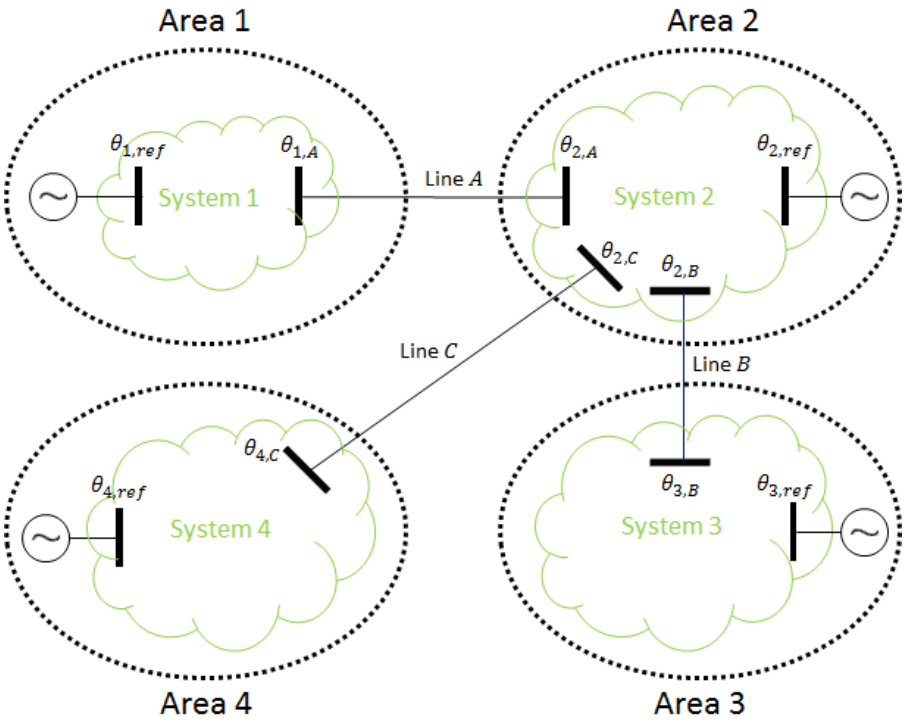

Figure 13. Slack referencing using only SCADA measurements for 4 areas

However, while it is possible to estimate $\Delta(\theta_{1,ref} - \theta_{2,ref})$, it is not possible to directly estimate $\Delta(\theta_{1,ref} - \theta_{3,ref})$ and $\Delta(\theta_{1,ref} - \theta_{4,ref})$ since Area 1 does not share any tie lines with Areas 3 and 4. Instead we have to estimate $\Delta(\theta_{2,ref} - \theta_{3,ref})$ and $\Delta(\theta_{2,ref} - \theta_{4,ref})$ first, using the same method we used to estimate $\Delta(\theta_{1,ref} - \theta_{2,ref})$. Then we can effectively obtain

$$\Delta(\theta_{1,ref} - \theta_{3,ref}) \cong \Delta(\theta_{1,ref} - \theta_{2,ref}) + \Delta(\theta_{2,ref} - \theta_{3,ref}) \quad (109)$$

$$\Delta(\theta_{1,ref} - \theta_{4,ref}) \cong \Delta(\theta_{1,ref} - \theta_{2,ref}) + \Delta(\theta_{2,ref} - \theta_{4,ref}) \quad (110)$$

Another way to visualize this idea is to consider the areas as a tree data structure (illustrated in Figure 14). In our example, Area 1 is the root of the tree, because it

contains the global slack bus. Area 1 has one child, Area 2. Area 2 has two children, Area 3 and Area 4.

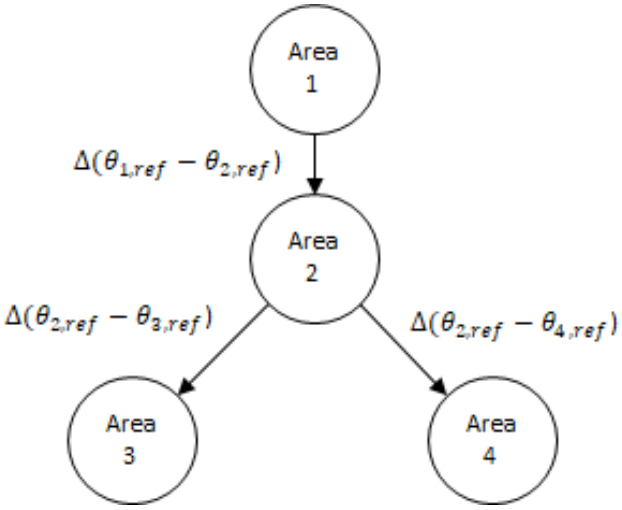

Figure 14. Tree representation of slack referencing for 4-area example

In a hierarchical state estimator architecture, the central coordinator would have access to a tree of how each sub-area or sub-problem is connected to the global slack area (the area containing the global slack bus). Using that tree, they can reference the states for each sub-area or sub-problem to the parent area until they reach the global slack area. Once slack referencing is complete, we can move onto solution integration.

3.5.4 Solution Integration

After the solutions for all sub-problems are referenced to the same global reference bus using the approach described in the previous section, the ADMM averaging step integrates the separate sub-problem solutions into the global SE solution. States that are not shared between sub-problems are taken directly from the solution of their respective sub-problem. Shared states are averaged across their sub-problems:

$$\bar{x}^{(k+1)} = \frac{1}{N}\sum_{i=1}^{N} x_i^{(k+1)} \tag{111}$$

3.5.5 Simulation Results

The fast decomposition-based SE algorithm was implemented serially on a single processor and tested on the IEEE 14 bus, 57 bus, 118 bus, and 300 bus systems. We made an assumption that there are a sufficient number of measurements in the system that observability is ensured even at the finest granularity of partitioning (a single bus). Practical applications would not use such high granularity, but this assumption was necessary to illustrate the full computational trend. Hence the measurements for each system were automatically generated from their power flow results at a global redundancy ratio, defined as the number of measurements divided by the number of estimated states, of approximately 2.9. Then each power flow value was perturbed by a randomly generated Gaussian error with a mean of 0 and a standard deviation of 0.01, which represents a measurement error standard deviation of 1%. Since the measurements are based on the power flow results, the true state is known for our simulated systems. Hence, the accuracy of each SE approach can be quantified using the absolute error between the estimated power system state and the true state.

Timing tests for decomposition-based SE were carried out with different numbers of partitions, ranging from two partitions to fully decomposed (where every partition contains a single bus). These times were compared against the respective time for central WLS SE. To ensure a fair comparison, the core implementation of decomposition-based SE was similar to that of central WLS SE with necessary changes added to support ADMM, as described in Section 3.1.3. Also, the same convergence criteria of $\varepsilon < 10^{-4}$ was used for both SE approaches. All times shown are averaged across 10 runs. The timing was performed using the *tic toc* function in MATLAB R2014b on a 64-bit workstation that has a 2.80 GHz Intel Xeon W3530 CPU and 6 GB of installed RAM.

Using METIS, the IEEE 14 bus system was partitioned into 2, 4, 6, 8, 10, 12, and 14 sub-areas. The IEEE 57 bus system was partitioned into 2, 4, 8, 16, 24, 32, 44, and 57 sub-areas. The IEEE 118 bus system was partitioned into 2, 4, 8, 16, 32, 48, 64, 88, and

118 sub-areas. The IEEE 300 bus system was partitioned into 2, 4, 8, 16, 32, 62, 113, 200, and 300 sub-areas. The ADMM penalty factor ρ was set to 1.

3.5.5.1 Computational Speed

The aggregate timing results for decomposition-based SE and central WLS SE (represented as a single partition) are shown in Figure 15. Among the tested number of partitions, the fastest average decomposition-based SE time for the IEEE 14 bus system was for 2 partitions, each containing 7 buses. For the IEEE 57 bus system, the fastest decomposition-based SE time was for 4 partitions. Of the 4 partitions, 2 sub-areas contained 15 buses each, 1 sub-area contained 14 buses, and 1 sub-area contained 13 buses. For the IEEE 118 bus system, the fastest time was for 8 partitions. Of the 8 partitions, 6 sub-areas contained 15 buses each, and 2 sub-areas contained 14 buses each. For the IEEE 300 bus system, the fastest time was for 32 partitions. Of the 32 partitions, 16 sub-areas contained 9 buses each, 11 sub-areas contained 10 buses each, 2 sub-areas contained 11 buses, one sub-area contained only 6 buses, and one contained only 4 buses.

The best case speedup factors, defined as the central SE computation time divided by the fastest decomposition-based SE time, for the 4 test systems ranged from 1.293 to 7.892, as shown in Figure 16. For the considered systems, larger best case speedups were observed for larger systems. Empirical results suggest that initially increasing the number of partitions decreases the computational time for decomposition-based SE. By partitioning the global SE problem into smaller problems, we reduce the size of each SE problem, and hence a decrease in overall time is observed for a serial implementation. As the number of partitions grows, we begin to observe an increase in the computation time of decomposition-based SE. This is due to an increase in the number of virtual tie line measurements in each local SE problem, which adds more states that need to be estimated. In the case of the 14 bus case, this caused the overall SE computation time for 12 and 14 partitions to exceed the global SE time. However, this is an effect seen only in

very small systems. It is important to note that for the three larger systems, decomposition-based WLS SE is always faster than central WLS SE regardless of the number of partitions, even for a serial implementation. Because large SE problems can always be automatically broken into smaller SE problems that are then solved, the decomposition-based SE is very scalable.

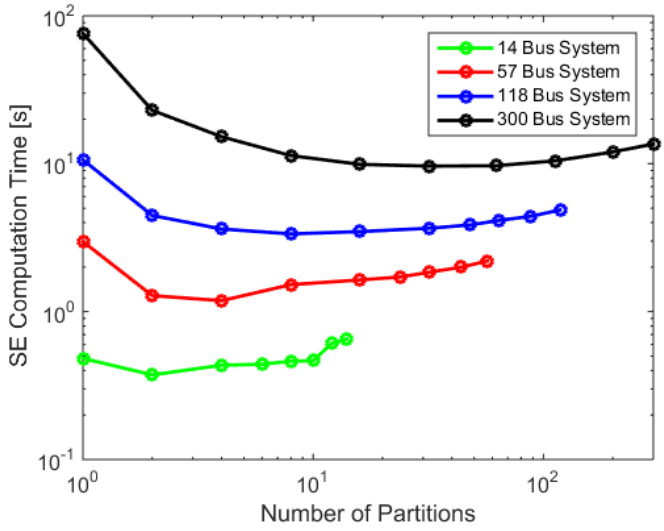

Figure 15. Timing results for 4 standard IEEE test systems

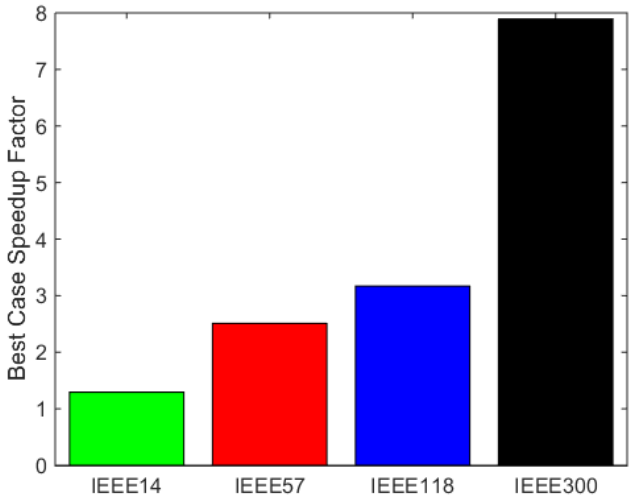

Figure 16. Speedup factors for 4 standard IEEE test systems

3.5.5.2 Accuracy

The emphasis of this work is on computational speed, so the discussion of accuracy is focused around the fastest case for each test system. A comparison of the solution accuracy for bus voltages is shown in Figure 17 for both central and decomposition-based SE. A similar comparison of the accuracy for bus angles is presented in Figure 18.

For both central and decomposition-based SE, the average absolute errors for voltage magnitudes and angles are on the order of 10^{-3}. Because nominal voltage magnitudes are 1 per unit while nominal angles are 0 radians, this level of error has a bigger relative impact on angles than voltage magnitudes. However, since the absolute error is small, it will have a negligible impact on most buses.

In general, the solution accuracy of decomposition-based SE is comparable to that of central SE but slightly worse. This is to be expected since each sub-problem receives only a subset of the global measurements. This slight loss in solution accuracy is acceptable when computational speed is of the utmost importance. For example, decomposition-based SE could be used to solve larger problems in real time that are currently intractable for central SE.

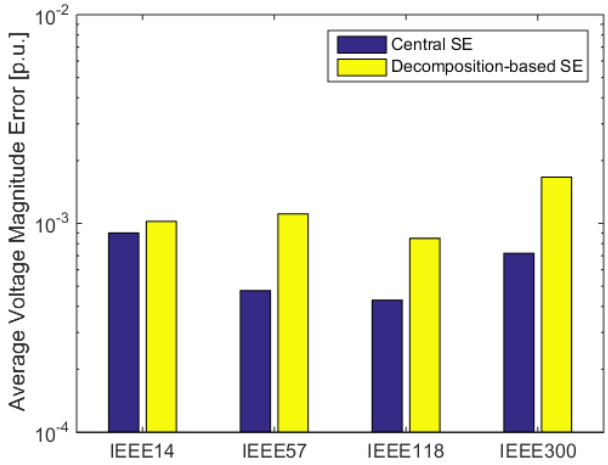

Figure 17. Average voltage magnitude errors for 4 standard IEEE test systems

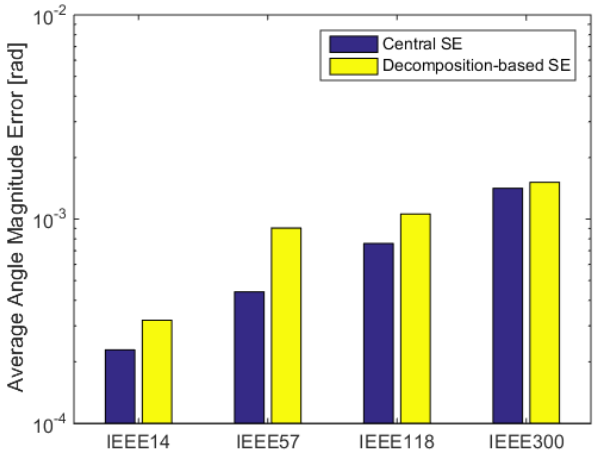

Figure 18. Average angle errors for 4 standard IEEE test systems

3.5.5.3 Implications for Larger Systems

Empirical results suggest that the best case speedup factor *sf* increases with the number of buses *n*. If we assume a linear relationship between the speedup factor and the system size, we can extrapolate that $sf = 0.0228n+0.9269$, as visualized in Figure 19. This extrapolated relationship suggests that the proposed decomposition-based SE approach has the potential to greatly increase the speed of SE for real power systems.

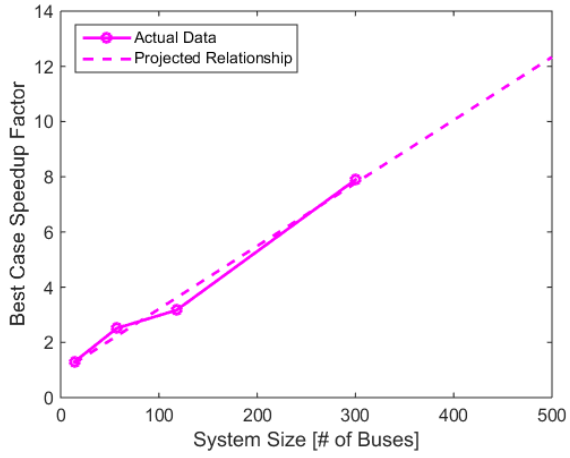

Figure 19. Potential speedup for larger systems

3.5.6 Observability and Bad Data Detection

Observability of each partition is a necessary condition for decomposition-based SE. If one or more partitions are unobservable, estimates for some states may be unobtainable by decomposition-based SE, even though centralized SE may be able to estimate them. This situation tends to arise when the following conditions are true: a) the partition is very small, b) it has few measurements relative to the number of states, and c) there are insufficient measurements at the partition boundary to make the states observable. Mechanisms to merge neighboring partitions with the unobservable partitions could be implemented to regain observability.

Because decomposition-based SE is built on the WLS estimation method, BDD and correction can occur after the estimation process. Since we implemented decomposition-based SE serially and in a central manner, BDD could also be performed in a central manner after obtaining the final joint SE result, using the largest normalized residual test. If it is desirable that BDD be performed in a decentralized way, the method of [36] could be applied.

3.5.7 Conclusions

This work used automatic graph partitioning to enable a scalable decomposition-based SE approach. Empirical results suggest that decomposition reduces computation time even for a serial implementation, except for very small cases. This implies that any SE problem should be decomposed beyond its physical boundaries. With the speed gains, decomposition-based SE could be used to process a greater number of measurements than is currently tractable for central SE.

3.6 Improving Solution Accuracy

If the accuracy of the state estimator solution is of the utmost importance as opposed to computational speed, the methodology proposed in this section is more

suitable. The work presented in this section differs from the work presented in Section 3.5 in the following ways:

- The focus of the work presented in this section is primarily empirical analysis of the effect of partitioning on computational *accuracy*, whereas the focus of the work presented in Section 3.5 was primarily empirical analysis of the effect of partitioning on computational *speed*.
- Automatic graph partitioning methods were applied to the power system *measurement graph* in this work, whereas they were applied to the power system *topology graph* in Section 3.5.
- The problem formulation presented in this section uses convex relaxation of AC state estimation via semidefinite programming. This approach is best if solution accuracy is prioritized over speed. The semidefinite problem could be solved in MATLAB using CVX and SeDuMi. The work in Section 3.5 was based on a DC SE approximation, which did not require the use of any optimization programs.
- The slack referencing would be more accurate with the availability of both PMU and SCADA measurements, whereas the slack bus referencing method described in Section 3.5.3 assumes the availability of only SCADA measurements.
- The algorithm presented in this section is fully distributed, whereas the algorithm from Section 3.5 relies on a hierarchical SE architecture.

3.6.1 Automatic Measurement Graph Partitioning

The power system measurement graph is always a subset of the power system topology graph. Because it is prohibitively expensive to meter every single component in a power system (i.e. every transmission line, bus, transformer, etc.), the typical redundancy ratio (which is the number of measurements divided by the number of

estimated quantities) of realistic power systems is around 1.5-2.0. In our previous work, we assumed a redundancy ratio of 3.0 so that we could achieve very fine partitioning granularity to test large numbers of subproblems. Hence there was no difference between partitioning the power system topology graph and the measurement graph. However, when full measurements are not available, the partitioning results would be different. The main change is the input to the METIS software tool, which would be a smaller graph whose edges and vertices are created based on available measurements. The rest of the process remains identical to Section 3.5.1.

3.6.2 Convex Relaxation of AC State Estimation Using Semidefinite Programming

To guarantee ADMM convergence for the nonconvex and nonlinear AC state estimation problem, the problem first needs to undergo relaxation. There are two main ways of achieving this goal. The first class of methods is second-order cone program (SOCP) relaxation, which has been applied to the AC optimal power flow problem. The second class is semidefinite programming (SDP) relaxation, which has been applied to the AC state estimation problem. SOCPs are nonlinear convex problems, which includes linear programs, convex quadratic programs, and quadratically constrained convex quadratic programs. SDP includes SOCP as a special case. Although SOCPs can be expressed as SDPs, solving them via an SDP method is not a good idea, because interior-point methods that directly solve the SOCP have a much better complexity than an SDP method under the worst case scenario [45], [46]. A review of the relevant papers that focus on convex relaxation of the AC state estimation problem is presented in Table 4. Most of the research literature has focused on SDP rather than SOCP. Several papers also discussed using distributed methods to help reduce the computational complexity of SDP.

Table 4. AC State Estimation Convex Relaxation Literature Review

Year	Author	Description	Method?	DSE?
2011-	Zhu and	SDP problem formulation for	SDP	ADMM

2014	Giannakis [47]–[50]	static state estimation, distributed optimization using ADMM		
2012-2015	Weng et al. [51]–[54]	SDP problem formulation for static state estimation, distributed optimization using Lagrangian dual decomposition	SDP	Lagrangian
2014	Kim et al. [55], [56]	SDP problem formulation for online static state estimation (does not use dynamic equations like extended Kalman filter)	SDP	No
2015	Kim [57]	SDP problem formulation for online static state estimation, use ADMM to reduce SDP complexity	SDP	ADMM
2017	Zheng et al. [58]	Bilinear state estimation (two-stage change of variables), not mathematically equivalent to Gauss-Newton but provides similar results in practice	Bilinear	ADMM
2017	Zhang et al. [59]	Penalized SDP and SOCP problem formulations for both WLAV and WLS estimators	SDP+SOCP	No
2018	Aghamolki et al. [60]	SOCP problem formulation for WLAV estimator, SDP cutting plane method used to strength SOCP relaxation	SDP+SOCP	No

According to [50] and [57], the convex relaxation of AC state estimation is solved as follows. Recall the AC state estimation problem formulation from Section 2.2. Let Y denote the bus admittance matrix, where the (m, n) entry of Y is

$$Y_{mn} = \begin{cases} -y_{mn} & \text{if } (m,n) \in \mathcal{E} \\ \bar{y}_{nn} + \sum_{v \in N_n} y_{nv} & \text{if } m = n \\ 0 & \text{else} \end{cases} \tag{112}$$

where $\mathcal{E}$ is the set of transmission lines, y_{mn} is the line admittance between bus m and bus n, $\bar{y}_{nn}$ is the shunt admittance for bus n, and N_n is the set of buses connected to bus n via transmission lines.

Let us express each quadratic measurement in z in terms of the outer-product matrix of $V = vv^H$. Define the new admittance matrices as

$$Y_n = e_n e_n^T Y \tag{113}$$

$$Y_{mn} = (\bar{y}_{mn} + y_{mn}) e_m e_m^T - y_{mn} e_m e_n^T \tag{114}$$

where $\{e_n\}_{n=1}^N$ is the canonical basis of $\mathbb{R}^N$. Their Hermitian counterparts are:

$$H_{P,n} = \frac{1}{2}(Y_n + Y_n^H) \tag{115}$$

$$H_{Q,n} = \frac{j}{2}(Y_n - Y_n^H) \tag{116}$$

$$H_{P,mn} = \frac{1}{2}(Y_{mn} + Y_{mn}^H) \tag{117}$$

$$H_{Q,mn} = \frac{j}{2}(Y_{mn} - Y_{mn}^H) \tag{118}$$

$$H_{v,n} = e_n e_n^T \tag{119}$$

Then the quadratic measurement model can be expressed linearly with respect to $V \in \mathbb{C}^{N\times N}$ as

$$P_n = Tr(H_{P,n} V) \tag{120}$$

$$Q_n = Tr(H_{Q,n} V) \tag{121}$$

$$P_{mn} = Tr(H_{P,mn} V) \tag{122}$$

$$Q_{mn} = Tr(H_{Q,mn} V) \tag{123}$$

$$|V_n|^2 = Tr(H_{V,n} V) \tag{124}$$

Then the new AC state estimation problem formulation becomes

$$\min_V \sum_{m=1}^{M} w_m (z_m - Tr\{H_m V\})^2 \tag{125}$$

$$\text{s.t. } V \geq 0 \text{ and } \text{rank}(V) = 1 \tag{126}$$

where $w_m = \frac{1}{\sigma_m^2}$. At this point, this formulation is still nonconvex, because the cost function in (125) has a degree of 4 with respect to the entries of V, and the rank constraint in (126) is nonconvex. Using the Schur complement, (125) and (126) can be rewritten in a linear form using an auxiliary vector $\chi \in \mathbb{R}^M$.

$$\{\hat{V}_2, \hat{\chi}_2\} = \underset{V,\mathrm{X}}{\operatorname{argmin}} \sum_{m=1}^{M} w_m \chi_m = \underset{V,\mathrm{X}}{\operatorname{argmin}}\, w^T \chi \tag{127}$$

$$\text{s.t. } V \geq 0 \text{ and} \tag{128}$$

$$\operatorname{rank}(V) = 1 \tag{129}$$

$$\begin{bmatrix} -\chi_m & z_m - Tr(H_m V) \\ z_m - Tr(H_m V) & -1 \end{bmatrix} \leq 0 \;\; \forall\, m \tag{130}$$

In this new formulation, only the rank constraint is still nonconvex. By relaxing (dropping) this constraint via SDP, we arrive at the final problem formulation, which consists of (127), (128), and (130). This convex relaxation has a worst-case computational complexity of $O\left(M^4\sqrt{N}\log\left(\frac{1}{\epsilon}\right)\right)$, where ϵ represents a given solution accuracy. The computational burden of the algorithm motivates the need for a distributed implementation, such as ADMM.

To recover the estimated voltage state $\hat{v}$ from $\hat{V} = \hat{v}\hat{v}^H$, we can use eigenvalue decomposition.

$$\hat{V} = \sum_{i=1}^{r} \lambda_i u_i u_i^H \tag{131}$$

where $= rank(\hat{V})$, the eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_r > 0$ are positively ordered, and their corresponding eigenvectors are denoted as $u_1, u_2, \ldots, u_r$. Because the best rank-one approximation for $\hat{V}$ is $\lambda_1 u_1 u_1^H$, we can choose the state estimate to be $\hat{v} = \sqrt{\lambda_1}u_1$.

