The U.S. Electrical Power Grid as a Cyber-Physical System: Understanding Exposed Attack Surfaces

Brown R. Farinholt
University of California, San Diego

Abstract

The U.S. power grid is one of the largest installations of critical infrastructure in the world, the failure of which can cause significant socioeconomic damage. Due to the complexities involved in providing reliable power to customers over long distances, the grid's engineers rely on industrial automation technology to make important operational decisions. This technology exposes the grid to the risk of electronic attack.

In this paper, I describe the operation of the power grid and the control loops which govern it. I examine the industrial automation system that maintains the grid, and use it to present a model of the power grid as a cyber-physical system. Using this model, I enumerate the surfaces the grid exposes to electronic attack, and delve into several specific theoretical attacks. I conclude with an assessment of the grid's overall susceptibility to electronic attack, and discuss attractive directions for future research.

1 Introduction

The electrical power grids of populous, industrialized countries like the United States are among the most complex and large-scale installations of critical infrastructure ever created. The U.S. power grid alone is comprised of 340,000 km of high-voltage transmission line, 15,000 generators, and over 150 control centers [6].

Provision of stable electrical power by the grid is considered a permanent fixture by much of the urban world. Nearly all of the systems critical to sustaining urban life rely on a continuous supply of electricity, often without alternative. As urbanization expert Stephen Graham [24] notes, "In an electrical blackout it is not just electric lighting that fails. Electrically-powered water and sewerage systems tend to grind to a halt. Public transportation stops. Food processing and distribution is disabled. Health care becomes almost impossible. Even the Internet ceases to function."

The dependence of urban society on the power grid and the expectation of its availability therefore precipitate significant socioeconomic consequences when the grid fails [39]. The recent U.S. blackouts of 2003 [2] and 2011 [22], each caused by accidental grid failures, cost hundreds of millions of dollars in damages, affected tens of millions of people, and led to numerous deaths.

1.1 Automation Technology in the Grid

The power grid is a delicate feedback system. It aims to maintain equilibrium between the power it supplies and its consumers' fluctuating demand; its failure to do so can be catastrophic. Given the complexities involved in coordinating power transmission between numerous generators and consumers, the grid's operators use automation technology to manage the process [55].

This technology used to regulate the power grid is known as the energy management system (EMS). The EMS is an industrial control system that serves as the grid's central command and monitoring hub [64]. It provides operators with real-time measurements throughout the grid using a distributed network of sensors, and allows operators to remotely control the grid's topology and generation to satisfy consumer load or avoid stressing equipment. The EMS also triggers alarms when dangerous conditions are met or forecast, and can automatically initiate control measures such as load-shedding, energy re-routing, and generation adjustment. Further, the EMS directly influences the energy market, providing the measurements used for real-time electricity pricing.

We will discuss the grid's integration with the energy management system in explicit detail in Section 3.

1.2 Motivation

The capabilities afforded to the grid by industrial control systems are critical to its performance and reliability. In

fact, the U.S. Department of Energy has proposed the expansion of the role of automation and remote operation technology in the grid as part of its "smart grid" initiative [5]. Such an expansion is logical, given the complications that electric vehicles and household renewable energy generation present in terms of demand forecasting.

Herein lies the problem. These same industrial control systems that are critical to the power grid's functionality also expose the grid to the risk of electronic attack.

The EMS is inherently vulnerable to such attacks. It is designed to allow remote operators to wield its full monitoring and control capabilities, in doing so exposing itself to the public Internet [18]. The sensor networks it uses for monitoring and the communication channels through which it issues commands are often wireless (and sometimes over significant distances), exposing further attack vectors [60].

Attacks on authentication portals and wireless networks are nothing novel, but the context in which these vulnerabilities exist is significant. As we have observed, the power grid is a **cyber-physical system** [72], "a system of collaborating computational elements controlling physical entities."

In a cyber-physical system, exploits against exposed electronic attack surfaces can allow attackers to directly influence, or even control, machinery that affects change in the *physical* world. As the power grid's physical components control critical operations like nuclear power generation, damage to which could cause environmental and economic disaster, the risks associated with electronic attack on the grid are considerable. These risks are amplified by the fact that the power grid presents one of the largest attack surfaces on Earth.