3.6.3 Slack Bus Angle Referencing Using PMU Measurements

As discussed in Section 3.4, there are two different methods for slack bus referencing. The approach described in this section is more accurate than the approach described in Section 3.5.3, but it relies on the availability of PMU measurements (i.e. both PMU and SCADA measurements are inputs to the state estimator). PMU measurements have several advantages over traditional SCADA measurements. First of all, they are synchronous, unlike SCADA measurements which are asynchronous. Also,

PMU measurements are linear functions of the bus voltage state vector, unlike SCADA measurements which are nonlinear functions of the state in AC WLS state estimation (as discussed in Section 2.2). To incorporate PMU measurements, the convex relaxed problem formulation from Section 3.6.2 needs to be expanded as follows.

Let B_{PMU} denote the set of buses that have an installed PMU. Let $z_{PMU} = H_{PMU}v + e_{PMU}$, where z_{PMU} is the vector of PMU measurements, H_{PMU} is the PMU measurement matrix, and e_{PMU} is the vector of PMU measurement noise. Each PMU measurement error is assumed to be independent of the SCADA measurement errors and Gaussian with a mean of 0 and a variance of σ_s^2.

In polar coordinates, the expression for H_{PMU} is very straightforward but does not capture useful line current information. In rectangular coordinates, this information can be included (see [61]) as

$$H_s = \begin{bmatrix} e_s^T & 0^T \\ 0^T & e_s^T \\ S_n Re(Y_{fl}) & -S_n Im(Y_{fl}) \\ S_n Im(Y_{fl}) & S_n Re(Y_{fl}) \end{bmatrix} \tag{132}$$

where e_s^T is the s-th row of the identity matrix, $Y_{fl} \in \mathbb{C}^{2N_l \times N_b}$ is the line-to-bus admittance matrix (N_l is the number of lines and N_b is the number of buses), and S_n is the binary $L_n \times 2N_l$ matrix that selects rows of Y_{fl} such that they correspond to the lines connected to bus s.

When we include PMUs, the augmented WLS cost becomes

$$\min_V \sum_{m=1}^{M} w_m \left(z_m - h_m(v)\right)^2 + \sum_{s \in B_{PMU}} w_s \|z_s - H_s v\|_2^2 \tag{133}$$

where $w_s = \frac{1}{\sigma_s^2}\ \forall\ s \in B_{PMU}$, z_s is the PMU measurement at bus s, and H_s is the PMU measurement matrix corresponding to z_s. Because of the legacy SCADA measurements, this augmented problem is also nonconvex. Hence we use convex relaxation again to obtain

$$\{\hat{X}, \hat{\chi}\} = \underset{X,V,v,\chi}{\operatorname{argmin}} w^T \chi + \sum_{s \in B_{PMU}} w_s [Tr(H_s^H H_s V) - 2Re\{z_s^H H_s v)] \tag{134}$$

$$\text{s.t.} X = \begin{bmatrix} V & v \\ v^H & 1 \end{bmatrix} \geq 0 \text{ and} \tag{135}$$

$$\text{rank}(X) = 1 \tag{136}$$

$$\begin{bmatrix} -\chi_m & z_m - Tr(H_m V) \\ z_m - Tr(H_m V) & -1 \end{bmatrix} \leq 0 \ \ \forall\, m \tag{137}$$

After solving the optimization problem for each sub-area, to set the corresponding slack bus voltage angle to 0, we need to rotate v_k by multiplying it with $\frac{V_{ref}^H}{|V_{ref}|}$, where V_{ref} is the complex bus voltage for each local slack bus. Then we can apply the same slack bus referencing method described in Section 3.5.3, but using the more accurate states estimated from hybrid PMU and SCADA measurements.

3.6.4 ADMM-Based AC State Estimation

Partition the global state estimation problem into multiple areas (in the literature, the boundaries are established by control area or arbitrarily for computational purposes), and reformulate the problem such that it can be solved using a distributed optimization algorithm. Each area solves the following modified WLS problem

$$f_k(W_k) = \sum_{l=1}^{M_k} w_l \left[z_k^l - Tr\left(H_k^l W_k\right) \right]^2 \tag{138}$$

$$\widehat{W}_k = \underset{W_k}{\operatorname{argmin}} \sum_k f_k(W_k) \tag{139}$$

$$\text{s.t. } W_k \geq 0, \forall k, W_k^j = W_j^k, \forall k, \forall j \in A_k \tag{140}$$

where v_k is the voltage vector for area k, $W_k = v_k v_k^H$, M_k is the number of measurements included in area k, z_k^l is each measurement included in area k, and H_k^l is the subset of rows and columns from H in Section 3.6.2 corresponding to buses in area k. For neighboring areas k and j, let S_{jk} denote their shared buses. Then W_k^j is the submatrix of

W_k that contains the rows and columns corresponding to S_{jk}, and W_j^k is the submatrix of W_j that contains the rows and columns corresponding to S_{jk}. The coupling equality constraint $W_k^j = W_j^k$ forces the voltages of the shared buses to be the same. An equivalent problem statement for each area is

$$\widehat{W}_k = \underset{W_k>0}{\operatorname{argmin}} \sum_k f_k(W_k) \tag{141}$$

$$\text{s.t. } Re\left(W_k^j\right) = R_{kj}, \forall j \in A_k\,, \forall k \tag{142}$$

$$Im\left(W_k^j\right) = I_{kj}, \forall j \in A_k\,, \forall k \tag{143}$$

where R_{kj} and I_{kj} are auxiliary matrices to handle the coupling constraints. The augmented Lagrangian function of (141) is

$$\begin{aligned}
&\mathcal{L}\left(\{W_k\},\{R_{kj}\},\{I_{kj}\},\{\Gamma_{kj}\},\{\Lambda_{kj}\}\right) \\
&= \sum_k \Bigg\{ f_k(W_k) \\
&+ \sum_{j\in A_k} Tr\left[\Gamma_{kj}\left(Re\left(W_k^j\right) - R_{kj}\right)\right] + \sum_{j\in A_k} \frac{\rho}{2}\left\|Re\left(W_k^j\right) - R_{kj}\right\|_F^2 \\
&+ \sum_{j\in A_k} Tr\left[\Lambda_{kj}\left(Im\left(W_k^j\right) - I_{kj}\right)\right] + \sum_{j\in A_k} \frac{\rho}{2}\left\|Im\left(W_k^j\right) - I_{kj}\right\|_F^2 \Bigg\}
\end{aligned} \tag{144}$$

Then the steps for the resulting ADMM algorithm can be summarized as follows:

Step 1 – Update W_k^{i+1} for each area *k*.

$$\begin{aligned}
W_k^{i+1} = \underset{W_k\geq 0}{\operatorname{argmin}}\, f_k(W_k) &+ \sum_{j\in A_k} Tr\left[\Gamma_{kj}^i\left(Re\left(W_k^j\right)\right)\right] + \sum_{j\in A_k} \frac{\rho}{2}\left\|Re\left(W_k^j\right) - R_{kj}^i\right\|_F^2 \\
&+ \sum_{j\in A_k} Tr\left[\Lambda_{kj}^i\left(Im\left(W_k^j\right)\right)\right] + \sum_{j\in A_k} \frac{\rho}{2}\left\|Im\left(W_k^j\right) - I_{kj}^i\right\|_F^2
\end{aligned} \tag{145}$$

Step 2 – Update the auxiliary variables for each area *k*.

$$R_{kj}^{i+1} = \frac{1}{2}\left[Re\left(W_k^j\right)^{i+1} + Re\left(W_j^k\right)^{i+1}\right] \tag{146}$$

$$I_{kj}^{i+1} = \frac{1}{2}\left[Im\left(W_k^j\right)^{i+1} + Im\left(W_j^k\right)^{i+1}\right] \quad (147)$$

Step 3 – Update the ADMM multipliers per area *k*.

$$\Gamma_{kj}^{i+1} = \Gamma_{kj}^{i} + \frac{\rho}{2}\left[Re\left(W_k^j\right)^{i+1} - Re\left(W_j^k\right)^{i+1}\right] \quad (148)$$

$$\Lambda_{kj}^{i+1} = \Lambda_{kj}^{i} + \frac{\rho}{2}\left[Im\left(W_k^j\right)^{i+1} - Im\left(W_j^k\right)^{i+1}\right] \quad (149)$$

One way to recover the actual voltage estimate v for each area k from $W_k = v_k v_k^H$ is to use eigenvalue decomposition. Then $W_k = \sum_{i=1}^{r} \lambda_i u_i u_i^H$, where $r = rank(W_k)$, $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_r > 0$ (the positive ordered eigenvalues), and u_i are the eigenvectors. The state estimate can be chosen to be $v(u_1) = \sqrt{\lambda_1} u_1$.

4 POWER SYSTEM CYBERSECURITY

Chapter 4 provides background information on power system cybersecurity. Section 4.1 introduces cybersecurity as an emerging area of concern for the power industry. Section 4.2 describes the existing information and communication architecture of energy control centers and their vulnerabilities. Section 4.3 assesses the potential impact of a major cyber-attack on the power grid in terms of number of people impacted and economic damage. Section 4.4 provides an overview of relevant cyber-attacks in recent years. It covers both cyber-attacks on the electric grid that have succeeded in disrupting electricity service to customers as well as other cyber-attacks that have penetrated electric utility networks. Section 4.5 describes the response of the United States government and the American power system industry. Section 4.6 reviews specific power system applications that are particularly vulnerable to cyber-attacks as discussed in the research literature.

4.1 Introduction

Currently cybersecurity is a major challenge faced by governments and businesses around the world. In recent years, concerns have arisen that cyber-attacks launched by organizations with large amounts of resources at their disposal might choose to target critical infrastructure such as industrial control systems (ICS), which would be potentially devastating in terms of its personal and economic impact. The grid is one example of an ICS. The first well-known cyber-attack on an industrial control system occurred as early as 2000. However, in the power industry, cybersecurity is a relatively new concern arising in the past decade, although research on the subject has existed since the mid-2000s. The literature has remained largely academic in nature since there have been few cyber-attacks on the electric grid to date that have succeeded in permanently damaging infrastructure or disrupt the service of electricity to end customers for long periods of

time. Due to the lack of historical information about grid cyber-attacks, the main focus of the existing literature has been on classifying different types of cyber-attacks, identifying general vulnerabilities in the control and operation of the grid, and formulating/detecting novel attacks.

In the past, cyber-attacks on industrial control systems were carried out by disgruntled employees or as a means for financial extortion (such as ransomware), but increasingly they are cyber-weapons deployed by nation-states. The most sophisticated cyber-attacks require months to years of reconnaissance, planning, and coordination, which means they require the investment of massive amounts of resources that only nation-states can afford to provide. In recent years, there has been a rash of cyber-attacks that targeted utilities and energy companies in the United States, although none have succeeded to date in shutting off power to customers. These attacks have been linked to hostile foreign nation-states. The end goal of the attackers is unclear, but the focus thus far has been to gather information. The intent of these attacks could be to simply commit industrial espionage, or it could be more insidious, such as setting the stage for future cyber-attacks that aim to do major infrastructure damage.

4.2 The Vulnerability of the Power Grid

Electric utilities use real-time SCADA systems to monitor and control their assets. A typical SCADA architecture for an energy control center is shown in Figure 20. There are five main components. RTUs are field-based devices that are used to manage the flow of power at an electrical substation. Substations are typically located in remote areas and are unattended by humans. Programmable logic controllers (PLCs), intelligent electronic devices (IEDs), and relays are used to automate tasks at each substation. A control center with central computers manages the remote equipment at the substations and processes, analyzes, and archives the collected real-time information. A communication system is used to convey monitoring information from substations back to the control center as

well as send commands from the control center to the substations, using a communication protocol (such as DNP3, the predominant standard used in distribution substations in North America). In addition to substation communication, a control center may also communicate with its neighboring control centers to coordinate the transfer of power from one area to another. Finally there is a human-machine interface (HMI) at the control center that allows power system operators to supervise and manage the flow of power. It is important to emphasize that the SCADA network at an energy control center is typically isolated from the corporate network and the Internet via firewalls. Engineering and IT functions can connect to the corporate network but cannot directly control and operate the power grid. (However, as discussed in Section 4.4, it is possible to bridge this type of system protection.)

These legacy SCADA systems were originally designed for safety and reliability. There is an implicit assumption of trust of all system components and communications, so these systems are vulnerable to the cyber-threat posed by malicious actors. As the connectivity between the Internet, different networks, and different systems increases, the power grid becomes increasingly susceptible to cyber-attacks. For example, generating turbines were once controlled and operated mechanically, but now they are mostly controlled by industrial control systems via remote automation. Also, grid modernization efforts to incorporate digital automation into protection, operation, and control have introduced more Internet protocol enabled points into the power system network, which further increases the likelihood of malicious intrusion. Furthermore, an increasing number of new vulnerabilities may be introduced into the SCADA system via system additions or integrations. Finally, greater requirements for sensitive data collection and exchange between parties such as utilities, coordinators, and customers create additional points-of-entry. For example, networks may be reconfigured to allow one-time access for a particular need, and that access is forgotten, leaving behind an open door for malicious actors [62].

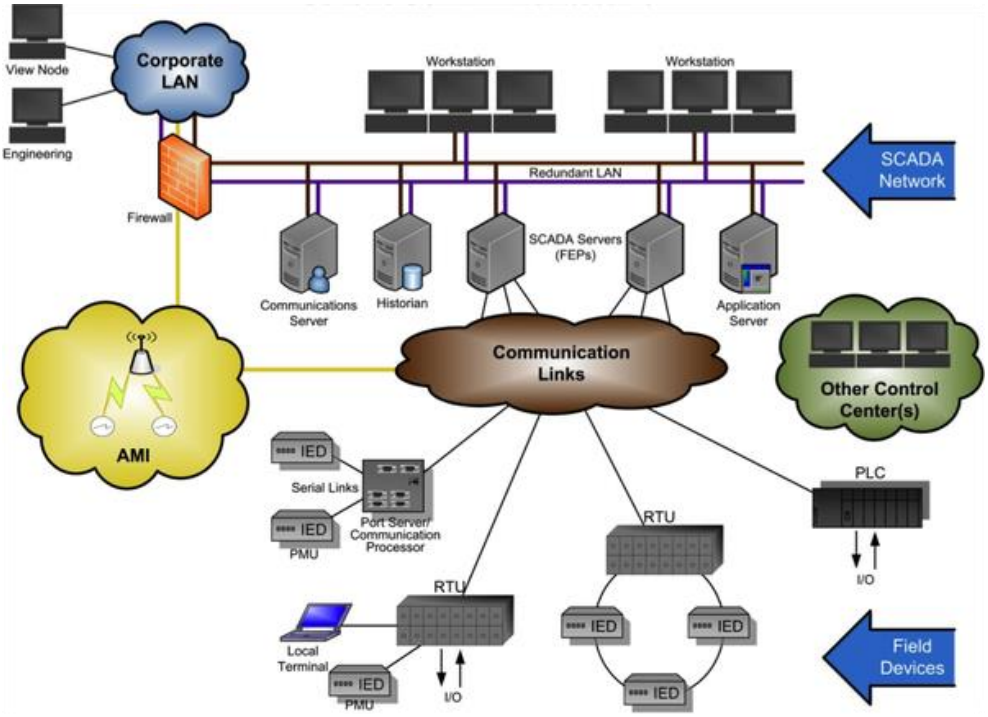

Figure 20. SCADA architecture for the control center [63]

Because the power grid is a mix of older legacy equipment and newer digital assets built over the span of decades, it is difficult to defend the entire system against cyber-attacks since there are so many potential vulnerabilities. Also, because substations are unmanned and geographically disbursed in remote locations, they need to be secured against physical threats. (Experts have noted that if a threat actor is able to physically access a substation, there is no limit to the amount of potential damage since malware can be introduced directly to devices, equipment can be physically destroyed, and relays can be manipulated [62].) Because a successful attack on the bulk generation and transmission systems would cause the most severe consequences, those should be the primary focus of power system cybersecurity research. Other essential grid functions that are at risk to cyber-attacks include:

- Distribution systems

- Backup systems
- Communication systems
- Blackstart capability

Some of these will be discussed in greater detail in Section 4.6.

4.3 Potential Impacts of a Grid Cyber-Attack

A 2015 joint report by Lloyd's, a specialist insurance and reinsurance underwriting company, and the University of Cambridge's Centre for Risk Studies [64] explored the implications of a simulated cyber-attack on the U.S. power grid in terms of potential economic impact and outage times. The scenario they created assumed a piece of malware was able to infect energy control centers in parts of the Northeast and seizes control of 50 generators (out of almost 700 generators in the region), forcing them to overload and burn out. This event had the potential to destabilize the Eastern Interconnection and trigger a blackout that impacts 93 million people. The estimated total impact to the US economy was $243B.

While cyber-attacks have not yet caused widespread power outages in the United States, we can examine other major blackouts that have occurred in recent history to get another estimate of their potential impact. The infamous Northeast blackout of August 2003 impacted approximately 50 million people and caused 100 deaths. (A satellite image of the event is shown in Figure 21.) The estimated economic damage was approximately $6B, and restoration took between 2 days and several weeks. Overall, this event led to 592M customer-hours of lost service [65]. In September 2017, Hurricane Maria devastated parts of Florida and the Caribbean. Puerto Rico was especially hard-hit since the storm passed directly overhead, wiping out most of their major infrastructure. This event is considered to be one of the longest power outages in human history (thus far) with 1248M+ customer-hours of lost service [65]. As of May 2018, there were still pockets of the island without electricity.

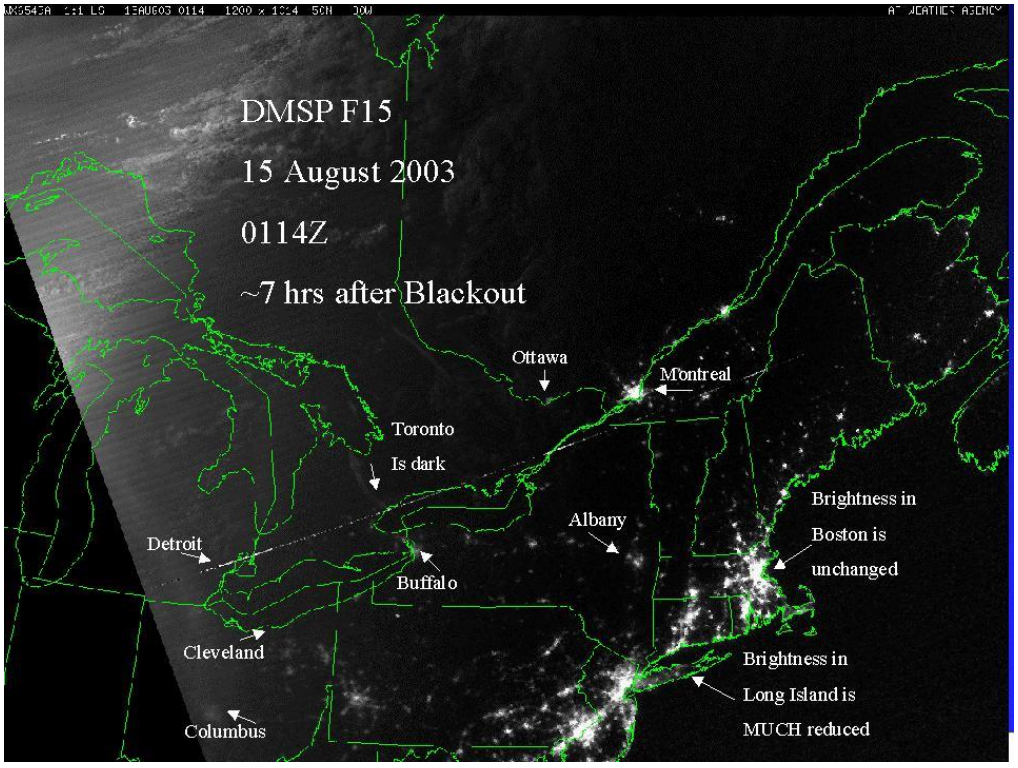

Figure 21. NOAA satellite image of the U.S. Northeast blackout in August 2003 [66]

4.4 Significant Cyber-Attacks

As of May 2018, there have only been two confirmed cyber-attacks on the power grid that have succeeded in disrupting the service of electricity. In recent years, there have been a number of other cyber-attacks on the grid (listed in Table 5) that, while they did not cause outages, did succeed in causing economic damage and/or the hacking of encrypted information, which could be used at a later date to launch more sophisticated attacks.

Table 5. Incomplete List of Recent Attacks on the Power Grid

Date	Area Impacted	Loss of Power?	Brief Description and Impact
03/2007	Aurora Generator Test (Demonstration) [67]	No	Idaho National Labs demonstrated how a cyber-attack that rapidly opens/closes a diesel generator's breakers out of phase can cause it to explode.
04/2013	Pacific Gas & Electric	No	Snipers using rifles fired at the Metcalf

	(US) [68]		transmission substation in California, knocking out 17 transformers and causing over $15M in damage.
12/2015	Prykarpattyaoblenergo, Chernivtsioblenergo, Kyivoblenergo (Ukraine) [7], [69]	Yes	Approximately 225,000 people lost power for up to 6 hours; control centers not fully functional even after 2 months. See Section 4.4.2 for more details.
12/2016	Ukrenergo (Ukraine)	Yes	One fifth of Kiev lost power for 1 hour. Possible large-scale test of new automated malware, using information gained from the 2015 attack. See Section 4.4.3 for more details.
12/2016	Burlington Electric (US)	No	Malware found on isolated laptop. No impact on grid operations.
04/2017	EirGrid (Ireland) [70]	No	Hackers gained access to the utility's communications by tapping their Vodafone network. No impact on grid operations.
05/2017	Wolf Creek Nuclear Operating Corp (US)	No	Hackers compromised computers by emailing fake resumes that contained malicious code to senior engineers. No impact on grid operations.

In this section, we will discuss several significant cyber-attacks that are relevant to power system cybersecurity. The first is the well-known Stuxnet attack in 2009-2010 that was able to bridge an air-gapped industrial control system in Iran that was used for uranium-enrichment. Although Stuxnet was not an attack on the power grid, it does demonstrate that air-gapping alone is not enough to protect industrial control systems. The second cyber-attack discussed in this section is the December 2015 cyber-attack in Ukraine. This incident is the first well-documented cyber-attack on a power system that succeeded in shutting off power to customers. It was almost certainly carried out by a nation-state, and it was incredibly sophisticated in nature. The third cyber-attack we will study is the December 2016 cyber-attack in Ukraine. Using information that was most likely obtained during the 2015 attacks, attackers were able to shut off power in Kiev, Ukraine for approximately an hour. This incident was potentially a dry run to test the latest malware on a large-scale power grid.

4.4.1 Stuxnet (2009)

Stuxnet was most likely a cyber-weapon developed by nation-states to target the uranium-enrichment facilities in Iran. (It reportedly ruined almost 20% of Iran's nuclear centrifuges, and 58% of the infected computers were in Iran [71].) It is noteworthy as the first confirmed example of malware that was specifically tailored to target industrial control systems used for power plants, dams, oil pipelines, and other critical infrastructure. Despite the fact that the facilities utilized air gaps, Stuxnet was able to bridge them via infected USB drives. (Air gaps are network security measures to ensure that a secure computer network is physically isolated from unsecured networks such as the Internet. They are typically used in military/government computer systems, financial computer systems, industrial control systems such as SCADA, and life-critical systems including nuclear power plant controls. They represent nearly the maximum protection one network could have from another network.) Once Stuxnet entered a system, the worm would infect and update other network computers that were not directly connected to the Internet. Its target was to look for computers and networks with very specific configurations.

First, it would infect all machines that were running Microsoft Windows. It was able to evade automatic detection systems, because its device drivers were digitally signed with the private keys of two certificates from JMicron and Realtek, which allowed it to install kernel-mode rootkit drivers without notifying users. Next, it would look for whether a given machine contained project files from Siemens's WinCC/PCS 7 (Step 7) SCADA control software and infect them. It intercepted communications between the WinCC software and the target Siemens PLC devices, so that it could install itself on PLCs unnoticed. Stuxnet's payload only attacked PLC systems with variable-frequency drives from two vendors: Vacon (Finland) and Fararo Paya (Iran), and it monitored the frequency of the attached motors, attacking only systems that spin between 807 Hz and 1210 Hz. When specific criteria were met, it modified the frequency to 1410 Hz, then to

2 Hz, and then to 1064 Hz, thus changing the rotational speed of the connected motors. It also installs a rootkit that masked these changes in rotational speed from monitoring systems, so that the centrifuges still appeared to be operating normally even when they were not [71]. The way the worm was spread is illustrated in Figure 22. Figure 23 illustrates the potential vulnerabilities in an industrial control system network.

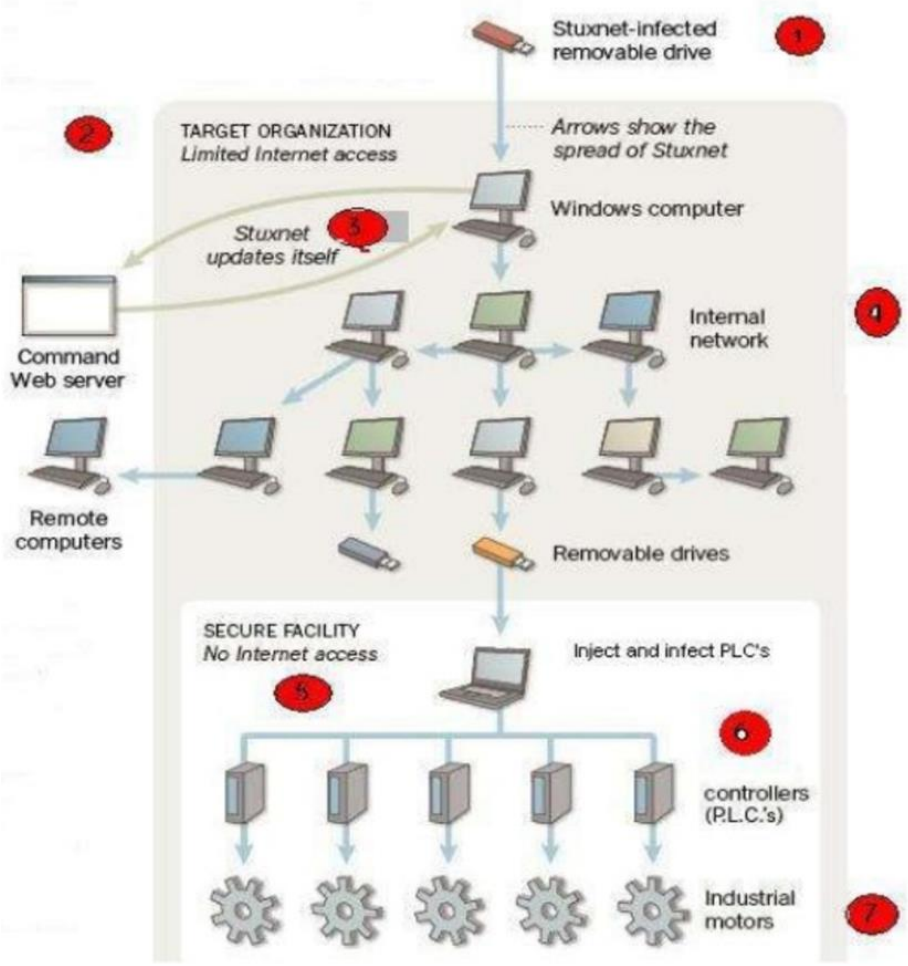

Figure 22. Attack procedure in Stuxnet ICS cyber-attack [72]

There are several key takeaways from Stuxnet. First, it specifically targeted industrial control systems. Next it was able to bridge air gaps, which are the industry standard for network protection. Finally it spied on fast-spinning centrifuges, altering their speed to cause them to tear themselves apart while providing false normal feedback

to monitoring systems. That final step was crucial in masking the attack and fooling human operators into believing that the centrifuges were still performing as intended. This real-world example demonstrates the level of coordination and sophistication that can be expected from a nation-state that decides to attack critical infrastructure such as the power grid. It also illustrates that such an attack would be difficult but not impossible to carry out, requiring vast amounts of resources, planning, and reconnaissance.

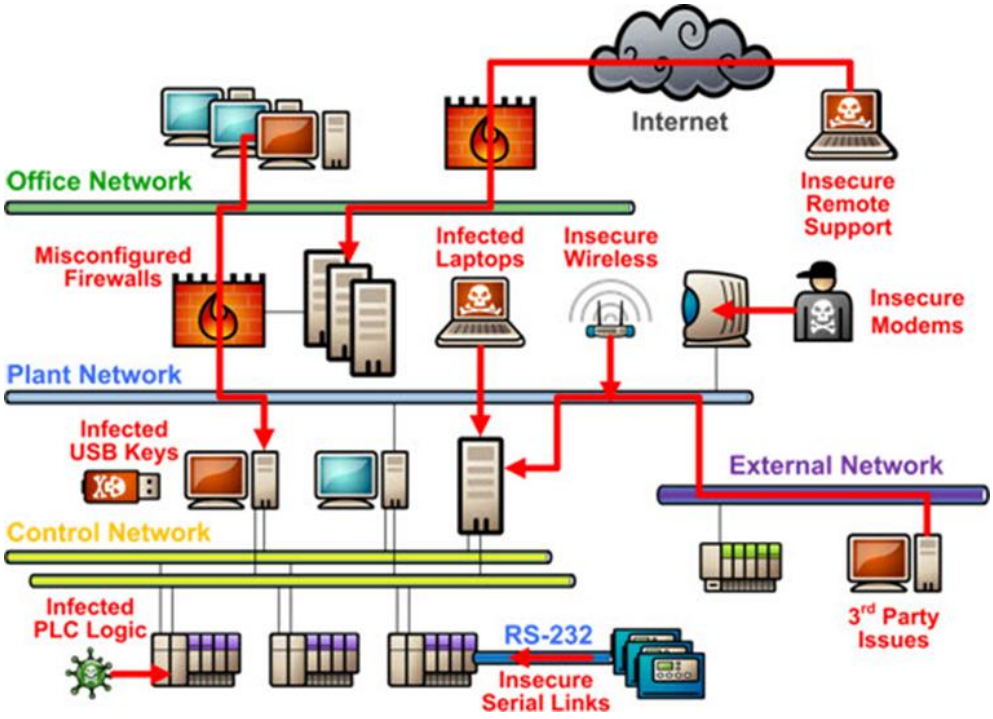

Figure 23. Potential vulnerabilities in an ICS cyber-attack [73]

4.4.2 Ukraine (2015)

The 2015 cyber-attack on the Ukraine power grid is the first known cyber-attack on the grid that succeeded in disrupting electricity service to end users. This attack was almost certainly engineered by a nation-state with the motivation being political in nature. The end result was that the control centers of 3 regional electricity distribution companies were directly affected (3 more experienced intrusion that did not impact operations). A total of 30 substations were disconnected, and approximately 225,000

customers (a total load loss of over 130 MW) lost power for 1 to 6 hours. More than two months after the attack, the control centers were still not fully operational. The attackers had overwritten firmware on critical devices at 16 substations, leaving them unable to respond to remote operator commands. Thus workers had to control these breakers manually.

According to [74], the events unfolded as follows. On December 23, 2015, Prykarpattyaoblenergo (one of the electric utilities in Western Ukraine) detected that power was out in the region's main city, Ivano-Frankivsk. The cause was unknown at first. Then the oblenergo discovered that a third party was able to illegally enter their computer and SCADA systems, using them to remotely disconnect substation breakers. Also, its call center was having technical difficulties due to an influx of calls. Two other utilities Kyivobnergo and Chernivtsioblenergo were attacked within 30 minutes of each other. The attackers disconnected seven 110 kV and twenty-three 35 kV substations for three hours. In order to respond to the attack, operators were forced to switch to manual mode since remote commands were not operational.

This attack utilized a number of sophisticated techniques. They included the following [7]:

- Spear-phishing emails to gain access to the oblenergos' business networks
- BlackEnergy 3 malware embedded in manipulated Microsoft Office documents (Excel and Word) to steal credentials from the business networks
- Modified KillDisk to erase master boot record, select logs, and system events
- Use of the Uninterruptible Power Supply (UPS) system to disconnect load with scheduled service outage
- Telephone DoS attack on the oblenergo's call center to deny outage reporting

Figure 24 illustrates how the attackers infiltrated the utilities' control centers.

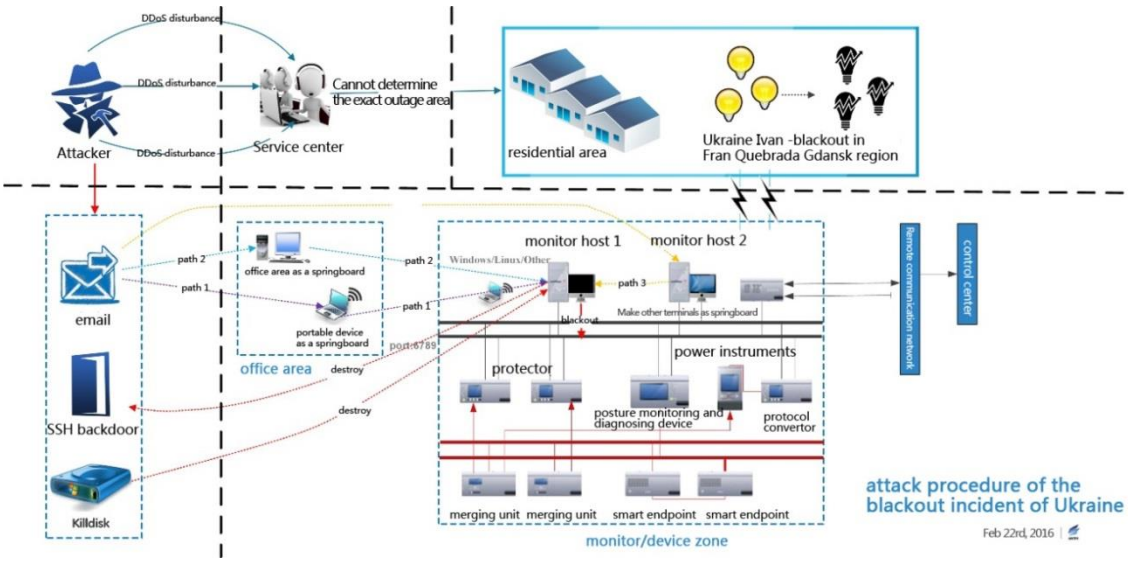

Figure 24. Attack procedure in Ukraine grid cyber-attack [75]

4.4.3 Ukraine (2016)

On December 17, 2016, a second cyber-attack on Ukraine's power grid occurred. This time the target was Ukrenergo (a regional electric utility), and the attack caused one-fifth of Kiev, the capital of Ukraine, to lose power for approximately one hour. It has been speculated that this second attack was a large-scale test of the latest malware dubbed alternately as Industroyer or Crash Override. Unlike in 2015 where the attackers gained access to the utility networks and manually turned off power to electrical substations, the 2016 cyber-attack was entirely automated. Fully automated cyber-attacks can cause blackouts faster, with less preparation, and with less human oversight, making them far more scalable. According to [76], if you estimate that it took more than 20 people to attack three utilities in 2015, that same team could target more than ten utilities at a time in 2016.

It is unclear how Crash Override infected Ukrenergo. ESET (a Slovakian anti-virus firm), one of the two cybersecurity firms that analyzed this incident, suspects that the same type of targeted phishing emails that were used in the 2015 attack may have been used in 2016 as well. Once Crash Override infected Windows machines on a

network, it is able to automatically map out the control system and locate the target equipment. Then it could launch one of four payload modules that communicate with equipment via different industrial communication protocols (IEC 60870-5-101, IEC 60870-5-104, IEC 61850, and OPC DA) [77]. The malware has several features designed to enable it to remain under the radar. For example, communication can be limited to non-working hours, and it can masquerade as the Notepad application. It can also erase crucial system registry keys and overwrite files to make the system unbootable and recovery more difficult. The architecture of this malware is shown in Figure 25.

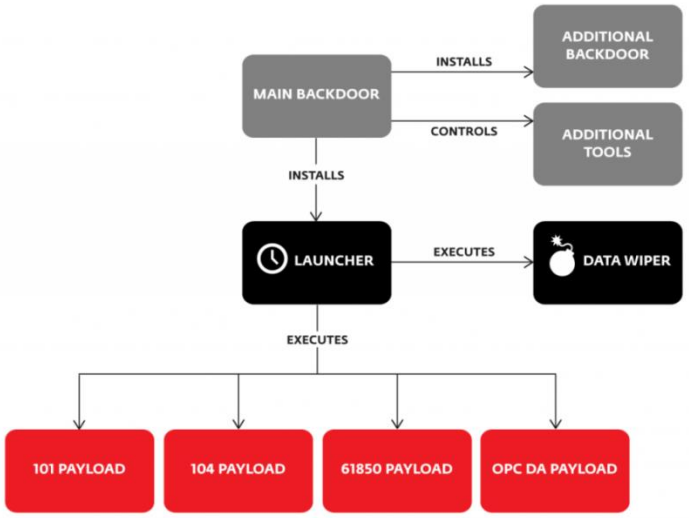

Figure 25. Crash Override malware architecture [77]

ESET suggests that Crash Override has an additional feature that attackers could use to cause physical damage to power equipment; it has a denial-of-service tool that exploits a known vulnerability in the Siemens Siprotec digital relay. By sending the relay a carefully crafted chunk of data, Crash Override could disable the device until it is manually rebooted. In July 2015, Siemens released a firmware update for its vulnerable Siprotec relays and suggests that owners patch their devices. The intent of the feature might be to cut off access to circuit breakers after the malware opens them. If used in combination with a feature that overloads critical components (such as transformers), it

could cause severe damage to that equipment since circuit breakers would be unable to prevent them from overheating.

Dragos Inc. independently verified some of ESET's findings [78]. In their defense recommendations, they suggested that electric utilities:

- Understand where IEC 104, IEC 61850 and DNP3 protocols are used in their systems, look for increased usage of the protocols against known environment baselines, and look for systems leveraging these protocols that have not used them before.
- Prepare incident response plans and perform table top exercises, including practice manual operations to recover compromised SCADA equipment.

They suggested that electric utilities should not:

- Rely on DNP3 as a protection mechanism simply because Crash Override did not have an explicit DNP3 module.
- Rely on air gapped networks, unidirectional firewalls, anti-virus in the ICS, and other passive defense/architecture changes. "No amount of security control will protect against a determined human adversary. Human defenders are required."

4.5 U.S. Response

4.5.1 Government Analysis and Response

In January 2018, the New York Times reported that "A newly drafted United States nuclear strategy… would permit the use of nuclear weapons to respond to… cyber-attacks" [79]. The U.S. government has identified cyber-attacks as a major threat to critical infrastructure, which includes the transmission power grid, and thus national security, so it is considering the use of nuclear weapons as an appropriate response to nation-states that carry out cyber-attacks against U.S. infrastructure. This action

illustrates the seriousness of the threat that cyber-attacks poses to our nation and the extent to which the U.S. government is willing to go to combat them.

In March 2018, the U.S. government issued a technical alert entitled "Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors [80]." In the alert, they discussed different tactics, techniques, and practices used by Russian government threat actors on American victim networks since at least March 2016. There are commonly two categories of victims: staging targets and intended targets. Staging targets are peripheral organizations such as third-party suppliers with less secure networks that are initially attacked so that the attackers can then reach their final intended targets such as organizational networks for government and critical infrastructure. To reach their intended targets, Russian threat actors will use practices that include spear-phishing emails from compromised legitimate accounts, watering-hole domains, credential gathering, open-source and network reconnaissance, host-based exploitation, and targeting industrial control system infrastructure.

The Lockheed-Martin Cyber Kill Chain model adopted by the Department of Homeland Security includes the following phases: reconnaissance, weaponization, delivery, exploitation, installation, command and control, and actions on the objective.

Stage 1 – Reconnaissance: Threat actors accessed publicly available information hosted by organization-monitored networks to look for information on network and organizational design, used compromised staging targets to download source code for intended targets' websites, and attempted to remotely access corporate email/VPN connections.

Stage 2 – Weaponization: Threat actors used Microsoft Office email attachments to obtain a credential hash and then use password-cracking techniques to get the plaintext password. Also threat actors developed compromised staging targets (e.g. trade publications, informational websites on infrastructure, ICS, etc.) into watering holes, where they embedded malicious code into those websites to harvest credentials.

Stage 3 – Delivery: Threat actors used generic PDF documents that contained a shortened URL which led users to a website that prompted them for their email address and password. In some cases, the attachments masqueraded as contracts, resumes, invitations, or policy documents to entice users to open them.

Stage 4 – Exploitation: To capture user credentials, threat actors used malicious .docx files, which connected to a command and control server owned by either the threat actors or a victim. When a user attempted to authenticate to the domain, the server was provided with the hash of the password. Local users received a graphical user interface prompt to enter a username/password, which the server received.

Stage 5 – Installation: Threat actors used compromised credentials to masquerade as authorized users in environments that only use single-factor authentication. Threat actors created local administrator accounts in the staging targets (sometimes disguised as legitimate backup accounts) and placed files within the intended targets. Once inside an intended target's network, threat actors downloaded tools from a remote server.

Stage 6 – Command and Control: Threat actors created web shells on intended targets' publicly accessible web and email servers.

Stage 7 – Actions on Objectives: Threat actors conducted reconnaissance in the intended target's network by identifying and browsing its file servers. They used privileged credentials to access the victim's domain controller. In multiple instances, they accessed workstations/servers on a corporate network that contained data from energy generation facilities. They accessed files related to SCADA systems, such as wiring diagrams or panel layouts. They targeted and copied profile and configuration information for accessing ICS systems on the network. Also they performed cleanup operations afterwards to cover up their tracks.

4.5.2 Industry Analysis and Response

Currently transmission substations are subject to mandatory North American Electric Reliability Corporation's Critical Infrastructure Protection (NERC CIP) cybersecurity standards and physical security standards (e.g. CIP-014-2 – Physical Security), while distribution substations are not (they are governed by individual state public utility commissions so cybersecurity standards and/or measures may vary widely). Hence, a threat actor is more likely to be able to access a distribution substation, which may lack even basic cybersecurity and physical security protection. Although transmission substations are protected by NERC CIP regulation, they are still susceptible to cyber-attacks at the distribution level. In general, a cyber-attack on a distribution substation could open circuit breakers to interrupt power or compromise SCADA operations to cause load instability, but these effects would only have impact on local electricity service. However, if a threat actor can physically access a distribution SCADA master via a substation, which allows access to other distribution and possibly transmission elements in the system, this could have much more serious consequences. Also, connection points between the transmission system and the distribution system, such as step-down transformers, do not always fall under NERC CIP regulation and may present additional cyber vulnerabilities [62].

There are currently eleven NERC CIP sections, ten of which are related to cybersecurity. CIP-002-5.1a covers (bulk electric system) cyber system categorization. CIP-003-6 covers security management controls. CIP-004-6 covers personnel and training. CIP-005-5 covers electronic security perimeters. CIP-006-6 covers physical security of cyber systems. CIP-007-6 covers system security management. CIP-005-5 covers incident reporting and response planning. CIP-009-6 covers recovery plans for (bulk electric system) cyber systems. CIP-010-2 covers configuration change management and vulnerability assessments. Finally, CIP-011-2 covers information protection.

In November 2017, more than 450 organizations and 6,500 individuals participated in NERC's fourth simulated power grid attack exercise known as GridEx (held every two years). The exercise was designed to provide an opportunity for different industry and government stakeholders to respond to simulated cyber and physical attacks. The baseline scenario was composed of a series of four "moves," each lasting four hours (detailed descriptions are provided in Figure 26). It is assumed that coordinated cyber-physical attacks were able to impact operations at bulk transmission and generation facilities, causing outages and some islanding conditions. The simulated cyber-attack options included watering hole, patch deployment, and remote access vulnerabilities that affected utility ICS. From this exercise, the following lessons were learned [81]:

- Lead planners need to be more proactive in communicating with local law enforcement or utility equipment vendors.
- Other critical infrastructure such as oil, natural gas, water, and telecommunications, which may have interdependencies with the power grid, also need to be notified in the event of an attack.
- The E-ISAC portal, which is designed to be a central hub for information, mitigation, and response related to grid cyber-physical attacks, needs to be more effective.
- Utilities need to develop better procedures for external communications in order to alert, educate, and inform customers about outages.
- Utilities need to increase communication resilience in the event that primary communication paths are disrupted so that they can maintain contact with internal personnel, emergency managers, and first responders.
- There needs to be increased participation in the Cyber Mutual Assistance (CMA) program (launched in 2016), which provides a pool of utility cybersecurity experts who volunteer to share their information with other utilities in the event of a cyber-emergency.

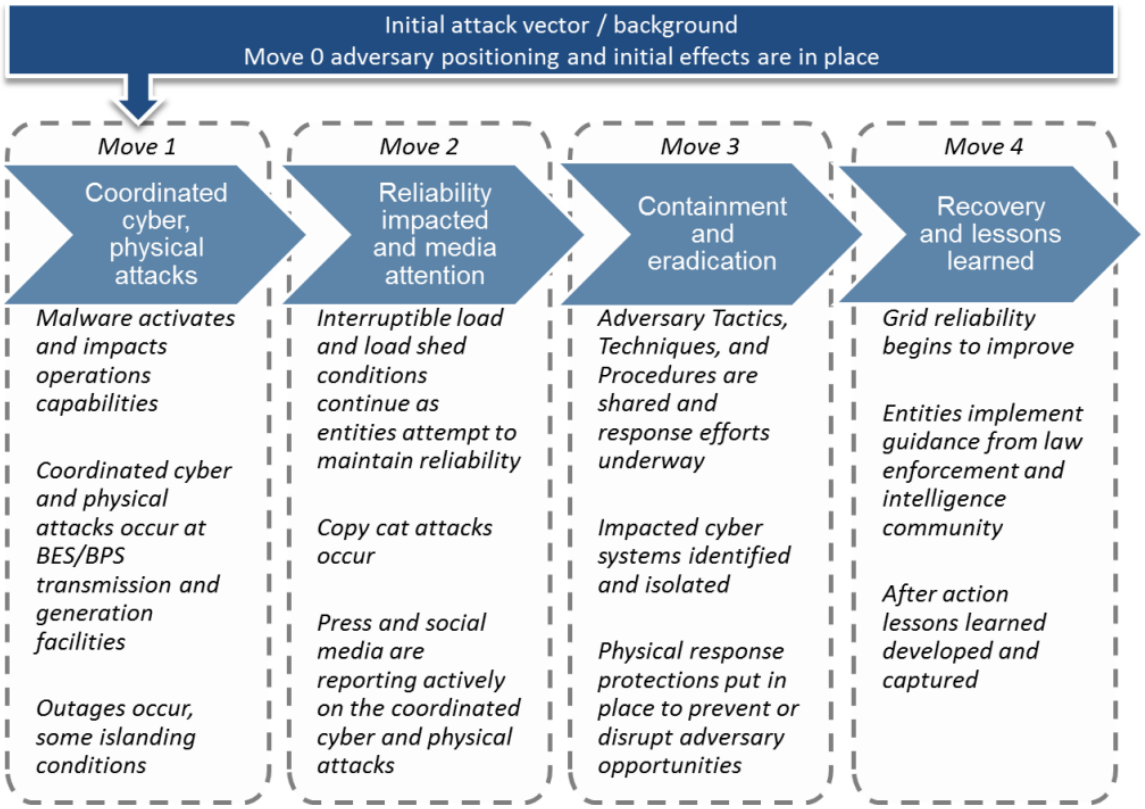

Figure 26. GridEx IV Moves [81]

4.6 Vulnerable Power System Applications

Most of the research literature on power system cybersecurity categorizes cyber-attacks on the grid as attacks that target availability, integrity, or confidentiality. Attacks that target availability, also known as denial-of-service (DoS) attacks, attempt to delay, block, or corrupt system communication. These can be targeted at a variety of communication layers, like channel jamming at the physical layer, spoofing at the MAC layer, distributed traffic flooding at the network and transport layer, and application-layer DoS attacks. Because the grid has stringent communication requirements (3 ms for IEC 61850), an attack that delays a time-critical message can lead to catastrophic results. Attacks that target integrity, such as false data injection (FDI), are aimed at the application layer and try to stealthily modify critical data so that it compromises the

secure operation of the grid. Attacks that target confidentiality eavesdrop on communication channels to acquire unauthorized information, such as customers' account numbers and electricity usage. The survey in [82] describes these three classes of cyber-attacks in greater detail and discusses vulnerabilities of 1) transmission and distribution operations and 2) the advanced metering infrastructure (AMI) and home-area networks for each attack class.

The literature presents a number of arenas in which cyber-attacks can negatively affect the electric grid. Xie et al. demonstrated that an undetectable attack can impact market operations and provide opportunities for financial arbitrage [83], [84]. Sridhar and Govindarasu showed that a data integrity attack can successfully fool power system operators into making the wrong control action via the automatic generation control (AGC) loop by maliciously altering the frequency and tie line flow measurements [85]. That work was extended in [86] to include how to detect anomalies in the ACE and mitigate the effects of such an attack. Sridhar et al. pinpointed other applications with cyber vulnerabilities in [87]. The authors divided these vulnerabilities into three categories: generation, transmission, and distribution. In generation systems, governor control and AGC are potential targets. In transmission systems, the state estimator and FACTS devices are vulnerable. In distribution systems, malicious changes to load shedding schemes are a concern, and smart meters could be targeted.

5 STATE ESTIMATION IN POWER SYSTEM CYBERSECURITY

The work presented in this chapter focuses on investigating and ultimately mitigating the impact of two specific types of cyber-attacks on the power grid. The first type is bad command injection, where the attacker intercepts a legitimate command from the control center and alters it to have a malicious impact on the power system. This type of attack is sophisticated, requiring the attacker to have at least partial knowledge of the system topology and operation. Also, it is undetectable under current power system practices (existing bad data detection tests, such as the largest normalized residual test described in Section 2.4, are incapable of identifying this type of attack). The second type is a network availability attack such as denial-of-service (DoS) on an RTU. This type of attack is less sophisticated and easier to detect, but it can still significantly hamper power system operations and operator situational awareness (as discussed in later sections).

Chapter 5 is organized as follows. Section 5.1 describes the process for traditional power system security assessment. In Section 5.2, we propose a new cyber-physical security assessment (CPSA) methodology and present a new co-simulation environment, which includes an enhanced state estimator, that can model the electrical and communication networks in a cyber-physical smart grid. We also present simulation results for mitigating a bad command injection attack. In Section 5.3, we introduce a novel 3D visualization that could be used to identify how a power system event or cyber-attack evolves over time. In Section 5.4, we present simulation results that demonstrate the devastating effect that DoS attacks could have on the state estimator. We also propose a new methodology for identifying the most vulnerable RTUs, which are good candidates for power system planners and IT staff to focus their resources on securing.