In this paper, we present a simplified model of the power grid as a cyber-physical system. We use this model to enumerate the attack surfaces that the power grid exposes, and then discuss select attacks against them which have been studied by researchers. We conclude by evaluating the attacks and their implications towards the grid's overall security, so as to determine potential directions for future research.

The structure of the paper is as follows. In Section 2, we present the overarching themes of the paper. In Section 3 of the paper, we provide background on the power grid's operation and model it as a cyber-physical system. In Section 4, we recognize the electronic attack surfaces which the grid exposes. In Section 5, we detail four interesting attacks on the grid and their potential consequences. In Section 6, we assess our findings and ponder the future directions of power grid security research efforts.

2 Themes

In this section, we highlight two overarching themes that appear throughout our examination of power grid security. These themes will ultimately guide our discussion of the current state of grid security in Section 6, as well as help us determine possible directions for future work.

2.1 Security Through Obscurity

Security through obscurity refers to the security obtained by hiding the design or implementation of a system that may otherwise be vulnerable. Security through obscurity is generally eschewed by the security engineering community, as it offers no actual security guarantees to a system; however, it is employed frequently in industry, the power grid being no exception.

The owners and operators of the grid have long subscribed to the notion of security through obscurity. Grid configuration and topological details, equipment manufacturer and model, and even communication medium information are guarded closely. According to a former employee at a major U.S. power utility [7], even personnel working on the grid are kept on a need-to-know basis with regards to operational details.

Throughout this paper, we will encounter security through obscurity. Some of the attacks we examine are hindered by lack of knowledge of the grid's topology; others explicitly circumvent this requirement. In Section 6, we will discuss the costs and benefits of security through obscurity to attackers and defenders of the grid alike.

2.2 Cyber to Physical Translation

Another prominent theme in power grid security research is translation between cyber-attacks and physical events; specifically, its difficulty. The grid's purpose as a cyber-physical system is to translate "cyber-actions" into physical outcomes (e.g. a grid operator's interaction with a GUI into a physical change in the amount of power flowing into a city). So when an attacker hijacks the grid electronically, it affords him/her the ability to affect change in the physical world – change to dangerous, critical components of the physical world at that.

Modeling the physical effects of a cyber-attack on the grid is challenging. The grid's components handle a massive amount of electrical energy, and are dangerous because of it: high-voltage transmission lines are enormous fire hazards; power plants are potential environmental disasters; capacitor banks can fail explosively. Also, the grid is sensitive to external conditions like weather and climate. A hot day increases the likelihood and severity of a power line failure considerably. It is quite difficult

to predict the outcome of an attack against a grid component by observing the grid alone.

As we study various attacks in Section 5, we will find that most authors neglect to *quantify* the potential outcomes of their attacks in the physical world, instead focusing on immediate effects to the grid and hypothetical outcomes in the physical world (often with reference to historical events).

3 Modeling the Power Grid Under Attack

In this section, we introduce the actors that run the power grid. We describe the main control loops that stabilize the grid, and how the EMS facilitates this. We then use a cyber-physical system model to define the different attack surfaces the grid's EMS presents to attackers.

3.1 Threat Model

We describe several categories of power grid attacker, drawing on work by Nicholson *et al.* [49].

Nation State: Will conduct reconnaissance for intellectual property theft and military preparation. Will attempt to sabotage the grid in times of war.

Terrorist: Wants to visibly attack the grid to cause damage and create fear. Conducts reconnaissance to these ends.

Organized Crime: Wants to steal power services from the grid, or financially exploit the grid in some way (blackmail, ransoming).

Competitor: Conducts reconnaissance to steal new technology or business models. Could also attempt sabotage the grid to achieve competitive advantage.

Employee / Insider Attacker: Wants to cause damage to the company/utility running the grid, normally through sabotage. Particularly dangerous due to level of access to sensitive controls.

Rogue / Script Kiddie: Possibly attacking the grid without a goal. Could want to gather information or cause general mayhem. Unlikely to be skilled enough to do the latter.

Hacktivist: Wants to disable the grid, likely through sabotage of control systems.