5.1 Traditional Power System Security Assessment

Traditional power system security refers to the ability of a power system to withstand disturbances (also referred to as contingencies) without unduly impacting the availability or quality of its electricity service to loads. Currently, security assessment functions in power system control and operations analyze the vulnerability of the system to a set of potential contingencies on a real-time or near real-time basis. At any given point in time, power system operators need to be prepared for possible contingencies such as the loss of a transmission line, equipment outage, generator failure, a sudden change in load demand, and a change in the system configuration.

The security of the power system can be classified into one of four states: normal secure, normal insecure, emergency, or restorative. The system is in a normal secure state when all system loads are supplied, no constraints are violated, and there are no contingency violations for a set of given contingencies (determined by the balancing authority and regulatory agencies). The system is in a normal insecure state when all system loads are supplied, no constraints are violated, but there are one or more contingency violations for a set of given contingencies. The system is in an emergency state when all of the loads in the system are supplied but one or more constraints are violated. Finally, the system is in a restorative state when there is a loss of load and no constraint is violated. Table 6 summarizes these classifications.

Table 6. Classification of Power System Security States

State	All loads served?	Operational violations?	Contingency violations?
Normal Secure	Yes	No	No
Normal Insecure	Yes	No	One or more
Emergency	Yes	One or more	One or more
Restorative	Loss of load	No	No

The state diagram in Figure 27 illustrates how a power system can transition between these four states, either due to a disturbance or a control action initiated by the operator. Assuming the system starts out in a normal secure state, a disturbance can cause the system to enter a normal insecure state. If preventive control actions are taken in time, the system can return to a normal secure state. If not, a disturbance could potentially push the system into an emergency state. At that point, corrective control actions need to be taken quickly to restore the system to a normal state. However, if they are not taken in time or if the best option is to shed load, then the system enters a restorative state. Once further control actions are taken, the system is finally restored to a normal secure state. Examples of control actions include opening or closing transmission lines and transformers, changing the setting of generation/load, bypassing or putting into service series capacitors, and changing phase-shifter angles.

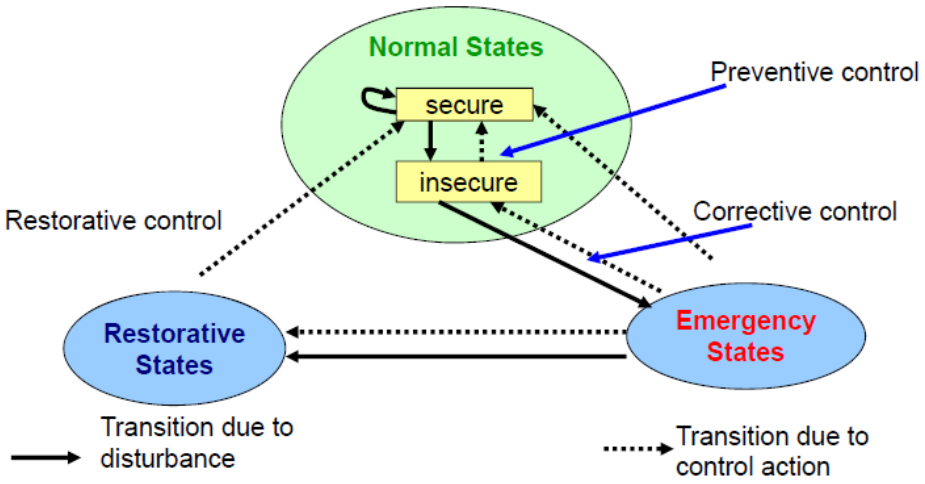

Figure 27. Power system security state diagram

Traditional power system security assessment (illustrated in Figure 28) only analyzes the impact of contingencies on the electrical network. The first step is security monitoring, during which the control center polls substation RTUs to obtain real-time electrical measurements. These raw measurements are then processed by the state estimator, which determines if the system is observable and eliminates or corrects

inconsistent measurements. These filtered measurements are then passed onto the real-time power flow application, which calculates the bus voltage magnitudes, angles, real/reactive injections, and real/reactive line flows. The power flow solution serves as the base case for real-time contingency analysis, which is an automatic what-if assessment of the impact of each contingency on the system. The base case serves as the reference scenario. Then for each contingency, a sequential simulation is run on the reference scenario, the post-contingency real/reactive line flows are calculated, and thermal limit violations are reported before the reference scenario is restored. The results of those violations are visualized using a large-scale display. The entire process runs automatically every few minutes.

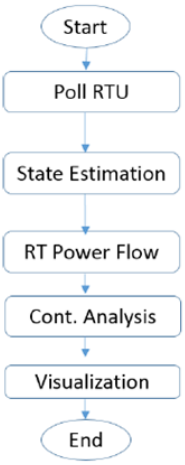

Figure 28. Traditional power system security assessment

5.2 Cyber-Physical Security Assessment

The smart grid is a cyber-physical system whose electrical components and communication system are strongly intertwined. Traditional power system security assessment tools only focus on analyzing the security of the physical power system itself. However, as the grid evolves to become more reliant on information and communication technology, it becomes increasingly important to analyze the cyber-physical interactions between the electrical system and the communication system, rather than focus only on

the electrical system. Hence, there is the need to create a co-simulation environment that is capable of capturing the behavior of the cyber-physical smart grid. Also, there needs to be cyber-physical security assessment tools that can identify when a bad command injection attack is in progress and alert power system operators to the potential impacts on the power system.

5.2.1 System and Threat Models

Assume there exists a cyber-physical smart grid with multiple electrical and communication nodes. Due to limited resources, it is infeasible to physically protect or secure every single node in the cyber-physical system. While conventional cybersecurity measures are available and deployed in the system, they are incapable of detecting data-related cyber-attacks on their own.

For the subsequent work, we assume a centralized power system architecture, where control, operation, and monitoring functions are performed by a single control center. The communication system between the control center and the substation RTUs is a star topology (as illustrated in Figure 29). Bidirectional communication is possible. The control center can poll the substation RTUs for measurements as well as send commands to them. Under normal operations, the RTUs can report equipment statuses and measurements back to the control center. Assume an Intrusion Detection System (IDS) has been mirrored at the connected port of each substation as well as at the control center.

In our scenario, we assume that a malicious actor compromises a legitimate control command issued by power system operators at the control center to a substation RTU. Through a man-in-the-middle attack, the attacker alters the command to be malicious. The RTU executes the malicious command to trip a circuit breaker (for a transmission line, generator, or load). The consequences of this attack are as follows. First, the malicious command has a detrimental effect on the operation of the power system. Secondly, the circuit breaker status at the control center is incorrect and

inconsistent with the actual status. Finally, the incorrect system topology is fed as inputs to downstream EMS applications.

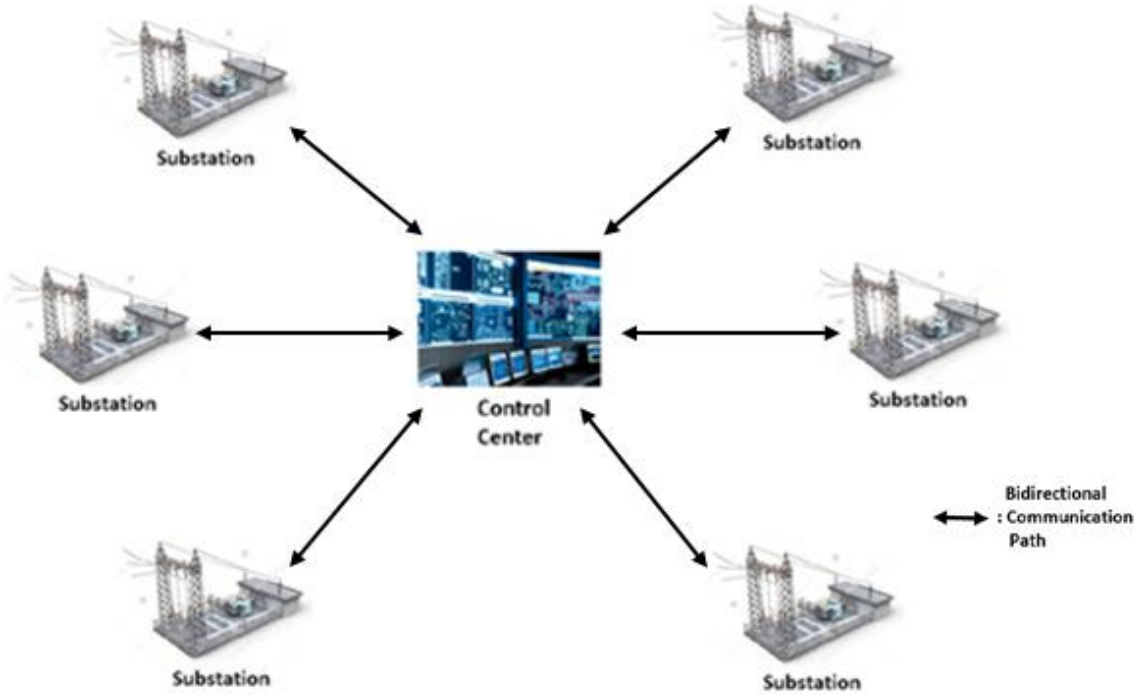

Figure 29. Cyber-physical smart grid topology

5.2.2 Use Case Scenarios

An adversary can perform a bad command injection attack by sending a false control command to a substation RTU. If the adversary has no knowledge of the system and simply injects a false command at random, the power system operator should be able to identify and stop the execution of the malicious command on the power system. However, if the adversary has at least partial knowledge (or full knowledge in the worst case) of the system, they can purposefully inject a specific malicious command to cause maximum damage to the overall system. If the adversary issues a bad command from an illegitimate source, the IDS will be able to detect the malicious nature of the command and prevent its execution on the power system. However, if the adversary has the

capability to model an “intelligent” malicious command, which is issued by a legitimate source and appears to be a routine operation but has undesirable consequences (for example, suddenly reducing power generation by 50%), the IDS most likely will be unable to correctly detect it. However, it could alert the operator by sending a notification of suspicious behavior (considering user-defined threshold values and historical changes in values). One example of a malicious command that could significantly impact the power system is the opening of the circuit breaker connected to the largest generator in the system.

Three specific use case scenarios are described below:

Use Case 1

- Description – Adversary impersonates the network and sends a false (unwanted) but legitimate command outside of the control center to the circuit breaker of the largest generator.
- Effects on the Communication Network – Under this attack, we can observe and monitor several effects on the communication system, such as:
    1) The IDS notifies the operator which command it received. The operator verifies whether the command is legitimate.
    2) A false command is issued to the substation RTU connected to the breaker of the targeted generator.
- Effects on the Power System – If this attack is successful, we can observe the following impacts on the power system:
    1) Insecure operation of the power system
    2) Possible shedding of electrical load
- Use Case Steps:
    1) The attacker targets a bad command as described in the threat model.

2) The IDS detects a suspected malicious command (based on its rules such as IP address, port number, etc.) and notifies the operator. The operator verifies that the control center did not issue this command.
3) CPSA performs power flow and cyber-physical contingency analysis to evaluate the effect of the command on the power system if it was allowed to go through and discovers that the system would become insecure, indicating that the command was malicious.
4) The operator discards the command. Secure system operation is restored.

Use Case 2

- Description – The adversary fabricates or modifies a legitimate command sent to a generator breaker over an insecure network.
- Effects on the Communication Network - Same as Use Case 1, except the legitimate command was modified over the network.
- Effects on the Power System – Same as Use Case 1.
- Use Case Steps:
  1) The attacker sends a command as described in the threat model.
  2) The IDS does not detect the command modification, but still sends a notification to the operator. The operator verifies that the control center did not issue the command.
  3) Same as Step 3 for Use Case 1.
  4) CPSA queries IT personnel for information. IT responds that they suspect a MITM attack.
  5) Same as Step 4 in Use Case 1.

Use Case 3

- Description – The adversary is an insider at the control center, who sends a legitimate but unwanted command to the generator breaker.

- Effects on the Communication Network – The operator receives a command notification from the IDS and finds the transmitted legitimate command was not issued by them. In the worst case scenario, the operator ignores the notification and allows the execution of the command on the power system.
- Effects on the Power System – Same as Use Case 1.
- Use Case Steps:
    1) The attacker sends a command as described in the threat model.
    2) The IDS does not detect the insider attack and notifies the operator that it is a legitimate command. The operator verifies that the received command is the same as what was issued from the control center.
    3) The generator breaker receives a false trip command and trips.
    4) CPSA runs contingency analysis and discovers that the system is in an insecure state, indicating that the command was legitimate but false (unwanted).
    5) CPSA queries IT personnel for information. IT responds that they suspect an insider attack. CPSA prompts the operator to reclose the breaker.
    6) If the breaker does not respond after 20 seconds, CPSA will prompt the operator to initiate the appropriate remedial action scheme, after which secure system operation is restored.

5.2.3 Co-simulator Overview

This section provides an overview description of the novel co-simulation environment used in this work to simulate a bad command injection attack. The design framework of the co-simulator is shown in Figure 30. The co-simulator is composed of three main modules. The first module, which was implemented using a combination of Java APIs (GridSim, Matlabcontrol, and Java Agent DEvelopment Framework, also known as JADE), models the communication network and acts as the simulation driver. The second module (described in greater detail in Section 5.2.4), which was implemented

in MATLAB, is the main computational engine that handles observability analysis, state estimation, power flow, contingency analysis, and cyber-physical security assessment. The third module deals with the topology information as well as visualization of simulation results. The interface between the communication network/simulator driver to the CPSA module is Java to MATLAB. The interface between the CPSA module and the visualization module is MATLAB to PowerWorld, a commercial power system solver/software package.

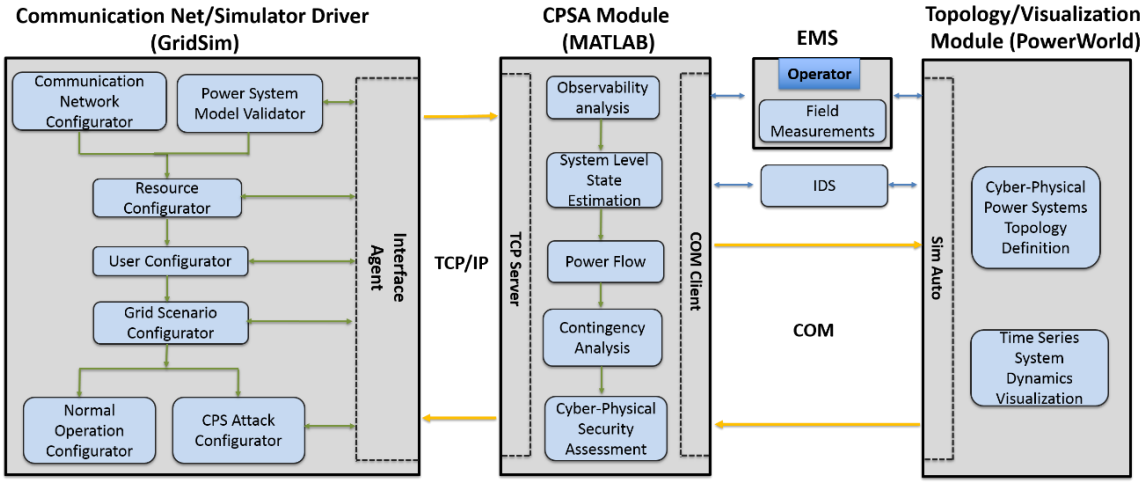

Figure 30. Overall Co-simulator Framework

The sub-modules of the communication network/simulation driver module (which resides in Java) are as follows:

- Communication network configurator – This sub-module is responsible for modeling the communication network between the control center and the RTUs, which includes the communication network topology, number of connected devices in the network, the baud rate, number of packets sent, the control center packet size, RTU packet size, and the propagation delay over the communication channel.

- Power system model validator – This sub-module is responsible for validating at the beginning of each simulation that the electrical network parameters in the co-simulator match the values in the original test case. This step is necessary to ensure that the test case from PowerWorld was not corrupted during the loading process.
- Resource configurator – This sub-module allows the user to configure GridSim resources instantiated during the simulation with the following attributes: number of processors and speed of processing. The resource speed and the job execution time are defined in terms of the Millions of Instructions per Second (MIPS).
- User configurator – This sub-module creates GridSim users (who communicate by sending/receiving passive event objects efficiently) and a simulation instance that contains an application description (list of simulation actions to be processed). For example, each RTU is considered a user, and the control center is also a user. The control center user instance is different from an RTU user instance in that only the control center can execute polling commands. Both the control center and the RTUs implement a time minimization strategy by immediately sending requested data as soon as they receive the request.
- Grid scenario configurator – This sub-module allows the user to select whether they want to simulate normal operations, a single cyber-attack, or a simultaneous combination of cyber-attacks. The result is passed onto either the normal operation configurator or the cyber-physical attack configurator.
- Normal operation configurator – This sub-module handles the event-driven simulation for normal system operations (no cyber-attack).
- Cyber-physical attack configurator – This sub-module handles the event-driven simulation for the system under a cyber-physical attack. The user can

select the duration of the simulated attack, the timestep at which the simulated attack occurs, as well as the type of attack. If a single bad command injection attack is selected, there is an attack modeler which can model a bad command which impacts a transmission line, bus, generator, load, transformer, or shunt capacitor.

The sub-modules of the topology/visualization module (which resides in PowerWorld and is accessible through the SimAuto COM object) are as follows:

- Cyber-physical power system definition – This sub-module provides power system information as input to the co-simulator, which includes power system measurements (real flow, reactive flow, real injection, reactive injection, and voltage magnitude), different parameters for components (such as transmission lines, buses, generators, loads, shunt capacitors, and transformers), and the system topology at the time of data acquisition.
- Time-series system dynamics visualization - This sub-module is responsible for visualizing the time-evolving behavior of the cyber-physical system, both for normal operation as well as when the system is under a cyber-attack. Traditional power system visualizations only display the electrical network while the hybrid cyber-physical visualization in Figure 31 shows the intertwined electrical and communication networks. The typical power system one-line diagram symbols are used to represent electrical components. The orange lines represent electrical connections, while the grey lines with blue squares indicate communication links. The small grey square next to each bus represents a router. Each router is connected to the control center. This novel type of visualization helps power system operators to quickly identify the source of a suspected problem. The time-evolving behavior of the system is illustrated using animation. An alternative approach to visualizing time-evolving behavior is discussed in Section 5.3.

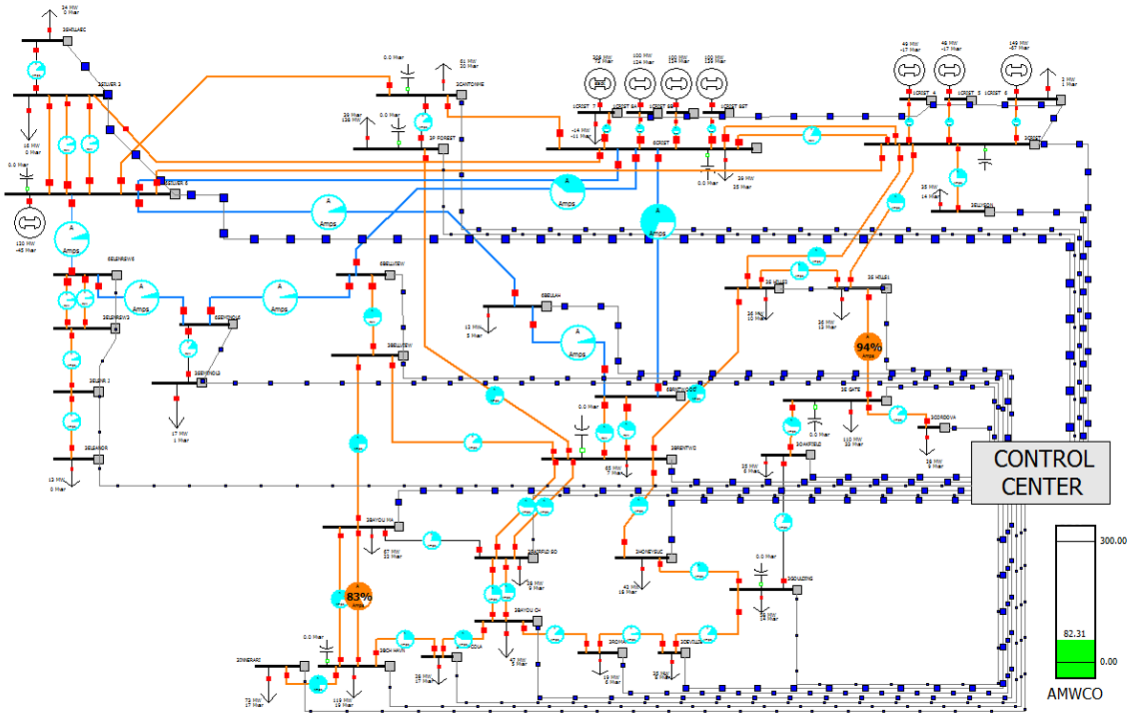

Figure 31. Co-simulator visualization of cyber-physical system

5.2.4 Proposed Methodology

We propose a new cyber-physical security assessment methodology (illustrated in Figure 32b) that considers both the electrical and communication networks. This methodology was encapsulated in the CPSA module and incorporated into the co-simulator described in Section 5.2.3. First we used the co-simulator to model the cyber-physical grid. After the control center polls substation RTUs, the real-time measurements are sent to the IDS. The measurements are passed through to the state estimator, which performs observability analysis, state estimation, and bad data detection to flag inconsistent measurements.

If the IDS suspects a cyber-attack, then it alerts the power system operator using a dialog. For example, assume the IDS suspects a bad command injection attack (also modeled using the co-simulator) on the circuit breaker for a generator. The IDS sends an alert and prompts the operator to simulate the impact of allowing the suspicious

command to be carried out. The CPSA power flow application simulates opening that circuit breaker. Then the CPSA contingency analysis application uses the scenario with the open generator breaker as the reference case. Once the contingency analysis is done, CPSA compares the metrics for the case when the command was rejected versus the case when the command was allowed. A user-defined threshold based on historical information determines the maximum allowable increased insecurity before more verification steps are necessary. If the system is known to be under threat, CPSA initiates modeling of the worst-case attack scenario. Then a large-scale time-evolving visualization of the cyber-physical system is presented to the operator.

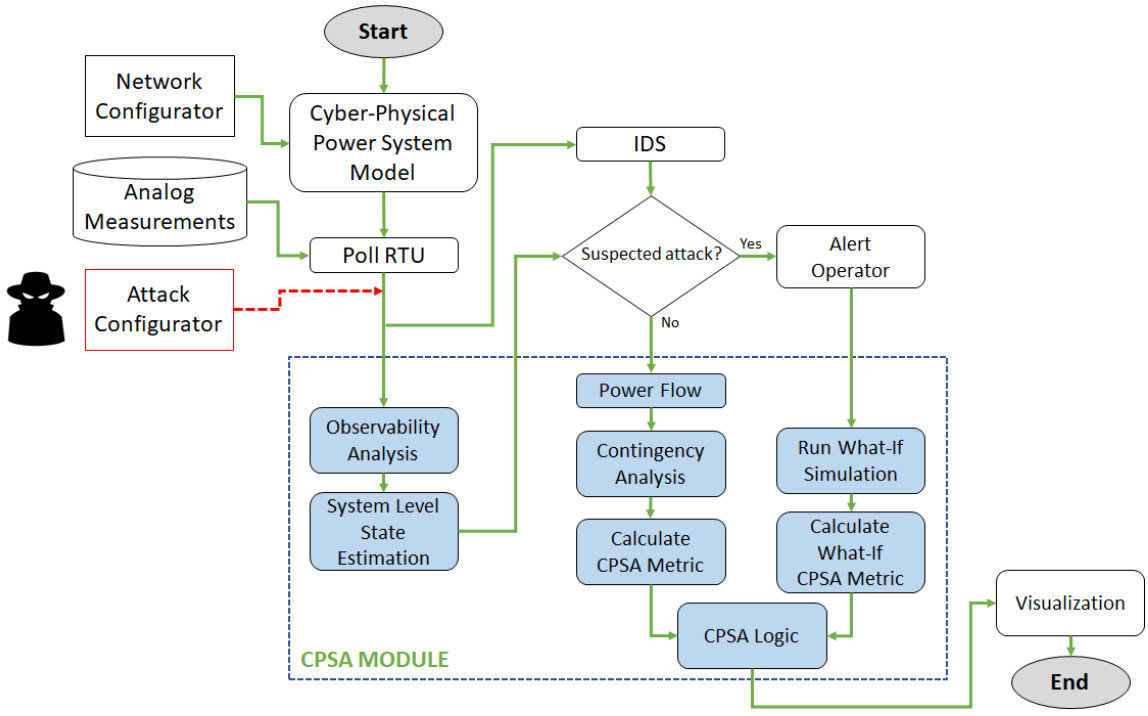

Figure 32. Proposed cyber-physical security assessment

The CPSA module, which encapsulates the new proposed CPSA methodology, runs automatically every few minutes to assess the current health of the cyber-physical system. The sub-modules of the CPSA module (which was implemented in MATLAB) are as follows:

- Observability analysis – This sub-module uses available power system measurements to identify whether the power system is observable (i.e. a unique estimate can be found). It uses the numerical observability analysis algorithm described in Section 2.3.1 to identify observable islands. If the entire system or a part of the system is found to be unobservable, then the worst case scenario is assumed for the unobservable portion(s).
- State estimation – This sub-module uses raw measurement values from RTUs to calculate the most likely system state. It uses the AC weighted least squares solution method described in Section 2.2.1. It uses the largest normalized residual test algorithm described in Section 2.4.1 to attempt to identify inconsistent electrical measurements.
- Power flow – Both legitimate measurements and suspicious measurements are inputs to this sub-module. The solution of the power flow (bus voltage magnitudes and angles) is used to calculate real/reactive transmission line flows and real/reactive bus injections. These results serve as the reference (pre-contingency scenario) for the next sub-module.
- Contingency analysis – This sub-module generates a list of cyber-physical contingencies. Then, using the power flow results from the previous sub-module as the reference scenario, each cyber-physical contingency from a user-defined set of contingency definitions is simulated on top of the reference scenario to evaluate the potential impact on the cyber-physical system. The worst contingencies (above a user-defined threshold) are identified and flagged for the power system operator.
- Cyber-physical security assessment – This sub-module assesses the cyber-physical security of the system using inputs from other sub-modules. First, the IDS sends an alert to this sub-module if it suspects a malicious command has been sent. Next the contingency analysis sub-module provides the outcomes

of its what-if scenarios. Then this sub-module computes a metric known as the system aggregate megawatt contingency overload (sysAMWCO), which is defined below, for the simulated scenario where the suspicious command is rejected as well as the simulated scenario where the suspicious command is allowed to execute. By comparing the difference in sysAMWCO, the power system operator can evaluate whether the suspicious command is actually malicious. Large sysAMWCO values indicate more system insecurity, while low sysAMWCO values indicate more system security. The suspicious command is rejected if the sysAMWCO difference is greater than a user-defined threshold (established using historical information and operator experience), and it is allowed to execute if the sysAMWCO difference is less than the threshold.

The sysAMWCO metric is defined as a function of the aggregate megawatt contingency overload (AMWCO), which is in turn a function of the aggregate percentage contingency overload (APCO).

$$APCO_{branch(j,k)} = \sum_{\substack{Contingencies\ that \\ overload\ branch(j,k)}} (\%\ Overload - 100) \tag{150}$$

$$AMWCO_{branch(j,k)} = APCO_{branch(j,k)} \times MVA\ Rating_{branch(j,k)} \tag{151}$$

$$sysAMWCO = \sum_{\forall\ branch(j,k) \in system} AMWCO_{branch(j,k)} \tag{152}$$

Although the primary use case for the CPSA module in the co-simulator is real-time analysis that is performed automatically every few minutes, it can also be used for faster-than-real-time analysis. For example, future predictive information such as load forecasts and renewable generation output forecasts (e.g. for the next 30 minutes) could be used as inputs into the co-simulator. Assuming the predictions are available in 1-

minute intervals, the co-simulator can evaluate the cyber-physical security of the system for the next 30 timesteps. The time required for the co-simulator to run is only a few minutes so we can effectively analyze the system in faster-than-real-time. Once that simulation is complete and the forecasts for the next 30 minutes are available, we can use the co-simulator to assess the cyber-physical security of the system during that window of time.

5.2.5 Simulation Results

We tested the proposed CPSA methodology on a 42-bus equivalent system of a real power grid in the United States. The electrical network topology of the 42-bus system is shown in Figure 33. The system consisted of 24 substations, 45 transmission lines, 16 transformers, 8 generators, and 27 loads (approximately 902 MW).

The communication network topology of the test system is shown in Figure 34. The communication system consists of two routers, a link from the control center to the first router, a connection between the two routers, and 24 links to each of the 24 substations.

The graphical user interface for the co-simulator is shown in Figure 35. The user inputs system simulation parameters into the GUI, which are then fed into the communication network/simulation driver module of the co-simulator (described in Section 5.2.3). For this test system, the baud rate was 1572864 bits/second, the propagation delay was 300 ms, the packet buffer size at the control center was 180 bytes, and the packet buffer size at the RTU was 1500 bytes. The propagation delay was estimated based on the number of bits transmitted in a packet (2553 kbit) over the network, which has a speed of 8.51 Mbps.

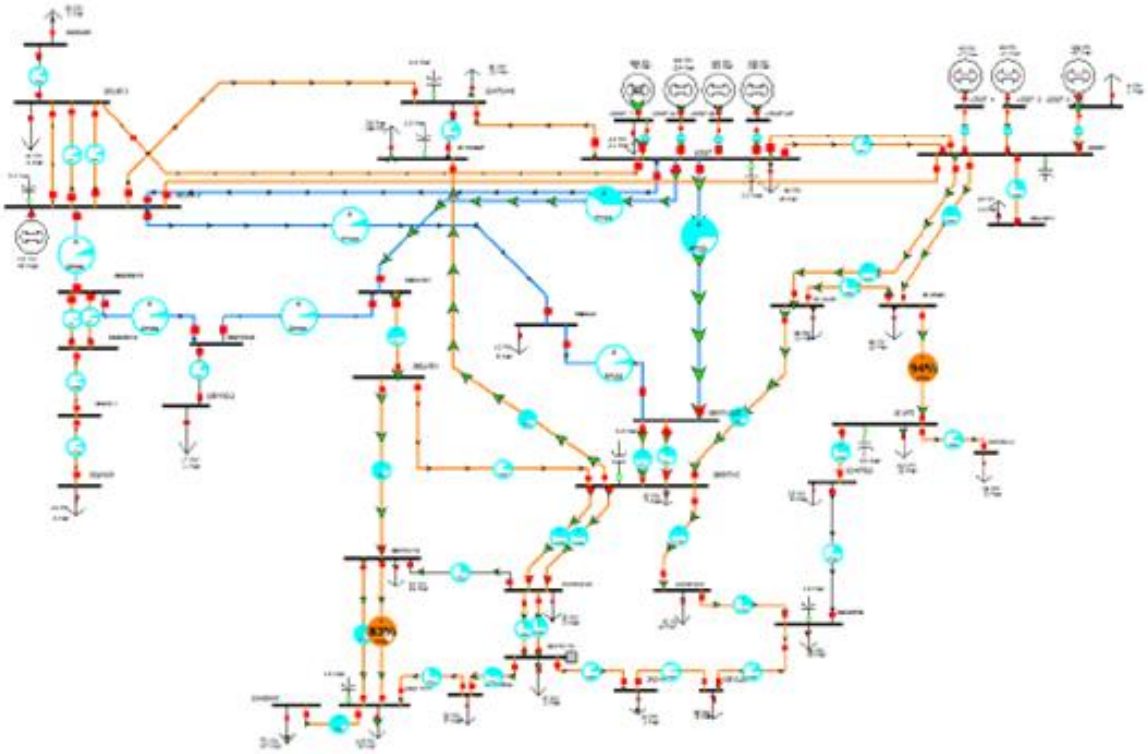

Figure 33. Electrical network topology for 42-bus test system

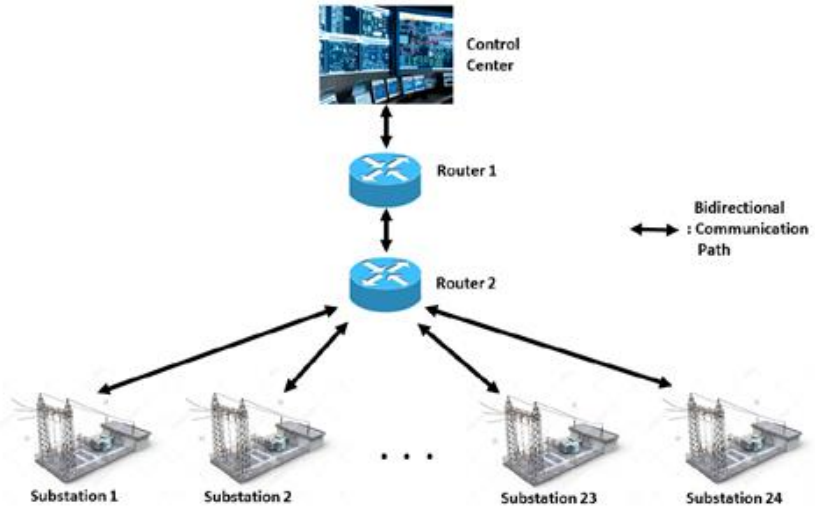

Figure 34. Communication network topology for 42-bus test system

Figure 35. Co-simulator GUI

The cyber-physical system operates normally from timestep $t = 1$ to $t = 4$. At $t = 5$, the adversary sends a command to one of the RTUs to open the largest generator. The IDS suspects that this command is not legitimate based on its rules. Hence, the IDS sends an alert to the operator at the control center. The operator decides to simulate the suspicious command and observe its effect on the power system. First, the operator uses the co-simulator to simulate the normal operation scenario (i.e. the command is not allowed to execute) given load and renewable generation forecasts. Then the operator simulates the scenario if the command is allowed to execute on the system.

These results are shown in Figure 36. The load profile is assumed to be the same for both the normal operation scenario and the scenario where there is a bad command injection attack (the attack only targeted generation). If the command is rejected, then the sysAMWCO at $t = 5$ would be approximately 200, and the sysAMWCO at $t = 6$ would be approximately 170. However, if the command was allowed to execute on the system, the sysAMWCO at $t = 5$ would increase to approximately 300, and the sysAMWCO at $t = 6$ would increase to approximately 225. That would be an increase of approximately 50% and 32%. The user-defined maximum allowable threshold was only 15% above the normal operation case, so the operator rejects the suspicious command and flags it as malicious for the IT department to investigate and take further action.

In addition to presenting the sysAMWCO plots to the operator, the co-simulator also provides a large-scale dynamic visualization to the operator that uses animation to show the time-varying behavior of the system under the bad command injection attack. When the system was in normal operation from $t = 1$ to $t = 4$, this visualization appeared with no electrical network violations and no communication link issues (see Figure 31). However, at $t = 5$, this visualization changes to indicate thermal limit violations on multiple transmission lines in the electrical network, and the communication links that are compromised are highlighted using a flashing animation (see Figure 37). Also, the sysAMWCO plot on the right side of the visualization changes from green to red when

the sysAMWCO metric is more than 85% of a user-defined threshold. This color change immediately draws the operator's attention and aids in quick identification of the source of the problem.

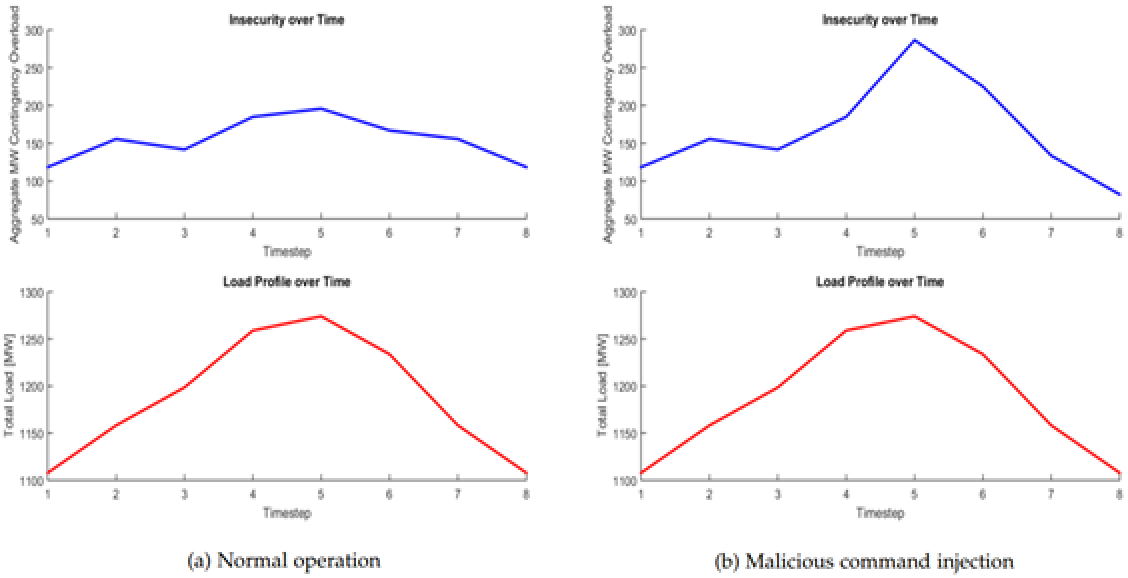

(a) Normal operation (b) Malicious command injection

Figure 36. Detection of a bad command injection attack by comparing the normal operation sysAMWCO against the sysAMWCO if the command is allowed to execute

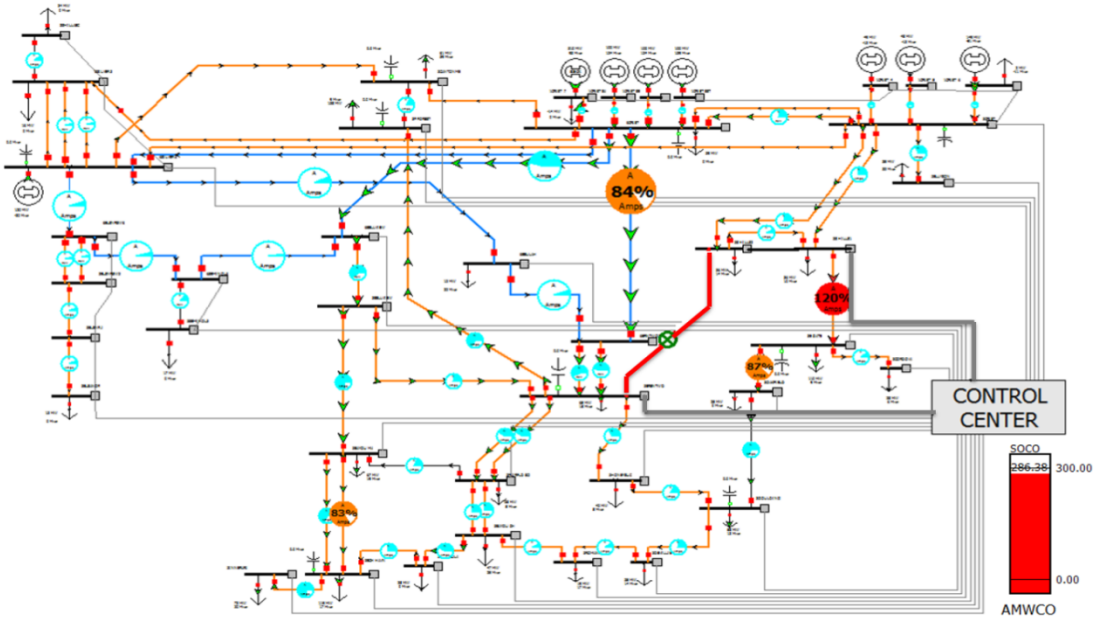

Figure 37. Co-simulator visualization of cyber-physical system

5.2.6 Conclusions

This work proposed a new cyber-physical security assessment methodology. Unlike traditional power system security assessment, it considers both the electrical and communication system. The methodology was tested on a 42-bus realistic power system using a co-simulator that was built in Java, MATLAB, and PowerWorld. The cyber-physical security of the system was assessed using a metric known as the system AMWCO. By comparing this metric for a normal operation scenario and the scenario where the suspicious command was executed, the operator can determine if the execution of a suspicious command will lead to system-wide instability. A large-scale dynamic visualization that illustrates the time-varying behavior of the system under attack aids the operator in quickly identifying the source of the problem.

5.3 Spatiotemporal Power System Visualization

In the work presented in Section 5.2, the co-simulator visualization used animation to illustrate dynamic time-varying power system behavior. This visualization method is commonly used in commercially available power system analysis tools, but it has several drawbacks. If the power system operator needs to reference the system topology while studying the time-varying behavior, they would need to replay the animation multiple times in order to gather the necessary information. Also, research has shown that static snapshots are generally easier for users to navigate rather than animations. We need a new visualization method that would allow the operator to both see the system topology while analyzing the time-varying behavior.

5.3.1 Introduction

Power system operators, planners, and managers need to quickly analyze large complex datasets in order to make effective decisions for the bulk electric grid. The

challenge is to present the relevant information in a manner that facilitates intuitive and rapid assessment of the system state [88]. To mitigate the information overload experienced by grid operators and decision makers, it is necessary to develop a visualization platform capable of quickly transforming large datasets into visual representations that are easy to understand at a glance. Of special importance is the ability to analyze temporal (both look-ahead and past) data. Look-ahead capabilities are crucial for understanding the emerging stochastic nature of the grid due to the integration of non-dispatchable renewable energy, demand response, and storage. These capabilities can be used by operators to assess the likelihood of possible future problems in the network as well as pinpoint how an event began.

Most of the existing work has focused on representing power system data for only a single point in time. This type of static large-scale 2-D visualization does not support the aforementioned requirements for temporal analysis. One way to fill this void is a visualization that allows users to see multiple time steps simultaneously. We propose a 3D approach that only renders the relevant time-varying data on top of a one-line diagram and allows navigation of various timesteps. This type of visualization could be used to show when the power system has entered an insecure state, either physical or cyber-physical. For example, it can pinpoint when problems such as thermal limit violations begin to emerge during a power system event. Also, it could be used to indicate when parts of the network become unobservable from loss of measurements, either due to instrumentation/communication error or from a cyber-attack.

5.3.2 Literature Review

Various visualization methods have been developed in the past to aid in the interpretation of power system data, although only some have made their way into the hands of practicing engineers. These techniques include animated flow arrows [89], bus voltage and transmission line contours [90], and 3D bar graphs [91], [92], which can all

be found in the commercial software package PowerWorld [93]. Using time sliders to explore time-varying networks is generally more effective than animation. Farrugia and Quigley [94] conducted studies to compare two common approaches: animation and static snapshots. They concluded that static snapshots are generally more effective in terms of task completion time. To tackle the limitation of the snapshot-based approach, DiffAni [95] integrated several interaction techniques into a snapshot-based visualization.

Recent research has investigated the use of three-dimensional space for visualizing time-varying networks. Itoh et al. [96] stacked 2-dimensional planes on a horizontal axis to represent time- varying networks. Tominski et al. [97] presented a similar approach for visualizing spatiotemporal data, but they used a vertical axis for time. We chose to apply this stacking approach to visualizing the evolution of power system states. Previous work has also attempted a hybrid approach of using 3-D and 2-D. MatrixCube represents a time-evolving graph as a 3D cube by stacking adjacency matrices. Each graph can be represented as a matrix [98]. This type of cube allows users to easily explore graph in various perspectives using simple operations, such as filtering by slicing the cube (e.g., show only for 3-6pm).

5.3.3 Proposed Approach

In this section, we propose a novel 3D spatiotemporal visualization prototype. Power system operators need the ability to simultaneously see the system topology as well as time-evolving trends in the power system. Our proposed solution is to impose data volumes on top of the system one-line for time-evolving system quantities such as bus voltages and line thermal limits. The same technique could be extended to generation or load levels. The major goal of this approach is to facilitate fast investigation and exploration of dynamic behavior in a power system. Because the changes in state are summarized by shape and color, the overall state of the system can be understood by the user without the need to read individual values. Many contemporary power system

visualization tools such as PowerWorld [93] only employ 2D or 3D layouts that summarize exact values at an instantaneous point in time. Usually these tools have to leverage animation to show how the system changes over time. Often comparison of system-wide behavior at two time points is challenging, and comparison of three or more time instances is nearly impossible. This is one of the natural limitations when using animation to show trends over time.

The initial prototype of this proposed solution is presented in Figure 38. The one-line, shown as a semi-transparent floor plan in the middle of the screen, is used to indicate the present time. The vertical z-axis indicates different system states over time. The +z direction represents future predicted system states. A point higher along the +z axis indicates a time further into the future. The –z direction represents historical system data. A point lower on the –z axis indicates a time further in the past. The stacked cylinders represent line thermal limits in the network at different time points. The size and color indicate the severity of the problem. The redder the cylinder and the larger its radius, the more severe the line overload is.

This 3D representation gives users the ability to view system conditions across a window of time so that they can determine whether a problem is present in the system, as well as when it began and when it is expected to end, which may change as the problem evolves. The interactive nature of this visualization allows the user to navigate between system states one step at a time. To help distinguish between the current time and any past or future states that the user is navigating, the cylindrical slice corresponding to the current time is separated from the overall stack.

The heat map on the left panel shows the user more detailed information about the network. Each column along the horizontal dimension shows an interval of time, the length of which depends on the use case. Each row indicates the state of a different power system quantity in the network, such as bus voltage or line limit. The color shows whether the value of the attribute at a specific time falls within the normal range.

Accounting for human factors, green was chosen for normal operation while red indicates that urgent action is required. Also, a small white arrow points to the present time.

A second more sophisticated visualization prototype is presented in Figure 39 and Figure 40. It was implemented using WebGL, a web browser based 3D rendering environment. WebGL was selected for its performance and its compatibility. By leveraging graphics hardware at a low level, it offers excellent 3D performance and allows graphically complex scenes to run smoothly, enabling real-time interaction and exploration. WebGL allows 3D environments with rich user interaction directly in a web browser. Currently most mainstream web browsers (e.g., Chrome, Firefox, Internet Explorer 11, Safari) can support it. By choosing WebGL, our 3D visualization can run on many systems with a wide variety of graphics hardware with very little additional setup.

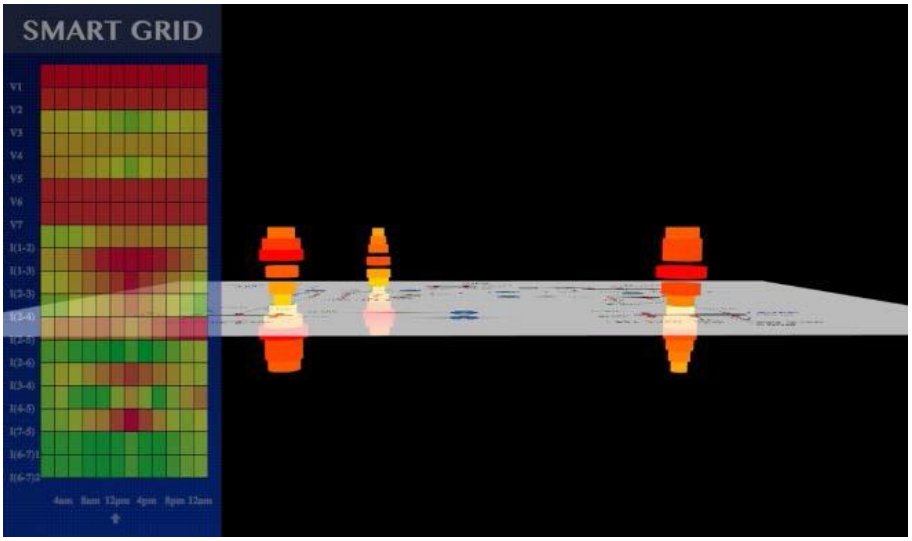

Figure 38. Initial 3D prototype representation

Each bus in the system is visualized as a vertical volume. The volumes are created by modifying the radius of a vertical column based on the desired quantity to be visualized at each time step, as shown in Figure 39. Transmission lines are visualized using colored lines that run between the buses they connect. The color of the line is based on the transmission line flow value. As the loading increases, the transmission

lines become more and more visible. In addition to the change in visibility, the color of the line also changes from grey to dull red once the thermal loading surpasses 90% of its rated limit. As the line flow value increases, the red color becomes brighter until it transforms into a stark vivid red at above 105%. When the percent thermal limit violation is very low, the associated line is not drawn. We chose to hide the lightly loaded lines by default, because they generally operate as expected and are less critical for the operators to notice than the lines that are near or above their maximum thermal limits. The lines in Figure 40 illustrate the changes in color and opacity as the system evolves over time. In this scenario, the shown line was over capacity for the entire study time. Currently the prototype supports panning, rotation, zooming, highlighting, and click-hiding [99].