We observe that there are three main types of attacker, categorized by goal: damage to the grid; reconnaissance against the grid; and financial exploitation and theft. In this paper, we will focus on the first of these categories, **attackers aiming to cause damage to, or disable, the power grid.** We note that often reconnaissance is required by these attackers as well.

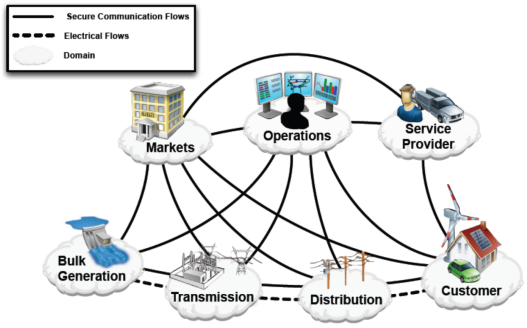

Figure 1: A conceptual model of the U.S. power grid [4]

3.2 Actors in the Grid

NIST traditionally models the power grid as a conceptual view of the actors involved in transmitting power from generators to customers, shown in Figure 1. We will describe the roles of each actor.

Bulk Generation: The primary producers of electrical power in the grid. Bulk generation encompasses all power plants, from nuclear-powered steam plants to hydro-electric dams to wind farms. The vast majority of customer demand for power is met by bulk generation.

Transmission: Carries the power produced by bulk generators at high voltage over large distances. Transmission networks consist of hundreds of thousands of miles of high-voltage power line, which transfer power from bulk generators to distribution substations or to other power grids, depending on demand and market contracts.

Distribution: Receives high-voltage power from transmission lines, converts it to a safe voltage, and delivers it directly to customers. Distribution networks include step-down substations for converting power from high-voltage lines, as well as mid-voltage power lines that transfer power throughout a limited geographical region to customers. Distribution networks also accept power generated by its customers; the power-provision relationship is not unidirectional.

Customer: Any consumer of power from the grid. Customers are typically partitioned into residential, commercial, and industrial domains [4]. They pay for access to power from the grid, but may also generate their own, often by harnessing renewable energy like solar power. Some customers even feed excess power back into the distribution network.

Markets: Despite being a public utility, the power grid is a for-profit entity. Markets exist to both set the price of electricity for customers, as well as to facilitate the trade of power between different sub-grids and their independent system operators (ISOs).[1]

[1]ISOs are the administrative operators of sub-grids. They ensure

Service Provider: Responsible for ensuring customers receive electricity, providing services like line repair and meter reading and connection.

Operations: Coordinate between the various entities in the power grid. Operations is responsible for ensuring the following:

- Power generated meets customer demand.
- Power is transmitted and distributed to customers efficiently.
- Equipment operates at safe conditions.
- Equipment failures are handled appropriately.
- Market decisions regarding pricing and trading of power are implemented.

In short, operations must ensure that the grid is always providing power to its customers within the limits of safety and generational capacity. It is important to note that operational decisions are influenced by the electricity market – particularly, pricing and trading decisions made by ISOs – and are not constrained *solely* by customer demand and equipment limitations.

3.3 Control Loops in the Grid

In order to uphold the guarantees made in the previous section, the grid's operators use several primary control loops to maintain the complex balance between the grid's generators and loads. The failure of any of these control loops to maintain their equilibria would result in the failure of the grid to provide its services to some extent. As such, **disruption of any operational control loop is the goal of an attacker[2] of the power grid**. We now describe each of the grid's primary control loops, enumerated by Sridhar *et al.* [62].

Automatic Voltage Regulation: Used to stabilize the output voltage of the grid's power generators [77]. It monitors the output voltage of one or more generators, and adjusts the input control voltage to the generators' exciters accordingly to keep output voltage levels stable.

Governor Control: Used to control the frequency of the grid's power generators [56]. It monitors the speed of an individual generator's turbine and adjusts the steam valve to account for deviations from its frequency setpoint. Frequency control is particularly important in the power grid, as coordinated generators expect a nominal frequency, and could be damaged if the actual frequency deviates from this too greatly [16].