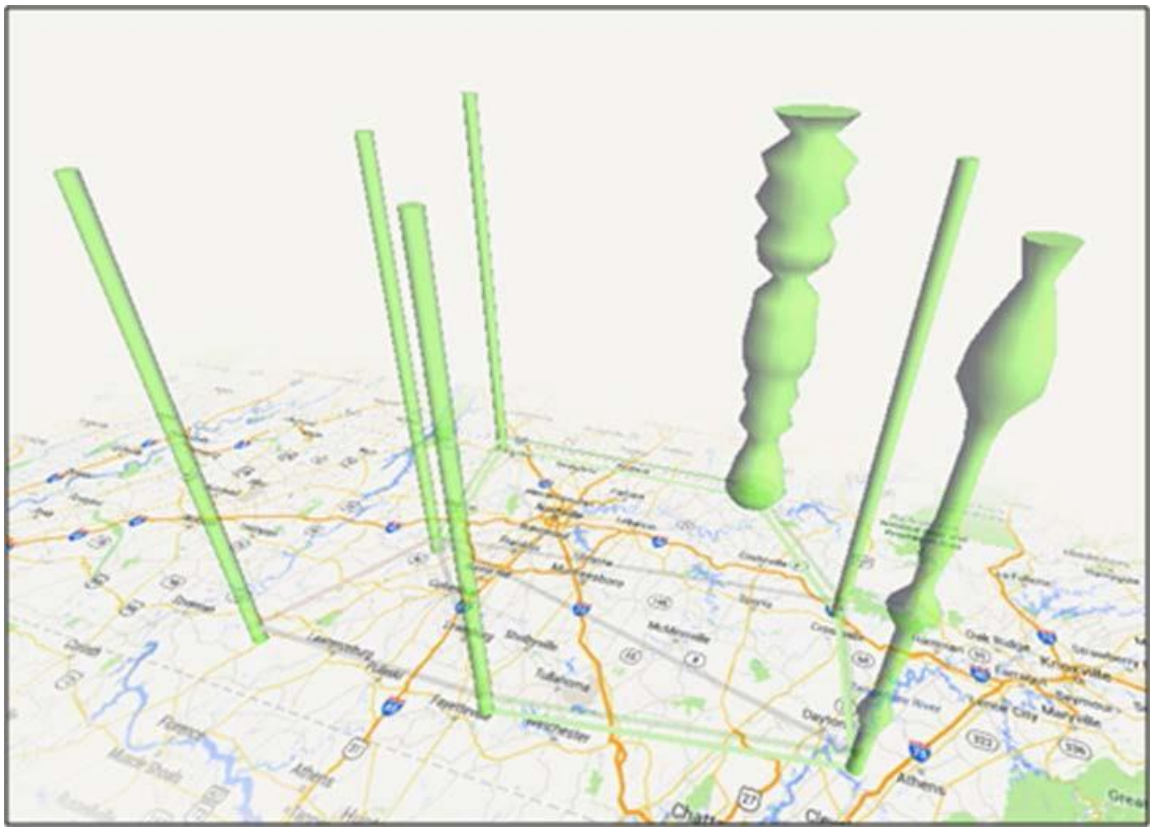

Figure 39. Visualization of bus voltages for a 7-bus test system

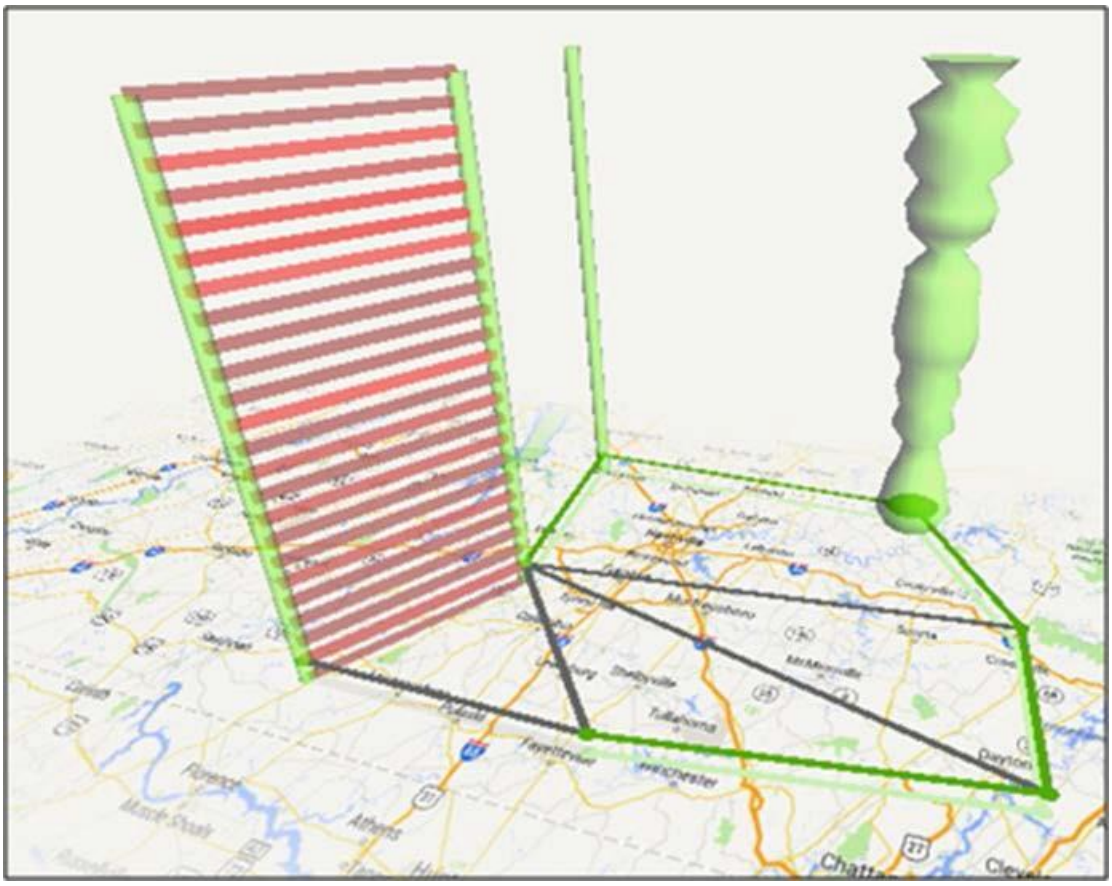

Figure 40. Visualization of line thermal limits for a 7-bus test system

5.3.4 Conclusions

This work proposed a novel 3D stacking visualization as an alternative to the animation-style visualizations currently used in commercial power system analysis tools. It has the benefit of allowing the power system operator to simultaneously view the power system one-line diagram while seeing time-varying behavior in the system. The ability to hide irrelevant data prevents the power system operator from being overloaded with too much information.

5.4 *N*-1 RTU Cyber-Physical Security Assessment Using State Estimation

Real-time SCADA systems use RTUs to monitor and manage the flow of power at electrical substations. As their connectivity to different utility and private networks

increases, RTUs are becoming more vulnerable to cyber-attacks. Some attacks seek to access RTUs to directly control power system devices with the intent to shed load or cause equipment damage. Other attacks (such as denial-of-service) target network availability and seek to block, delay, or corrupt communications between the RTU and the control center. In the most severe case, when communications are entirely blocked, the loss of an RTU can cause the power system to become unobservable. It is important to understand how losing an RTU impacts the system state (bus voltage magnitudes and angles). The system state is determined by the state estimator and serves as the input to other critical EMS applications. There is currently no systematic approach for assessing the cyber-physical impact of losing RTUs. This paper proposes a methodology for *N*-1 RTU cyber-physical assessment that could benefit power system control and operation. We demonstrate our approach on the IEEE 14-bus test system as well as on a synthetic 200-bus system.

5.4.1 Introduction

Electric utilities use real-time SCADA systems to monitor and control their assets. These systems were originally designed for safety and reliability, so there is an implicit assumption of trust of all system components and communications. This implicit trust makes them vulnerable to cyber-threats posed by malicious actors. As its connectivity to different networks and systems increases, the power grid grows increasingly more susceptible to cyber-attacks.

A typical SCADA system includes devices such as RTUs, programmable logic controllers (PLCs), intelligent electronic devices (IEDs), and relays located in electrical substations. RTUs are field-based devices that are used to monitor and manage the flow of power at a substation. PLCs, IEDs, and relays are used to automate tasks at each substation. A control center with central computers manages the remote substation equipment. The control center also processes, analyzes, and archives the collected real-

time information. A communication system is used to convey that information from substation RTUs back to the control center as well as send commands from the control center to substation RTUs, using a protocol such as DNP3. Once they arrive at the control center, raw system measurements are sent to the state estimator, which processes them to determine the most likely system state (bus voltage magnitudes and angles). The SE solution serves as the input to downstream EMS applications, such as contingency analysis and optimal power flow. Hence, it is key that the network is observable (i.e. a unique SE solution can be found) and that the solution is accurate.

The research community has long been aware of the potentially devastating impact of a cyber-attack on the state estimator. In 2009, Liu et al. introduced the false data injection (FDI) attack, where an attacker can introduce arbitrary state variable errors without being detected by existing bad data detection algorithms by manipulating measurements [100]. Since then, there has been a growing body of literature focusing on stealth deception attacks on the state estimator [83], [84], [101]. However, there is not as much work on network availability attacks (NAA), such as denial-of-service (DoS). Zhang et al. formulated an optimal DoS attack strategy considering an energy constraint [102]–[104], but it was not applied to power systems. Liu et al. studied the effect of a DoS attack on load frequency control [105]. Vukovic and Dan considered a DoS attack on a fully distributed state estimator as well as detection/localization of the attack [106]. None of these works focus on systematically studying the impact of losing RTUs.

In this work, we present a novel methodology for RTU cyber-physical security assessment. For simplicity, we assume that only one RTU is under attack at a given time (N-1), although this approach could be generalized to the loss of multiple RTUs. When an RTU is under a NAA, the measurements from that RTU will be unavailable to the state estimator. In some cases, the system may become unobservable, and the solution will contain gross errors. During the loss of telemetry from a cyber-attack, the estimated system state could diverge significantly from the actual system state, endangering the

grid's secure operation. Offline studies using our proposed approach could help power system planners and IT staff to identify the most critical RTUs and use that information to build a system that is more resilient against NAA. Online studies using the same approach would detect when the system is vulnerable to the loss of an RTU.

5.4.2 System and Threat Models

We assume that the power system has a centralized architecture with a single control center that receives and processes telemetry information from the field. We also assume there is a star communication topology between the control center and the substation RTUs (i.e. the substations can communicate with the control center but not with each other). We assume that the attacker is not an insider and thus has no specialized knowledge of the system topology/architecture but is capable of attacking an RTU via either physical or remote access.

5.4.3 Methodology

The objective of this work is to present a methodology to identify and rank vulnerable RTUs in an electrical power system so that utilities can bolster the security of these critical assets in order to minimize the impact of a potential cyber-attack. The idea is comparable to *N*-1 power system contingency analysis. However, traditional power system contingency analysis only considers the loss of physical elements such as a generator or transmission line. For our analysis, we focus on the loss of communication between an RTU and the control center, meaning the measurements associated with that RTU will be unavailable to the state estimator. We quantify the severity of that loss based on the number of observable islands created, the number of unobservable buses, and the estimation errors. (Refer to Section 2.2.1 for the weighted least squares state estimation algorithm and Section 2.3.1 for the nodal variable-based algorithm for numerical observability analysis.)

Our proposed algorithm is as follows:

Step 1. Assess the normal operation (no attack) scenario. Identify the observable islands and the unobservable buses.

Step 2. For each RTU, simulate the loss of the measurements connected to that RTU. Perform observability analysis (OA) and SE.

Step 3. If the global slack bus is connected to one of the unobservable buses, assign the slack to another bus.

Step 4. Record the observable island groups, the unobservable buses, the SE solution, and the SE solution errors in polar coordinates. The SE solution is defined as follows. The voltage at bus n in polar coordinates is $V_{n,SE} = |V_{n,SE}|\angle V_{n,SE} = x_{n+N}\angle x_n$, after remapping x (refer to Section 2.2.1) to include the slack bus angle which is 0.

Step 5. Restore to the no attack scenario. If not all included RTUs are processed, return to Step 2.

Step 6. To rank the severity of losing each RTU, we propose a straightforward metric, which is the time-average of the L2-norm of the difference between the SE voltage vector $V_{SE} \in \mathbb{C}^N$ (where N is the number of buses in a given power system) from Step 4 and the known power flow voltage $V_{PF} \in \mathbb{C}^N$:

$$\text{avg}\left\|V_{diff}\right\|_2 = \frac{\sum_{t=t_{start}}^{t_{start}+T_{attack}} \|V_{SE}(t) - V_{PF}(t)\|_2}{T_{attack}} \tag{153}$$

5.4.4 Simulation Results

We generated a realistic load profile and added random Gaussian noise to synthetic measurements created from the test systems' power flow results. Based on [19], we assumed the line flow, power injection, and voltage magnitude measurement error $\sigma = [0.008, 0.01, 0.004]^T$. For the IEEE 14-bus system, we used a graph partitioner to automatically divide the system into 4 substations and assign the appropriate

measurements to each substation RTU [33]. For the larger and more realistic Illinois 200-bus system, we used the existing substation assignments. We assumed every measurement (real and reactive line flows, real and reactive bus injections, and bus voltage magnitudes) is available to the state estimator when the system is in normal operation. When an RTU is simulated as being under attack, the measurements for that RTU are removed from the set of measurements that serve as the input to the state estimator. We chose a typical convergence tolerance ε of 10e-4.

IEEE 14-Bus Test System

The IEEE 14-bus test system contains 14 buses, 20 branches, 5 generators, and 11 loads. We assume there are 4 RTUs. The cyber-physical one-line diagram for the system is shown in Figure 41. We assume RTU 1 collects measurements for buses 9, 10, 13, 14; RTU 2 for buses 5, 6, 11, and 12; RTU 3 for buses 1, 2, and 3; RTU 4 for buses 4, 7, and 8. This assignment is shown in Table 7. The simulation considers 30 timesteps. The load profile is randomly generated and changes from t = 0 to t = 30. The system operates normally from t = 0 to t = 4. From t = 5 until t = 24, a simulated attack is performed on each RTU. The normalized load and attack profiles are illustrated in Figure 42.

Table 7. RTU Assignment for IEEE 14-Bus System

RTU #	Bus #
1	9, 10, 13, 14
2	5, 6, 11, 12
3	1, 2, 3
4	4, 7, 8

When the system is operating normally (no attack), bus 1 is the system slack, and OA shows that all buses belong to the same observable island (as shown in Figure 43), i.e. the entire system is observable. The measurement graph for this scenario is shown in Figure 44. The solid blue lines indicate that all measurements are available. Bus 1 is the slack bus for the entire system and is highlighted in green. When an RTU is under attack,

the system is no longer fully observable (shown in Figure 45). Through our analysis process, we can identify which RTU loss has the biggest impact on the state estimator. For example, if RTU 1 is under attack (see Figure 46a), the measurements associated with buses 9, 10, 13, and 14 become unavailable (illustrated using dashed lines). Bus 1 is the global slack bus (illustrated as a green dot). When the measurements are lost, the system splits into two observable islands. One island contains buses 1-13, and bus 14 is its own island (illustrated as a black square) since there is not enough measurement information to deduce its behavior.

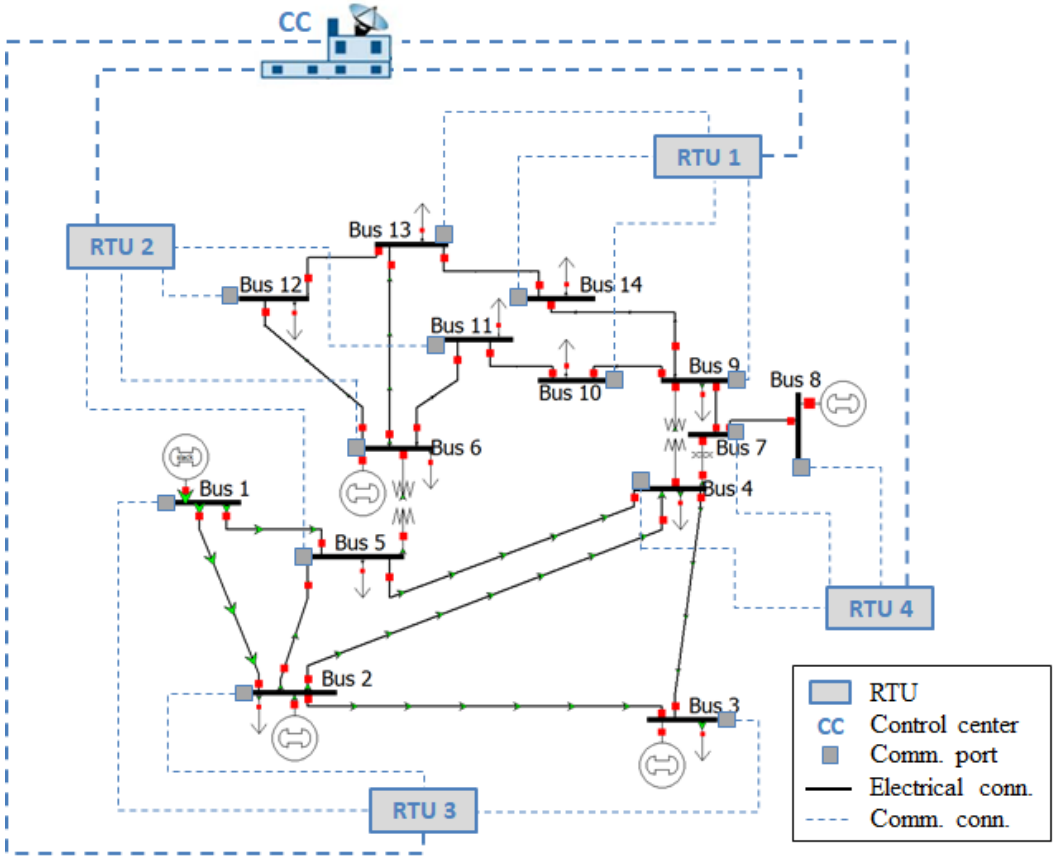

Figure 41. Cyber-physical one-line diagram for IEEE 14-bus test system

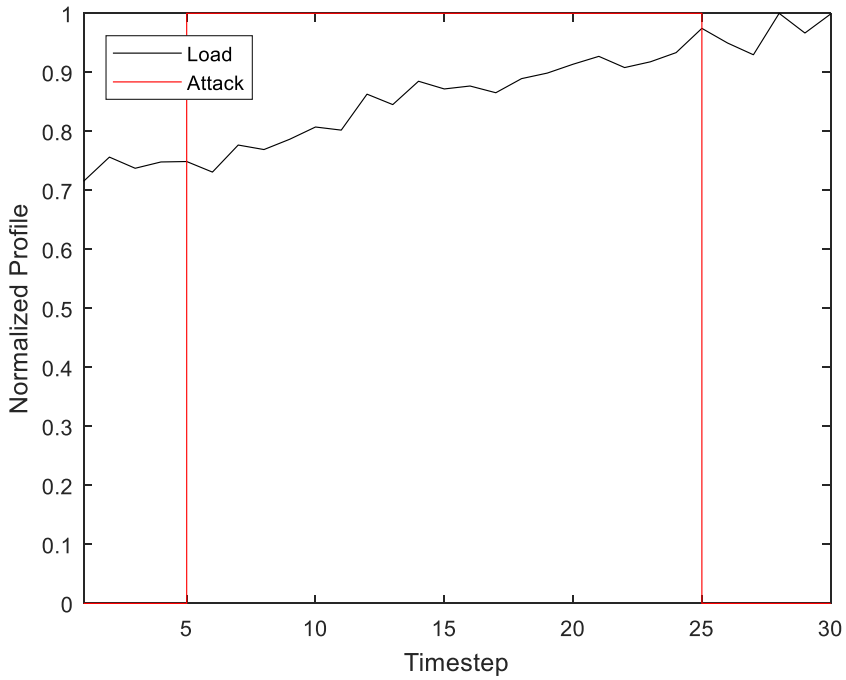

Figure 42. Normalized load profile and attack profile for simulation horizon

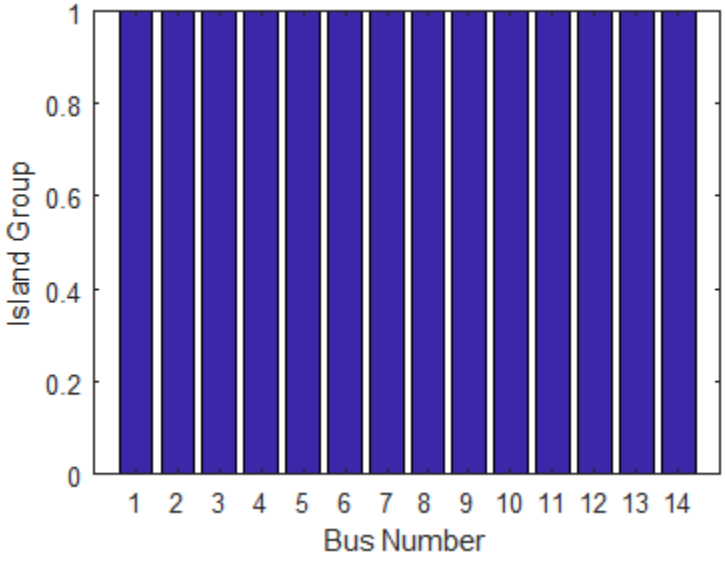

Figure 43. Observable island grouping for normal operations

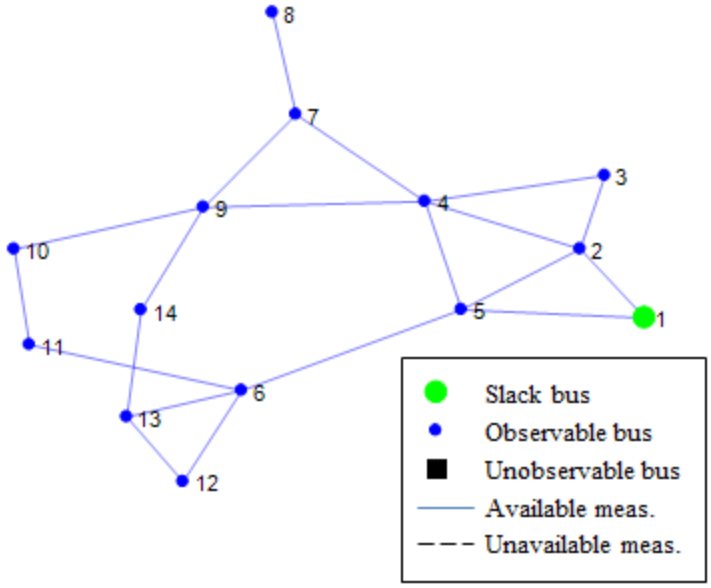

Figure 44. Measurement graph for normal operations

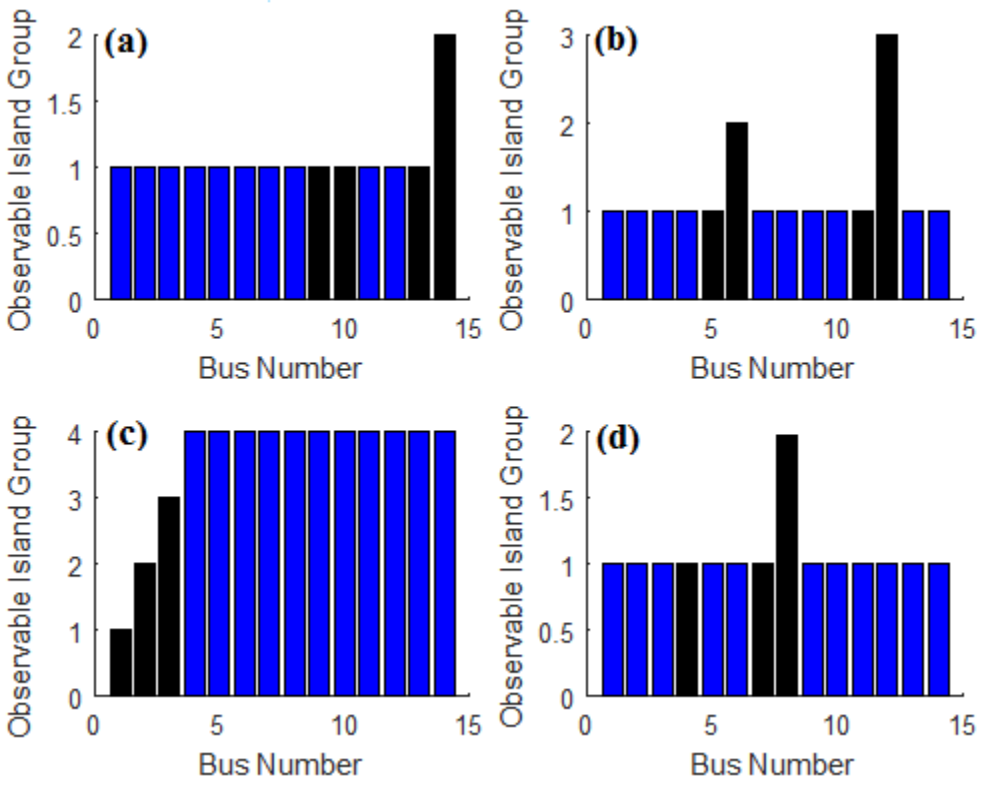

Figure 45. Observable island groupings for a DoS attack on: (a) RTU 1; (b) RTU 2; (c) RTU 3; (d) RTU 4 – black indicates unavailable measurements

If RTU 2 is under attack (see Figure 46b), measurements for buses 5, 6, 11, and 12 are unavailable. The system splits into three observable islands: bus 6, bus 12, and the remaining buses. Note that the global slack bus in this scenario is reassigned from bus 1 to bus 8. Because bus 1 is connected to bus 5 whose measurements are unavailable (illustrated using a black dot), the Gain matrix is ill-conditioned, and the solution error becomes large if bus 1 was selected as the slack. This issue can be avoided if the slack bus is reassigned to the bus furthest away from the unobservable region, bus 8.

If RTU 3 is under attack (see Figure 46c), measurements from buses 1, 2, and 3 are unavailable. Even though RTU 3 has only 3 buses, it has the largest negative impact, because its loss creates 4 observable islands and 3 unobservable buses. Once again, the global slack bus needs to be reassigned, since bus 1 is associated with RTU 3. Bus 8 is selected again since it is one of the furthest buses from the unobservable region. However, even after the slack reassignment, the Gain matrix is singular, so the solution error is very large.

If RTU 4 is under attack (see Figure 46d), the measurements for buses 4, 7, and 8 are unavailable, creating two observable islands: bus 8 and the remaining buses. Because bus 1 is far from the region with unavailable measurements, the slack bus does not need to be reassigned.

Figure 47 shows a plot of the voltage angle solution errors for five different scenarios: normal operation and the respective loss of RTU 1, RTU 2, RTU 3, and RTU 4. Figure 48 shows a plot of the voltage magnitude errors for those same five scenarios. Based on the results, it is clear that an attack on RTU 3 would have the most severe impact on solution accuracy. In particular, note the voltage magnitude error. At $t = 20$, the average magnitude error spikes almost as high as 800, which is not feasible considering that the voltage magnitude should be close to 1.0 pu.