Automatic Generation Control: Also used to control the frequency of the grid's power generators [71]. It is responsible for ensuring that the entire grid's load is met by its generators, while the power being traded between sub-grids meets its market-set values. It monitors the frequency and power flow through tie-in lines,[3] as well as the load on the system as a whole, and uses this information to adjust the output of each generator in the system accordingly.

State Estimation: Used to monitor and regulate the overall state of the power grid's transmission and distribution networks [48]. System state is represented by a set of variables, each of which is a value (like voltage magnitude or phase angle) at a specific point in the grid. State estimation receives measurements from every sensor in the grid, applies error correction to this set of measurements to account for faulty devices, and uses the result to "estimate" the grid's state. Based on this estimate, adjustments are made to the grid's topology and generator output to increase routing efficiency, protect equipment from damage, and satisfy shifting load.

Static VAR[4] Compensation: Used to improve the efficiency of power transmission in the grid by regulating line voltage [76]. It monitors the reactive load on individual transmission lines. If it determines the load is leading, it uses reactors and synchronous condensers to absorb reactive power from the line; if the load is instead lagging, it triggers capacitor banks to increase the line voltage. Both reactions attempt to adjust the line's power factor to unity for maximum efficiency and stability.

Wide Area Monitoring: Used to monitor the health of the grid's transmission and distribution networks and to protect them from failure [74]. Wide area monitoring uses phasor measurement units (PMUs) to measure AC phasors throughout the grid and synchronize them with GPS, providing a real-time view of the entire grid's state in terms of voltage and phase angle. Wide area monitoring determines grid state more rapidly than state estimation. As such, it is used to automatically make split-second adjustments to grid topology to prevent failures from occurring or escalating, decisions which a human operator might not be able to make quickly enough [8].[5] Wide area monitoring has the potential to replace state estimation and automatic generation control, but is not currently widely deployed [65].

Load Shedding: Used to maintain the balance between power generated and power consumed, and to prevent system collapse in emergencies [73]. It monitors the power output of the grid's generators and the load on the

that bulk generation meets customer demand economically, trading power wholesale with other sub-grids when it is more profitable to do so.

[2]Recall that we define our attackers' goal to be damaging/disabling the power grid.

[3]Tie-in lines are power transmission lines which connect two separate sub-grids which wish to exchange power between one another.

[4]Volt-amperes reactive (VAR), the standard unit of reactive power

[5]The 2011 Southwest blackout occurred in less than 60 seconds. Reaction time is critical to the grid's protection.

entire grid, and trips breakers to disconnect various loads that it determines cannot be served safely.

Demand Response: Smart meters and advanced metering infrastructure (AMI) allow for load shedding at a fine granularity, a process known as demand response [73]. Like load shedding, it monitors the grid's generational output and capacity, as well as the load on the grid *per customer*. When generation falls short of satisfying load, demand response can disconnect individual customers, or even customers' selected appliances, to shed load without causing large-scale blackouts.

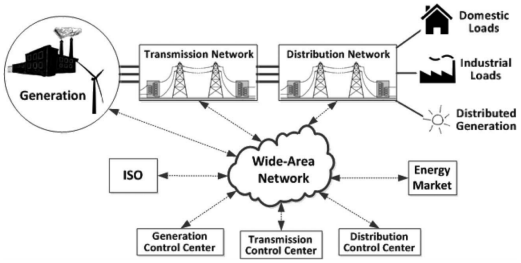

Figure 2: Communication network in the power grid's EMS system [62]

3.4 Energy Management System

The previous two subsections depict a grid that requires the coordination of multiple distinct actors to maintain several distinct equilibria [25] [57]. The EMS (which we introduced in Section 1.1) makes this coordination possible.

The EMS is fundamentally a SCADA (supervisory control and data acquisition) system.[6] It is an amalgamation of **sensors**, **actuators**, **control systems**, and a **communication network** that connects them all. It is the nervous system of the power grid, allowing the grid's operators to maintain its critical control loops. We now describe the components that compose the EMS:

Sensors: The EMS relies on a network of distributed sensors to monitor the status of all the components in the grid. Sensors known as IEDs (intelligent electronic devices) and PMUs record the voltage, phase angle, and frequency on power lines and at substations [20] [62]. RTUs (remote terminal units) also sometimes act as sensors at substations. In areas where AMI is deployed, smart meters measure customer power consumption and load. And all generators use sensors to monitor internal processes, like rotary speed sensors to measure turbine speed [4].

Control Systems: The EMS must act on the information gathered by its sensor network; this is the role of its control systems. Control systems are software suites running in the grid's control centers. They collect data from the grid's sensor network, process it using certain models like the state estimation model, and either act on or display the processed information. In some situations, they make decisions automatically. For example, when sensory data indicates impending equipment failure, the control systems can automatically disconnect equipment from the grid, initiate load shedding, and raise alarms. Control systems running generators, referred to as "local" control systems, automatically maintain certain conditions like output frequency. But control systems are not entirely automated; they allow human operators to observe the grid's state and manually issue commands through an HMI (human-machine interface), which can be accessed in person at a control center or remotely through an online interface. Control systems also provide data to the electricity market.

Actuators: When the EMS's control systems make decisions, they issue commands to actuators in order to affect physical change. The primary actuators in the grid's transmission and distribution networks are circuit breakers, which can connect or disconnect lines, substations, and entire distribution networks to or from the grid. Modifying grid topology is useful for load shedding, power routing, trading power between grids, and protecting failing equipment. Smart meters also contain circuit breakers that allow for connection and disconnection of individual customers, eliminating the need for service technicians to manually modify a customer's connection to the grid. Further, generators have a diverse array of actuators which allow the modification of its connection to the grid, as well as maintenance of internal conditions like pressure, turbine speed, and heat.

Communication Network: The transmission of sensor data to control systems and commands to actuators occurs on the EMS communication network. In transmission and distribution, sensor data is sent to RTUs at substations, where it is collated and sent to control systems. Commands issued from the control systems to actuators in the transmission and distribution networks follow the same route, traveling from the control systems to substations, and then to any actuator not physically in the substation [60]. In AMI, smart meters and control systems exchange data and commands directly [12]. In generation, actuation commands and sensor data ordinarily pass through a programmable logic controller (PLC), which serves as a middle-man between the control systems and the generator's actuators and sensors [23]. Figure 2 depicts the grid's EMS communication network, while Figure 3 shows a the typical communication flow inside a generator.

[6] SCADA systems are industrial control systems that provide their operators monitoring and control capabilities over equipment through sensors and remote-controlled actuators, respectively.

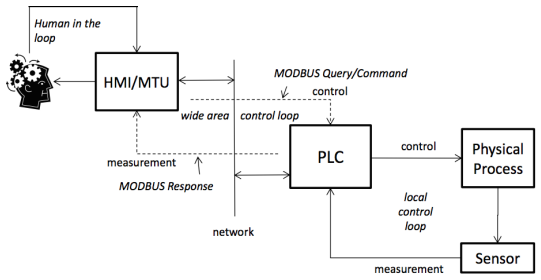

Figure 3: Communication flow inside a generator [23]

EMS Summary: The EMS's control systems use its network of sensors to observe the state of the grid as a whole and to determine when problems arise. When changes to the grid are required to preserve any one of its control loop equilibria,[7] commands are issued to actuators in the grid, which affect the desired change. The EMS is a large-scale feedback control system [69].

3.5 Cyber-Physical System Model

We now use our understanding of the EMS as a feedback control system to construct a cyber-physical system model of the power grid. The discreet entities involved in our model are **physical devices** (sensors, actuators, meters), **control systems** (all software used to make control decisions), and two communication channels: **data acquisition** and **control**. The data acquisition channel relays sensor data to the control systems, and the control channel conveys commands to actuators. We depict this model in Figure 4, and note that it is a prototypical illustration of a cyber-physical feedback control system.[8]

The cyber-physical system model of the power grid has an advantage over conceptual models of the grid in terms of identifying exposed attack surfaces. The conceptual model is partitioned by function, but some domains – transmission and distribution, in particular – have significant overlap in terms of electronic attack vectors.