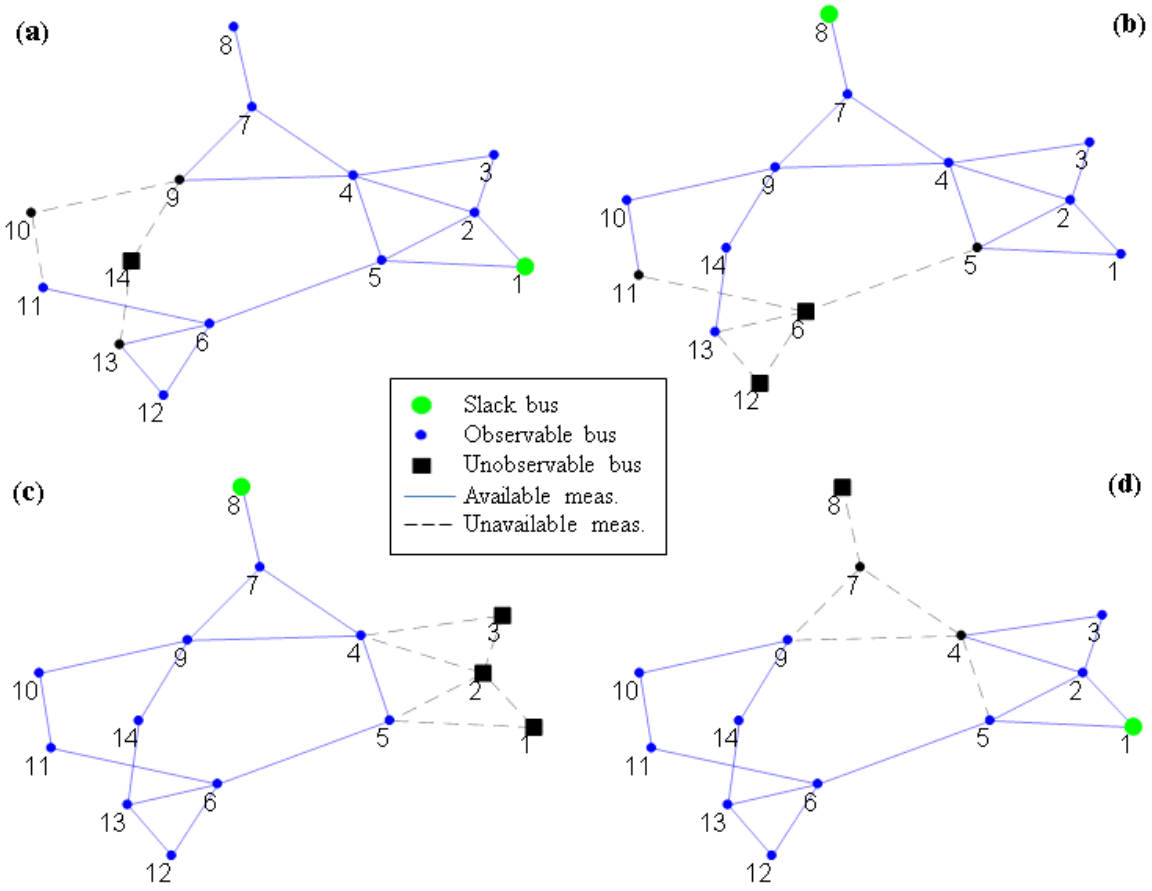

Figure 46. Measurement graph for attacks on: (a) RTU 1; (b) RTU 2; (c) RTU 3; (d) RTU 4

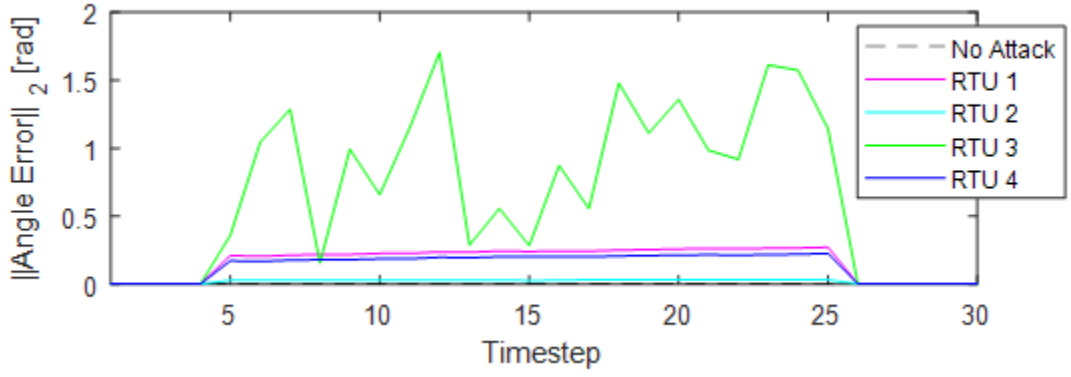

Figure 47. L2-norm of voltage angle errors for normal scenario versus an attack on RTU 1, RTU 2, RTU 3, and RTU 4

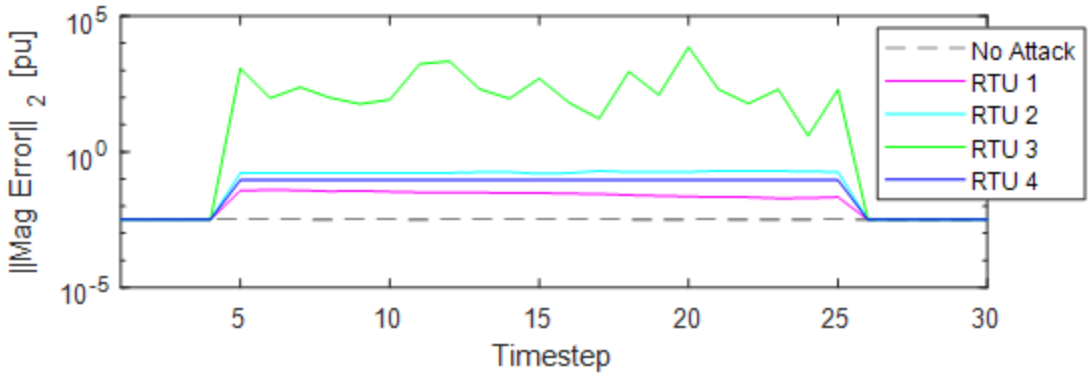

Figure 48. L2-norm of voltage magnitude errors for normal scenario versus RTU attack scenarios – (a) decimal scale; (b) logarithmic scale

When we do a closer inspection of the results for t = 20 in Figure 49, we see that the absolute value of the voltage angle and magnitude errors for each bus would be on the order of 10^{-3} if there was no attack. However, the L2-norm of the voltage angle error can be as high as 1.7057 at t = 12, and the L2-norm of the voltage magnitude error can be as high as 7203.23 at t = 20 for bus 1 when RTU 3 is under a simulated attack. The state estimator solution becomes meaningless since too many buses are unobservable.

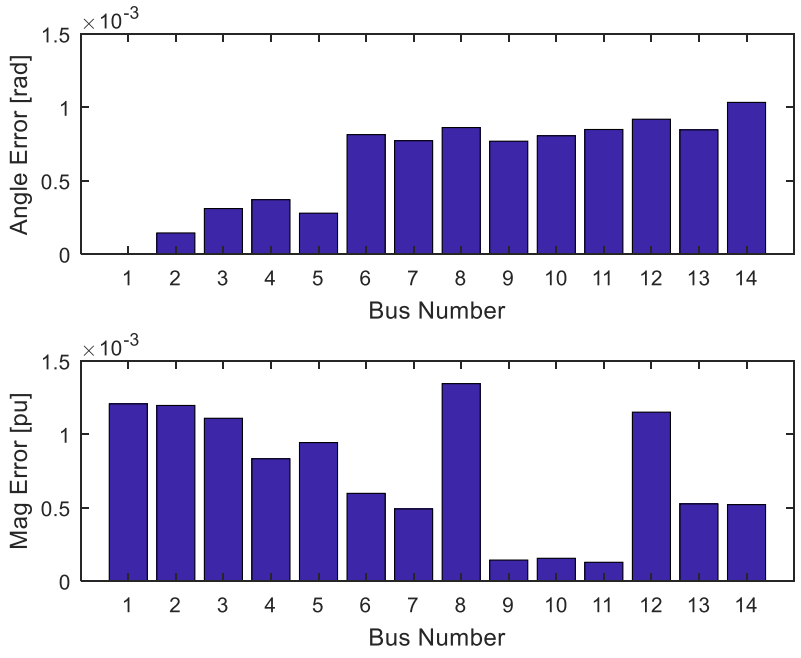

Figure 49. Voltage angle and magnitude error for normal scenario at t = 20

From Figure 50, we can see that the large voltage angle errors at t = 20 correspond to the buses that are unobservable. For a simulated attack on RTU 3, bus 1 has the largest angle error, followed by buses 2 and 3. Similarly, the large voltage magnitude errors also correspond to the unobservable buses (see Figure 51). Bus 3 has the largest voltage magnitude error, followed by bus 1 and bus 2. Table I ranks the 4 RTUs from the most severe to least severe loss. With its avg$\|V_{diff}\|_2$ being several orders of magnitude larger than that of the other RTUs, we can see the loss of RTU 3 has the biggest impact on the state estimator. Note that although RTU 3 is not the RTU with the most number of buses (it has three instead of four like RTU 1 and RTU 2), it is still the most critical one.

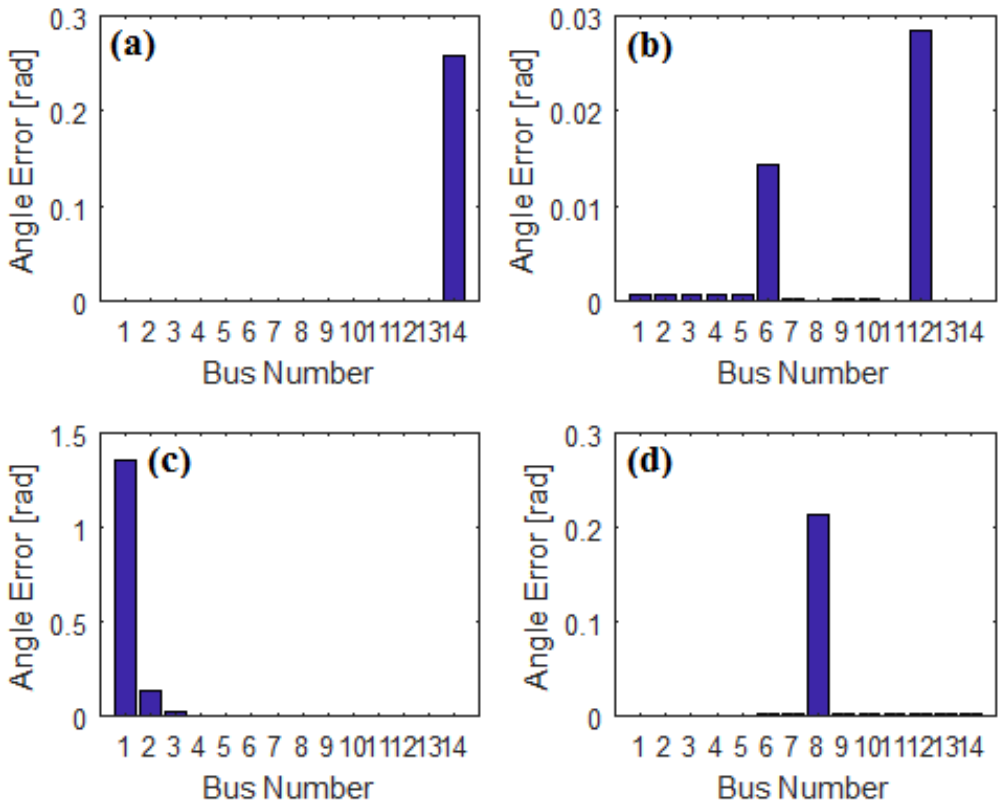

Figure 50. SE voltage angle error for an attack at t = 20 on: (a) RTU 1; (b) RTU 2; (c) RTU 3; (d) RTU 4

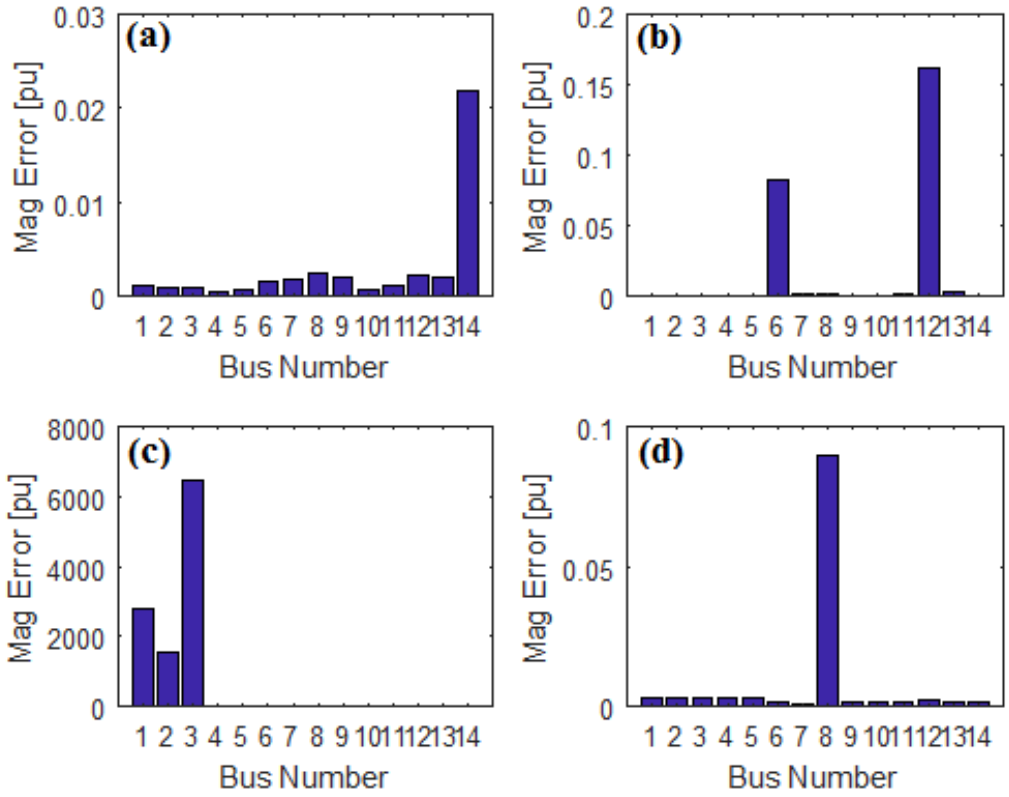

Figure 51. SE voltage magnitude error for an attack at t = 20 on: (a) RTU 1; (b) RTU 2; (c) RTU 3; (d) RTU 4

Table 8. *N*-1 RTU Ranking Results for IEEE 14-Bus System

Severity Ranking	RTU #	# Sub Buses	# Obs. Islands	Avg ||Ang Err||$_2$	Avg ||Mag Err||$_2$	Avg $\|\|V_{diff}\|\|_2$
1	3	3	4	0.94731	762.555	762.841
2	1	4	2	0.23922	0.02926	0.24390
3	4	3	2	0.19784	0.09034	0.22523
4	2	4	3	0.03052	0.17607	0.17864
Normal	-	-	1	0.00251	0.00318	0.00412

Illinois 200-Bus Test System [107]

The Illinois 200-bus test system has 200 buses, 246 branches, 49 generators, 160 loads, and 111 substations. There are 2 substations with 8 buses, 1 substation with 7 buses, 1 with 6 buses, 3 with 5 buses, 4 with 4 buses, 8 with 3 buses, and the remaining substations with 2 or fewer buses. Using the methodology described in Section 5.4.3, we

perform an assessment of all 111 substation RTUs. We show select results in Fig. 9 and Fig. 10 for the 8 worst scenarios as well as the results for the loss of RTU 34, which is one of the two largest RTUs, and the loss of RTU 102, which has the slack bus. Fig. 9 shows the L2-norm of the SE voltage angle errors, and Fig. 10 shows the L2-norm of the SE voltage magnitude errors over the course of the simulation horizon.

Table 9 shows a ranked list of select RTUs. It includes the avg$\left\|V_{diff}\right\|_2$ metric as well as the time-average L2-norm of only the angle error, which ranged as high as 0.47693, and the time-average L2-norm of only the voltage error, which ranged as high as 0.08619. Note that once again losing the largest RTU does not necessarily cause the most severe impact. For example, RTU 34 has 8 buses, but it is only ranked 13th while RTU 15 with 3 buses is ranked 5th. Empirically the avg$\left\|V_{diff}\right\|_2$ metric seems to depend on the system topology and RTU assignment, so it is important to run an offline systematic assessment on each power system to identify the most critical RTUs. Also, the effect of losing a single RTU impacts large systems less than small systems.

Table 9. *N*-1 RTU Cyber-Physical Assessment Results for Illinois 200-bus System

Severity Ranking	RTU #	# Sub Buses	# Obs. Islands	Avg ||Ang Err||$_2$	Avg ||Mag Err||$_2$	Avg $\|\|V_{diff}\|\|_2$
1	88	8	8	0.47693	0.08619	0.49159
2	23	6	6	0.30681	0.08904	0.32511
3	8	5	5	0.25693	0.03415	0.26128
4	44	4	4	0.22517	0.06256	0.23740
5	15	3	3	0.22326	0.02622	0.22662
6	13	3	3	0.22287	0.01724	0.22458
7	108	3	3	0.16059	0.05091	0.17104
8	10	2	2	0.16315	0.03538	0.16955
13	34	8	8	0.08549	0.10607	0.16529
23	102	3	2	0.09611	0.04071	0.10616
Normal	-	-	1	0.00345	0.00728	0.00811

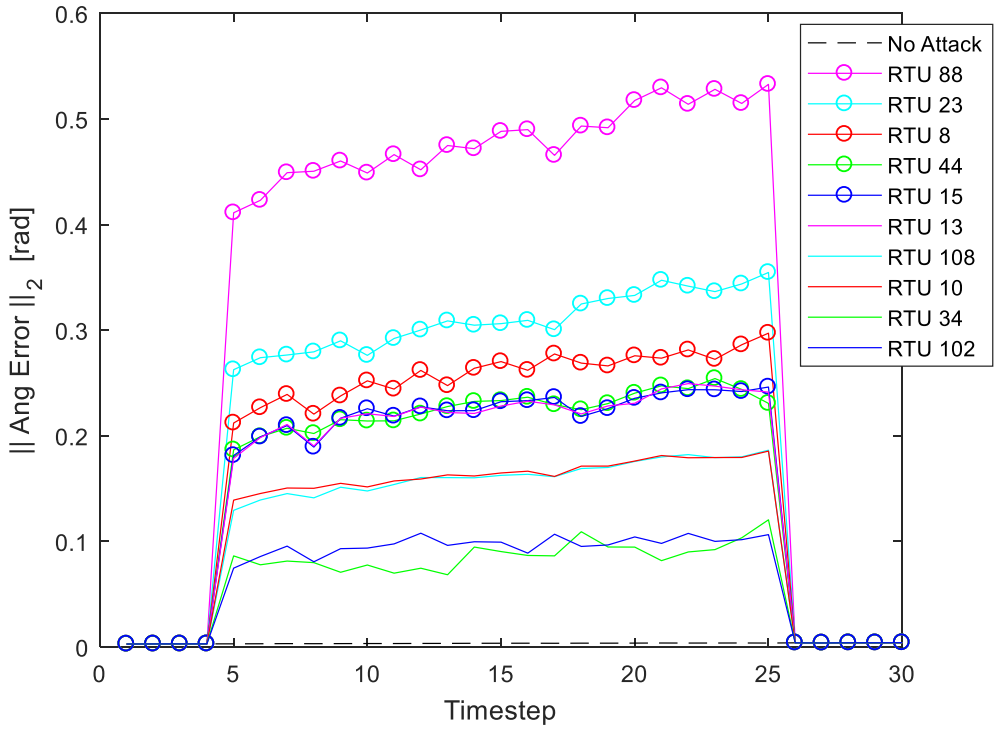

Figure 52. L2-norm of SE voltage angle errors for normal operations and select scenarios from the Illinois 200-bus system

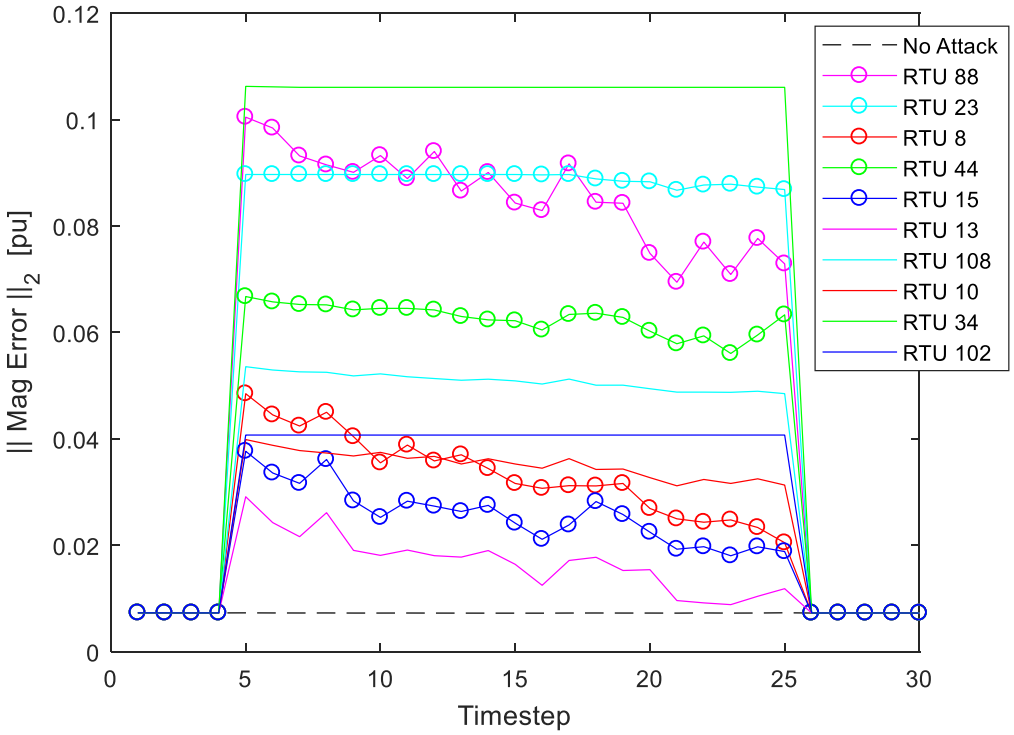

Figure 53. L2-norm of SE voltage magnitude errors for normal operations and select scenarios from the Illinois 200-bus system

5.4.5 Conclusions

In this paper, we proposed a novel methodology for *N*-1 RTU cyber-physical security assessment that systematically studies the impact of losing an RTU on the state estimator and a metric for ranking the severity of RTU loss. This approach could be used by power system operators to identify the most vulnerable RTUs. We tested our approach on two synthetic systems. Empirically we observed that losing the RTU with the most number of buses does not necessarily have the most severe impact, demonstrating the need for a systematic assessment approach. For future work, this approach could be generalized to the loss of multiple RTUs and to assess the impact of the number of critical measurements on the errors.

6 CONCLUSIONS

The work from this dissertation addressed two challenges for the state estimator of the future: 1) more data to process than ever before and 2) the growing threat of cyber-attacks. To address the first issue, we used automatic graph partitioning to enable a scalable decomposition-based state estimation approach. Empirical results suggest that decomposition reduces computation time, except for very small cases. This implies that any SE problem should be decomposed beyond its physical boundaries. With the speed gains, decomposition-based SE could be used to process a greater number of measurements than is currently tractable for central SE.

To address the second issue, we created a co-simulator that can simulate the combined effects of the electrical network and the communication network. By studying the cyber-physical system as a whole rather than as separate entities, we can see how attacks on one part of the system impacts the rest of the grid. We used the co-simulator to study the impact of bad command injection attacks on the power grid. In order to validate whether a command is malicious or not, we need information from the physical system as well as from the intrusion detection system. Also, we created a novel 3D visualization that can show how a power system event or cyber-attack evolves over time. 3D visualizations allow power system operators to see time-evolving trends in a single snapshot instead of needing to rely on animations. Finally, we presented a new cyber-physical security assessment methodology that evaluates the impact of network availability attacks on RTUs. We showed that losing even a single RTU can cause network observability issues. We also showed that the RTU with the most number of measurements is not necessarily the most critical RTU to lose.

6.1 Summary of Contributions

The contributions of the work in this dissertation are as follows:

1. Decentralized state estimation using automatic graph partitioning and ADMM
    a. Introduced automatic graph partitioning as a systematic method for virtually dividing a power system into an arbitrary number of sub-areas.
    b. Proposed a fast decomposition-based SE approach, an enhancement to traditional WLS SE, which uses the partitioning results to automatically decompose the global SE problem into smaller sub-problems and solve them in unison using ADMM.
    c. Explored empirically the impact of the number of sub-problems on the computational speed of the global SE problem for a serial implementation, and extrapolated the speedup seen for the considered IEEE test cases to larger power systems.
2. Cyber-physical security assessment for the smart grid
    a. Presented a co-simulator environment that is capable of simultaneously simulating the electrical network and the communication network of a given power system.
    b. Used the co-simulator to evaluate the impact of malicious commands on the smart grid.
    c. Proposed a metric for evaluating whether a command is malicious or not.
3. 3D spatiotemporal stacking visualization
    a. Presented a novel 3D stacking visualization as an alternative to the animation-style visualizations currently used in commercial power system analysis tools.
    b. Allows the power system operator to simultaneously view the power system one-line diagram while seeing time-varying behavior in the system.
    c. Provides the ability to hide irrelevant data, which prevents the power system operator from being overloaded with too much information.
4. *N*-1 RTU cyber-physical security assessment using state estimation

a. Proposed a novel methodology for *N*-1 RTU cyber-physical security assessment that systematically studies the impact of losing an RTU on the state estimator.
b. Proposed a metric for ranking the severity of RTU loss.
c. Observed that losing even a single RTU can cause network unobservability issues.
d. Observed that losing the RTU with the most number of measurements does not necessarily have the most severe impact, demonstrating the need for a systematic assessment approach.