The cyber-physical model is, rather, partitioned by role in the grid's control feedback loop. This model mimics the form of the control loops which the grid's attackers are attempting to disrupt, and allows for attacks to be classified by the components of the control loops they affect. It is thus a logical way to examine and dissect attacks on the grid.

[7]Changes to the grid could include the desire to change the route of power supplied, to adjust the amount of power generated, to trade power with other grids, to load shed, to relieve equipment, to connect or disconnect customers, or to deal with any number of problems that arise during ordinary grid operation.

[8]Sridhar *et al.* [62] employ a similar model of the power grid.

4 Exposed Attack Surfaces

An attack surface is the aggregation of all the individual vectors by which an attack may be launched [70]. In this section, we consider each component of the cyber-physical system model of the grid to be an individual attack surface. We describe the vectors on each surface which the EMS has exposed to electronic attack.

4.1 Out-of-Scope

It is useful to first define the scope of the attack vectors and attacks that we are considering in this paper. The purpose of this paper is to explore various electronic attacks on the power grid. We therefore disregard the following:

Physical Damage: Causing physical harm to, or destruction of, components in the power grid are not viable attack vectors, as these types of attacks are obviously not electronic. Though we will not consider them in this paper, such attacks are quite effective at disrupting the grid [68].

Market-based Attacks: Attacks that involve manipulating the electricity market as the vector through which to affect the operation of the grid are out-of-scope, as they do not take advantage of any attack surface presented by the EMS. We *will*, however, be considering attacks where manipulating the electricity market is the end goal in Section 5.1.

4.2 Communication Channels

Both the data acquisition channel and the control channel utilize the same communication infrastructure, so they share the same exposed attack surfaces. As grid components are frequently distributed over large geographical areas, the data acquisition and control channels present some of the most accessible attack surfaces in the entire grid.

Remote Communication: There is no standard infrastructure for the data acquisition and control communication channels between the grid's distributed components[9] and its control systems. As the U.S. power grid is in reality a patchwork of interwoven sub-grids operated by different utilities, it follows that each utility has implemented these channels to meet its own set of requirements. For example, a utility that controls a remote hydroelectric power plant may prefer wireless microwave communication, while a utility controlling a natural gas

[9]Distributed components are smart meters, sensors, switches, and even substations. Remote communication refers to the communication between these entities, as opposed to internally in any individual component.

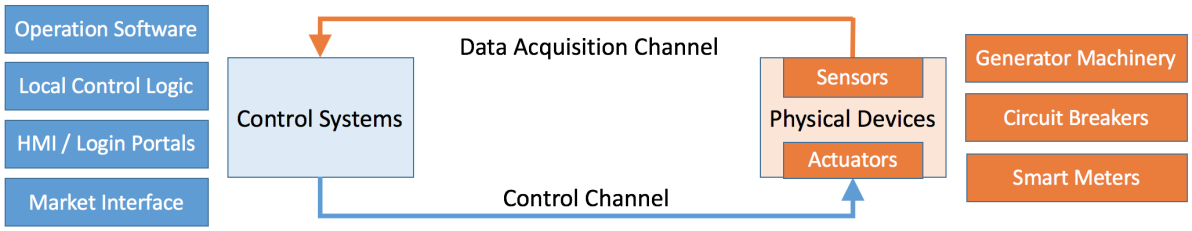

Figure 4: Power grid EMS system as a cyber-physical system

plant that serves its surrounding community might employ efficient, but costly, fiber optic cable.

The following are commonly used as infrastructure to facilitate the grid's long distance communication channels [61]: leased telephone lines; private telephone lines; power-line carrier systems;[10] microwave communication systems; satellite communications; 900MHz point-to-multipoint radio; 900MHz mesh radio; fiber optic cable; public Internet; and public cellular services.

Analog and digital microwave communication systems are the most prevalent in the U.S. due to their considerable range and low cost. AMI is an exception, using public cellular services. But given the variety of choices and implementations, the utility companies have agreed on several communication protocol standards so that their sub-grids can interface. DNP3, Modbus, and IEC 61850 are the common networking protocols used in the U.S. power grid [47] [61].