6.2 Publications

1. N. Saxena, **L. Xiong**, V. Chukwuka, and S. Grijalva, "Impact evaluation of malicious control commands in cyber-physical smart grids," *IEEE Transactions on Sustainable Computing*, 2018 (accepted).

2. S. Grijalva, **L. Xiong**, and S. Vejdan, "Resilience analysis of modular controllable transformers," *IEEE Texas Power and Energy Conference*, College Station, TX, 2018.

3. N. Saxena, V. Chukwuka, **L. Xiong**, and S. Grijalva, "CPSA: A cyber-physical security assessment tool for situational awareness in smart grid," *ACM Workshop on Cyber-Physical Systems Security & Privacy (CPS-SPC)*, Dallas, TX, 2017.

4. **L. Xiong** and S. Grijalva, "Fast decomposition-based state estimation using automatic graph partitioning and ADMM," *IEEE PES General Meeting*, Boston, MA, 2016.

5. R. Pienta, **L. Xiong**, S. Grijalva, P. Chau, and M. Kahng, "STEPS: A spatio-temporal electric power systems visualization," *ACM International Conference on Intelligent User Interfaces*, 2016.

6. Y. Yu, S. Grijalva, J.J. Thomas, **L. Xiong**, P. Ju, and Y. Min, "Oscillation energy analysis of inter-area low-frequency oscillations in power systems," *IEEE Transactions on Power Systems*, vol. 31, no. 2, pp. 1195–1203, 2015.

7. A. Bose, T. Overbye, S. Grijalva, Y. Wang, F. Ye, K. Shetye, R. Macwan, F. Borth, and **L. Xiong**, "Seamless Bulk Electric Grid Management," *PSERC Project S-62G Final Report*, 2015.

8. S. Grijalva, A. Bose, **L. Xiong**, and J. McCalley, "Seamless Energy Management Systems Part I: Assessment of Energy Management Systems and Key Technological Requirements," *PSERC Project S-53G Final Report*, 2014.
9. S. Grijalva, **L. Xiong**, A. Bose, R. Pienta, M. Kahng, P. Chau, Y. Wang, and P. Yemula, "Seamless Energy Management Systems Part II: Development of Prototype Core Elements," *PSERC Project S-53G Final Report*, 2014.
10. **L. Xiong** and T. J. Overbye, "Visualizing market power in real-time electric power systems," *41st Annual North American Power Symposium*, Starkville, MS, 2009.

Pending:

1. **L. Xiong** and S. Grijalva, "*N*-1 RTU cyber-physical security assessment using state estimation," *IEEE PES General Meeting*, 2018 (submitted).
2. **L. Xiong** and S. Grijalva, "Distributed state estimation using automatic graph partitioning and ADMM" (journal paper in preparation).

REFERENCES

[1] Gilstrap and Matt, “United States Electricity Industry Primer,” 2015.

[2] FERC, “Regional Transmission Organizations.” 2015.

[3] “U.S. electric system is made up of interconnections and balancing authorities - Today in Energy - U.S. Energy Information Administration (EIA).” [Online]. Available: https://www.eia.gov/todayinenergy/detail.php?id=27152. [Accessed: 16-Mar-2018].

[4] NERC, “NERC Balancing Authorities As of October 1, 2015.”

[5] S. Grijalva, *ECE 6320: Power System Control and Operation Lecture Notes*. .

[6] Y.-F. Huang, S. Werner, J. Huang, N. Kashyap, and V. Gupta, “State Estimation in Electric Power Grids: Meeting New Challenges Presented by the Requirements of the Future Grid,” *IEEE Signal Process. Mag.*, vol. 29, no. 5, pp. 33–43, Sep. 2012.

[7] “Analysis of the Cyber Attack on the Ukrainian Power Grid: Defense Use Case,” Mar. 2016.

[8] R. Ebrahimian and R. Baldick, “State estimation distributed processing [for power systems],” *IEEE Trans. Power Syst.*, vol. 15, no. 4, pp. 1240–1246, Nov. 2000.

[9] CAISO, “Readiness Criterion Identifier Readiness Category Criteria Measurable Elements Threshold Owner Status Evidence Tariff Mapping.”

[10] NYISO, “Manual 12 Transmission and Dispatch Operations Manual Transmission and Dispatch Operations Manual | ii Transmission and Dispatch Operations Manual | iii,” 2018.

[11] ERCOT, “State Estimator Observability and Redundancy Requirements,” 2004. [Online]. Available: http://www.ercot.com/content/meetings/tac/keydocs/2004/1104/TAC11042004-18.doc.

[12] M. Boddeti, D. Obadina, F. Garcia, N. D. R. Sarma, Y. Wu, and V. Kanduri, “ERCOT’s Experiences in Using Pseudo Measurements in State Estimation,” Detroit, MI, 2011.

[13] E. Litvinov, “Locational Marginal Pricing,” 2011.

[14] MISO, “Synchrophasor Data and State Estimation NASPI Workshop-MISO,” 2015.

[15] SPP, “Southwest Power Pool Independent Coordinator of Transmission RELIABILITY PLAN for Entergy in the Southeastern Electric Reliability Council,” 2009.

[16] F. C. Schweppe and D. B. Rom, “Power System Static-State Estimation, Part II: Approximate Model,” *IEEE Trans. Power Appar. Syst.*, vol. PAS-89, no. 1, pp. 125–130, Jan. 1970.

[17] F. C. Schweppe, “Power System Static-State Estimation, Part III: Implementation,” *IEEE Trans. Power Appar. Syst.*, vol. PAS-89, no. 1, pp. 130–135, Jan. 1970.

[18] A. Garcia, A. Monticelli, and P. Abreu, “Fast Decoupled State Estimation and Bad Data Processing,” *IEEE Trans. Power Appar. Syst.*, vol. PAS-98, no. 5, pp. 1645–1652, Sep. 1979.

[19] A. Abur and A. Gómez Expósito, *Power System State Estimation: Theory and*

Implementation, vol. 24. CRC Press, 2004.
[20] H. M. Merrill and F. C. Schweppe, "Bad Data Suppression in Power System Static State Estimation," *IEEE Trans. Power Appar. Syst.*, vol. PAS-90, no. 6, pp. 2718–2725, Nov. 1971.
[21] Q. Li and V. Vittal, "The convex hull of the AC power flow equations in rectangular coordinates," in *2016 IEEE Power and Energy Society General Meeting (PESGM)*, 2016, pp. 1–5.
[22] G. Krumpholz, K. Clements, and P. Davis, "Power System Observability: A Practical Algorithm Using Network Topology," *IEEE Trans. Power Appar. Syst.*, vol. PAS-99, no. 4, pp. 1534–1542, Jul. 1980.
[23] K. A. Clements, G. R. Krumpholz, and P. W. Davis, "Power System State Estimation with Measurement Deficiency: An Algorithm that Determines the Maximal Observable Subnetwork," *IEEE Trans. Power Appar. Syst.*, vol. PAS-101, no. 9, pp. 3044–3052, Sep. 1982.
[24] K. Clanents, G. Krutnpholz, and P. Davis, "Power System State Estimation with Measurement Deficiency: an Observability/Measurement Placement Algorithm," *IEEE Trans. Power Appar. Syst.*, vol. PAS-102, no. 7, pp. 2012–2020, Jul. 1983.
[25] K. Clements, G. Krumpholz, and P. Davis, "Power System State Estimation Residual Analysis: An Algorithm Using Network Topology," *IEEE Trans. Power Appar. Syst.*, vol. PAS-100, no. 4, pp. 1779–1787, Apr. 1981.
[26] V. H. Quintana, A. Simoes-Costa, and A. Mandel, "Power System Topological Observability Using a Direct Graph-Theoretic Approach," *IEEE Trans. Power Appar. Syst.*, vol. PAS-101, no. 3, pp. 617–626, Mar. 1982.
[27] R. R. Nucera and M. L. Gilles, "Observability analysis: a new topological algorithm," *IEEE Trans. Power Syst.*, vol. 6, no. 2, pp. 466–475, May 1991.
[28] H. Mori and S. Tsuzuki, "A fast method for topological observability analysis using a minimum spanning tree technique," *IEEE Trans. Power Syst.*, vol. 6, no. 2, pp. 491–500, May 1991.
[29] H. Mori and H. Tanaka, "A genetic approach to power system topological observability," in *1991., IEEE International Sympoisum on Circuits and Systems*, pp. 1141–1144.
[30] A. Jain, R. Balasubramanian, S. C. Tripathy, B. N. Singh, and Y. Kawazoe, "Power system topological observability analysis using artificial neural networks," in *IEEE Power Engineering Society General Meeting, 2005*, pp. 2687–2692.
[31] A. Jain, R. Balasubramanian, S. C. Tripathy, and Y. Kawazoe, "Topological observability analysis using heuristic rule based expert system," in *2006 IEEE Power Engineering Society General Meeting*, 2006, p. 6 pp.
[32] S. Vazquez-Rodriguez, A. Faina, and B. Neira-Duenas, "An Evolutionary Technique with Fast Convergence for Power System Topological Observability Analysis," in *2006 IEEE International Conference on Evolutionary Computation*, pp. 3086–3090.
[33] L. Xiong and S. Grijalva, "Fast decomposition-based state estimation using automatic graph partitioning and ADMM," in *2016 IEEE Power and Energy Society General Meeting (PESGM)*, 2016, pp. 1–5.
[34] D. M. Falcao, F. F. Wu, and L. Murphy, "Parallel and distributed state estimation," *IEEE Trans. Power Syst.*, vol. 10, no. 2, pp. 724–730, May 1995.

[35] A. Gómez Expósito, A. de la Villa Jaén, C. Gómez-Quiles, P. Rousseaux, and T. Van Cutsem, "A taxonomy of multi-area state estimation methods," *Electr. Power Syst. Res.*, vol. 81, no. 4, pp. 1060–1069, Apr. 2011.
[36] V. Kekatos and G. B. Giannakis, "Distributed Robust Power System State Estimation," *IEEE Trans. Power Syst.*, vol. 28, no. 2, pp. 1617–1626, May 2013.
[37] A. P. S. Meliopoulos, G. J. Cokkinides, C. Hedrington, and T. L. Conrad, "The supercalibrator - A fully distributed state estimator," in *IEEE PES General Meeting*, 2010, pp. 1–8.
[38] L. Xie, D. H. Choi, and S. Kar, "Cooperative distributed state estimation: Local observability relaxed," in *2011 IEEE Power and Energy Society General Meeting*, 2011, pp. 1–11.
[39] L. Xie, D.-H. Choi, S. Kar, and H. V. Poor, "Fully Distributed State Estimation for Wide-Area Monitoring Systems," *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1154–1169, Sep. 2012.
[40] H. Karimipour and V. Dinavahi, "Parallel Domain Decomposition Based Distributed State Estimation for Large-scale Power Systems," *IEEE Trans. Ind. Appl.*, pp. 1–1, 2015.
[41] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, *Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers*, vol. 1. 2011.
[42] P.-O. Fjällström, "Algorithms for Graph Partitioning: A Survey," Aug-1998. .
[43] G. Karypis and V. Kumar, "Multilevelk-way Partitioning Scheme for Irregular Graphs," *J. Parallel Distrib. Comput.*, vol. 48, no. 1, pp. 96–129, Jan. 1998.
[44] G. Karypis, "METIS - Serial Graph Partitioning and Fill-reducing Matrix Ordering." .
[45] M. S. Lobo, L. Vandenberghe, S. Boyd, H. Lebret, and D. P. O 'leary, "Applications of second-order cone programming '," *Elsevier Linear Algebr. its Appl.*, vol. 284, pp. 193–228, 1998.
[46] F. Alizadeh and D. Goldfarb, "Second-order cone programming," *Math. Program., Ser. B*, vol. 95, pp. 3–51, 2003.
[47] H. Zhu and G. B. Giannakis, "Estimating the state of AC power systems using semidefinite programming," in *2011 North American Power Symposium*, 2011, pp. 1–7.
[48] H. Zhu and G. B. Giannakis, "Robust power system state estimation for the nonlinear AC flow model," in *2012 North American Power Symposium (NAPS)*, 2012, pp. 1–6.
[49] H. Zhu and G. B. Giannakis, "Multi-area state estimation using distributed SDP for nonlinear power systems," in *2012 IEEE Third International Conference on Smart Grid Communications (SmartGridComm)*, 2012, pp. 623–628.
[50] H. Zhu and G. B. Giannakis, "Power System Nonlinear State Estimation Using Distributed Semidefinite Programming," *IEEE J. Sel. Top. Signal Process.*, vol. 8, no. 6, pp. 1039–1050, Dec. 2014.
[51] Yang Weng, Qiao Li, R. Negi, and M. Ilic, "Semidefinite programming for power system state estimation," in *2012 IEEE Power and Energy Society General Meeting*, 2012, pp. 1–8.
[52] Yang Weng, Qiao Li, R. Negi, and M. Ilic, "Distributed algorithm for SDP state

estimation," in *2013 IEEE PES Innovative Smart Grid Technologies Conference (ISGT)*, 2013, pp. 1–6.
[53] Y. Weng, B. Fardanesh, M. D. Ilic, and R. Negi, "Novel approaches using semidefinite programming method for power systems state estimation," in *2013 North American Power Symposium (NAPS)*, 2013, pp. 1–6.
[54] Q. Li, Y. Weng, R. Negi, and M. D. Ilić, "Convexification of bad data and topology error detection and identification problems in AC electric power systems," *IET Gener. Transm. Distrib.*, vol. 9, no. 16, pp. 2760–2767, Dec. 2015.
[55] S.-J. Kim, G. Wang, and G. B. Giannakis, "Online semidefinite programming for power system state estimation," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 6024–6027.
[56] G. Wang, S.-J. Kim, and G. B. Giannakis, "Moving-horizon dynamic power system state estimation using semidefinite relaxation," in *2014 IEEE PES General Meeting | Conference & Exposition*, 2014, pp. 1–5.
[57] S.-J. Kim, "Online power system state estimation using alternating direction method of multipliers," in *2015 IEEE Power & Energy Society General Meeting*, 2015, pp. 1–5.
[58] W. Zheng, W. Wu, A. Gomez-Exposito, B. Zhang, and Y. Guo, "Distributed Robust Bilinear State Estimation for Power Systems with Nonlinear Measurements," *IEEE Trans. Power Syst.*, vol. 32, no. 1, pp. 499–509, Jan. 2017.
[59] Y. Zhang, R. Madani, and J. Lavaei, "Conic Relaxations for Power System State Estimation with Line Measurements," *IEEE Trans. Control Netw. Syst.*, pp. 1–1, 2017.
[60] H. G. Aghamolki, Z. Miao, and L. Fan, "SOCP Convex Relaxation-Based Simultaneous State Estimation and Bad Data Identification," Apr. 2018.
[61] V. Kekatos, G. B. Giannakis, and B. Wollenberg, "Optimal Placement of Phasor Measurement Units via Convex Relaxation," *IEEE Trans. Power Syst.*, vol. 27, no. 3, pp. 1521–1530, Aug. 2012.
[62] I. N. L. Mission Support Center, "Cyber Threat and Vulnerability Analysis of the U.S. Electric Sector (Mission Support Center Analysis Report)," 2016.
[63] R. Joshi, "How PNNL and RTI Built a Secure Industrial Control System with Connext DDS," *RTI Blog*, 2014. [Online]. Available: http://blogs.rti.com/2014/06/05/how-pnnl-and-rti-built-a-secure-industrial-control-system-with-connext-dds/. [Accessed: 21-May-2018].
[64] Lloyd's and University of Cambridge, "Business Blackout," 2015.
[65] P. Marsters and T. Houser, "America's Biggest Blackout | Rhodium Group," *Rhodium Group*, 2017. [Online]. Available: https://rhg.com/research/americas-biggest-blackout-2/. [Accessed: 18-May-2018].
[66] NOAA, "NOAA Posts Images Online of Northeast Blackout," 2003. [Online]. Available: http://www.noaanews.noaa.gov/stories/s2015.htm. [Accessed: 18-May-2018].
[67] E. E. Bernabeu and F. Katiaei, "Aurora Vulnerability: Issues & Solutions Hardware Mitigation Devices (HMDs)," *Dominion) Farid Katiraei, Ph.D. (Quanta Technology)*. 2011.
[68] P. W. Parfomak, "Physical Security of the U.S. Power Grid: High-Voltage Transformer Substations Specialist in Energy and Infrastructure Policy," 2014.

[69] E-ISAC, “NERC Analysis of the Cyber Attack on the Ukrainian Power Grid,” 2016.
[70] Cathal McMahon, “Exclusive: EirGrid targeted by ‘state sponsored’ hackers leaving networks exposed to ‘devious attack’ - Independent.ie.” [Online]. Available: https://www.independent.ie/irish-news/news/exclusive-eirgrid-targeted-by-state-sponsored-hackers-leaving-networks-exposed-to-devious-attack-36003502.html. [Accessed: 29-Apr-2018].
[71] N. Falliere, L. O. Murchu, and E. Chien, “W32.Stuxnet Dossier,” 2011.
[72] Forensic Engineering, “Understanding the Stuxnet Worm « Forensic Engineering Hub,” 2011. [Online]. Available: http://www.armstrongforensic.com/blog/index.php/2011/02/10/understanding-the-stuxnet-worm/. [Accessed: 30-May-2018].
[73] E. Byres, “Air Gaps won’t Stop Stuxnet’s Children | Tofino Industrial Security Solution,” *Tofino Security*, 2012. [Online]. Available: https://www.tofinosecurity.com/blog/air-gaps-won’t-stop-stuxnet’s-children. [Accessed: 30-May-2018].
[74] K. Zetter, “Everything We Know About Ukraine’s Power Plant Hack | WIRED,” *Wired*, 2016. [Online]. Available: https://www.wired.com/2016/01/everything-we-know-about-ukraines-power-plant-hack/. [Accessed: 12-May-2018].
[75] Antiy Labs, “Comprehensive Analysis Report on Ukraine Power System Attacks - Antiy Labs | The Next Generation Anti-Virus Engine Innovator,” 2016. [Online]. Available: http://www.antiy.net/p/comprehensive-analysis-report-on-ukraine-power-system-attacks/. [Accessed: 21-May-2018].
[76] A. Greenberg, “Crash Override Malware Took Down Ukraine’s Power Grid Last December | WIRED,” *Wired*, 2017. [Online]. Available: https://www.wired.com/story/crash-override-malware/. [Accessed: 15-May-2018].
[77] R. Lipovsky and A. Cherepanov, “Industroyer: Biggest threat to industrial control systems since Stuxnet,” *Welivesecurity*, 2017. [Online]. Available: https://www.welivesecurity.com/2017/06/12/industroyer-biggest-threat-industrial-control-systems-since-stuxnet/. [Accessed: 17-May-2018].
[78] Dragos Inc., “CRASHOVERRIDE: Analysis of the Threat to Electric Grid Operations,” 2017.
[79] D. E. Sanger and W. J. Broad, “Pentagon Suggests Countering Devastating Cyberattacks With Nuclear Arms - The New York Times,” *New York Times*, 16-Jan-2018. [Online]. Available: https://www.nytimes.com/2018/01/16/us/politics/pentagon-nuclear-review-cyberattack-trump.html. [Accessed: 18-May-2018].
[80] United States Computer Emergency Readiness Team, “Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors | US-CERT,” *TA18-074A*, 2018. [Online]. Available: https://www.us-cert.gov/ncas/alerts/TA18-074A. [Accessed: 21-May-2018].
[81] NERC, “NERC Grid Security Exercise GridEx IV: Lessons Learned,” 2018.
[82] W. Wang and Z. Lu, “Cyber security in the Smart Grid: Survey and challenges,” *Comput. Networks*, vol. 57, no. 5, pp. 1344–1371, Apr. 2013.
[83] L. Xie, Y. Mo, and B. Sinopoli, “Integrity Data Attacks in Power Market Operations,” *IEEE Trans. Smart Grid*, vol. 2, no. 4, pp. 659–666, Dec. 2011.

[84] L. Xie, Y. Mo, and B. Sinopoli, “False Data Injection Attacks in Electricity Markets,” in *2010 First IEEE International Conference on Smart Grid Communications*, 2010, pp. 226–231.
[85] S. Sridhar and G. Manimaran, “Data integrity attacks and their impacts on SCADA control system,” in *IEEE PES General Meeting*, 2010, pp. 1–6.
[86] S. Sridhar and M. Govindarasu, “Model-Based Attack Detection and Mitigation for Automatic Generation Control,” *IEEE Trans. Smart Grid*, vol. 5, no. 2, pp. 580–591, Mar. 2014.
[87] S. Sridhar, A. Hahn, and M. Govindarasu, “Cyber-Physical System Security for the Electric Power Grid,” *Proc. IEEE*, vol. 100, no. 1, pp. 210–224, Jan. 2012.
[88] T. J. Overbye, D. A. Wiegmann, and R. J. Thomas, “Visualization of Power Systems,” *PSERC Final Rep.*, vol. 10, pp. 22–25, 2002.
[89] T. J. Overbye, “Visualization enhancements for power system situational assessment,” in *2008 IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century*, 2008, pp. 1–4.
[90] J. D. Weber and T. J. Overbye, “Voltage contours for power system visualization,” *IEEE Trans. Power Syst.*, vol. 15, no. 1, pp. 404–409, Feb. 2000.
[91] Y. Sun and T. J. Overbye, “Visualizations for power system contingency analysis data,” *IEEE Trans. Power Syst.*, vol. 19, no. 4, pp. 1859–1866, Nov. 2004.
[92] T. J. Overbye and J. D. Weber, “Visualization of power system data,” in *Proceedings of the 33rd Annual Hawaii International Conference on System Sciences*, 2000, p. 7 pp.-.
[93] “PowerWorld » The visual approach to electric power systems.” .
[94] M. Farrugia and A. Quigley, “Effective Temporal Graph Layout: A Comparative Study of Animation Versus Static Display Methods,” *Inf. Vis.*, vol. 10, no. 1, pp. 47–64, Jan. 2011.
[95] S. Rufiange and M. J. McGuffin, “DiffAni: Visualizing Dynamic Graphs with a Hybrid of Difference Maps and Animation,” *IEEE Trans. Vis. Comput. Graph.*, vol. 19, no. 12, pp. 2556–2565, Dec. 2013.
[96] M. Itoh, M. Toyoda, and M. Kitsuregawa, “An Interactive Visualization Framework for Time-Series of Web Graphs in a 3D Environment,” in *2010 14th International Conference Information Visualisation*, 2010, pp. 54–60.
[97] C. Tominski, H. Schumann, G. Andrienko, and N. Andrienko, “Stacking-Based Visualization of Trajectory Attribute Data,” *IEEE Trans. Vis. Comput. Graph.*, vol. 18, no. 12, pp. 2565–2574, Dec. 2012.
[98] B. Bach, E. Pietriga, and J.-D. Fekete, “Visualizing Dynamic Networks with Matrix Cubes,” in *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems*, 2014, pp. 877–886.
[99] R. Pienta, L. Xiong, S. Grijalva, D. H. (Polo) Chau, and M. Kahng, “STEPS: A Spatio-temporal Electric Power Systems Visualization,” in *Companion Publication of the 21st International Conference on Intelligent User Interfaces*, 2016, pp. 32–35.
[100] Y. Liu, P. Ning, and M. K. Reiter, “False Data Injection Attacks Against State Estimation in Electric Power Grids,” in *Proceedings of the 16th ACM Conference on Computer and Communications Security*, 2009, pp. 21–32.
[101] G. Hug and J. A. Giampapa, “Vulnerability Assessment of AC State Estimation

With Respect to False Data Injection Cyber-Attacks," *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1362–1370, Sep. 2012.
[102] H. Zhang, P. Cheng, L. Shi, and J. Chen, "Optimal Denial-of-Service Attack Scheduling With Energy Constraint," *IEEE Trans. Automat. Contr.*, vol. 60, no. 11, pp. 3023–3028, Nov. 2015.
[103] Heng Zhang, Peng Cheng, Ling Shi, and Jiming Chen, "Optimal DoS attack policy against remote state estimation," in *52nd IEEE Conference on Decision and Control*, 2013, pp. 5444–5449.
[104] H. Zhang, Y. Qi, J. Wu, L. Fu, and L. He, "DoS Attack Energy Management Against Remote State Estimation," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 1, pp. 383–394, Mar. 2018.
[105] Shichao Liu, X. P. Liu, and A. El Saddik, "Denial-of-Service (dos) attacks on load frequency control in smart grids," in *2013 IEEE PES Innovative Smart Grid Technologies Conference (ISGT)*, 2013, pp. 1–6.
[106] O. Vukovic and G. Dan, "Detection and localization of targeted attacks on fully distributed power system state estimation," in *2013 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2013, pp. 390–395.
[107] A. B. Birchfield, T. Xu, K. M. Gegner, K. S. Shetye, and T. J. Overbye, "Grid Structural Characteristics as Validation Criteria for Synthetic Networks," *IEEE Trans. Power Syst.*, vol. 32, no. 4, pp. 3258–3265, Jul. 2017.