We consider all of these communication system implementations to be exposed attack surfaces, as they all exist in public spaces, are ordinarily unprotected, and are distributed across the country [1]. However, some are more accessible to attackers than others. All wireless communication systems, including the widely used microwave communication, are significantly more vulnerable to sniffing, spoofing, and denial of service than their wired counterparts.[11] The standardization of communication protocols, combined with their mandated open sourcing by FERC [61] and general lack of authentication [47], lower the barrier for attacks on these networks significantly.

Local Communication: We also consider the communication channels between control systems and PLCs inside generators to be attack surfaces. These channels are often wired and unexposed, but not always. Maynor and Graham [40] provide anecdotal evidence of U.S. power plants whose internal SCADA systems involve wireless communication.

[10]Communications piggy-back on actual power lines in power-line carrier systems.

[11]As physical attacks are out-of-scope, wired communication systems hardly present an attack surface at all.

4.3 Control Systems

The EMS's control systems are proprietary software suites that run on commodity computers inside the grid's control centers. They are designed to be accessed by operators physically present in the control center as well as remote operators, and therefore expose authentication portals to both utility intranets as well as the public Internet. The public Internet-facing authentication portal is an exposed attack vector due to its potential availability to unknown parties on the Internet (especially if misconfigured), while the intranet authentication portal is an attack vector to any compromised device with access to the utility intranet.

Authentication portals are not the only concern. Control systems which expose network interfaces leave themselves open to remote exploit, whether against the control software itself or the operating system supporting it. Proprietary software running on outdated commodity computers and operating systems is likely to have some vulnerabilities. Additionally, using commodity computers and operating systems exposes the control systems to infection by malware; even if the control system computers are air-gapped, they could be infected by USB drives and removable media.

Finally, operators themselves are considered viable attack vectors, whether through social engineering and deceit or as an insider attacker.

4.4 Physical Devices

We consider the attack surfaces exposed by the two distinct classes of physical devices in the grid: those distributed throughout the grid, and those used inside power generation plants.

Remote Components: As we have discussed, the grid's EMS relies on a distributed network of sensors and actuators to perform its SCADA operations. As of 2010, the U.S. power grid had "tens of thousands, if not hundreds of thousands" of remote IEDs in use alone [61]. Including RTUs, PMUs, and smart meters, this number would reach into the millions.

These sensors and actuators are small, embedded de-

vices distributed geographically throughout the grid, and are frequently physically unprotected. Although we do not consider destruction of equipment to be in-scope, physical access to embedded devices allows for legitimate attacks like firmware rewriting [14]. Additionally, many embedded devices allow for direct wireless or wired connection for debugging and configuration over channels which can be accidentally left unsecured in production [15]. Given that these devices are low-cost, mass-produced embedded devices, such expectations of relative insecurity are reasonable.

As such, we consider all of the following remotely-distributed components to be viable attack vectors, both through physical access and direct wireless access: IEDs, PMUs, RTUs, smart meters, circuit breakers, and even specialized equipment like dynamic transformers.

Local Components: Inside power generators, local components like sensors, actuators, and PLCs control and safeguard power generation. Unlike their remote counterparts, local sensors and actuators are hard-wired to the PLCs that control them and are embedded in machinery, so we do not consider them viable attack vectors. However, the PLCs themselves are accessible, both physically and remotely. An insider attacker with physical access to a PLC can exploit operating system backdoors or vulnerable running services, both of which are known to exist [44] [75]. Further, some PLCs communicate wirelessly with their control systems, a channel over which exploits could be attempted by a remote attacker. And finally, an attacker that has compromised the local control system itself inherently has access to all its PLCs. We therefore consider the PLCs inside power generators to be viable attack vectors.

5 Attacks

We now explore several chosen attacks that demonstrate the vulnerability of each of the EMS's exposed attack surfaces, and the potential consequences associated with their success.

5.1 False Data Injection Attacks

State estimation [48] and automatic generation control (AGC) [71] are currently the two most pervasive control loops in the power grid. Combined, their job is to constantly monitor the state of the grid, providing power to satisfy changing demand, modifying grid topology (using circuit breakers) to route energy more efficiently or prevent line failure, and providing a detailed view of the grid to operators and markets alike. This large-scale monitoring is made possible by the distributed network of sensors providing complete coverage of the grid. These sensors report power flow information to the grid's control systems, where it is filtered for bad data, and then used by state estimation and AGC.

False data injection attacks refer to the addition of false data to the information being reported to the control systems by the grid's sensor network. An important distinction between false data injection attacks and any other attack in which spoofed data is sent to the control systems is that the data in a false data injection attack is intended to be stealthy; that is, the entities processing the false data should not be able to distinguish it from legitimate sensor data. We elaborate upon the first published false data injection attack, and then browse the ways in which this class of attack has been expanded by subsequent studies.

5.1.1 Attack Vectors

The attack vectors of false data injection attacks are the communication channels between sensors and control systems (falling under the Data Acquisition Channel attack surface). Nearly all remote sensors in the grid[12] report their measurements over wireless microwaves using insecure protocols like DNP3, making these communication channels easily susceptible to attack and manipulation.

Another possible attack vector is direct compromise of the sensors themselves, so as to corrupt their readings (part of the Physical Devices attack surface). This is not as scalable as attacking the sensors' communication channels, but is still feasible given the general insecurity of IEDs.

5.1.2 Methodology

Liu *et al.* [37] first introduce the false data injection attack, specifically against the state estimation control loop. To understand the attack, and specifically why it is undetectable, we must explore how bad measurement detection works in state estimation.

Bad Measurement Detection in State Estimation:

State estimation has always been sensitive to bad data. Low-cost embedded sensors are unreliable, and long-distance wireless communication prone to corruption, so bad sensor data needs to be detected and discarded before an accurate model of the power grid can be developed.

The following DC power flow model of the grid is used in bad measurement detection due to its simplicity over AC models:

$$z = Hx + e$$

In this model, H is the DC power flow matrix, z is the vector of system measurements, x is the vector of the actual system state, and e is error. Below, we show how $\hat{x}$,

[12] This does not include local sensors inside generator plants.

an estimator for x, is determined using system measurements and meter error variance W:

$$\hat{x} = (H^T W H)^{-1} H^T W z$$

The estimator $\hat{x}$ is used detect the presence of bad measurements as follows. The measurement residual $z - H\hat{x}$ is the difference between observed and expected (estimated) measurements. If its L_2 norm is greater than a preset threshold τ, as shown below, then bad measurements are present.

$$||z - H\hat{x}|| > \tau$$

False Data Injection Attack:

Unfortunately, bad data detection was meant to catch sensor data that deviated randomly. As Liu *et al.* demonstrate, measurement data can be corrupted in specific ways so as to pass the detection threshold. Let z_a be the measurements reported by the sensors, with a being the malicious data the attacker adds to the real signal.

$$z_a = z + a$$

As the measured data is corrupted, so too is the estimator. Let $\hat{x}_{bad}$ be the estimator forged from the false measurement data:

$$\hat{x}_{bad} = \hat{x} + c$$

In this equation, c is a vector representing the deviation of the estimator from its expected value induced by the modified measurement data. The authors propose crafting an attack vector a as such:

$$a = Hc$$

With an attack vector fitting this form, the measurement residual falls within expected ranges, ensuring that the corrupted z_a will be accepted as a legitimate set of measurements. We show this below:

$$||z_a - H\hat{x}_{bad}|| = ||z + a - H(\hat{x} + c)|| = ||z - H\hat{x}|| \leq \tau$$

Requirements:

The strongest requirement of this attack is the attacker's knowledge of the DC power matrix H, which requires full knowledge of the topology of the grid. However, configuration documents containing this information are frequently stored on substation control system computers, which are a known target of APT campaigns [54]. Given a resourceful attacker, this requirement is by no means insurmountable.

A false data injection attack can also be limited by the number of sensors the attacker can control, due to access restrictions or resource limitations. Likewise, the attacker may only wish to target certain state variables. Based on the attacker's resource limitations and targeting preferences, a possible attack may not exist.

Overview of Improvements:

Kosut *et al.* [31] [32] [33] verify the attacks that Liu *et al.* propose. They develop metrics to categorize worst-case attacks given a set of restrictions, as well as the least number of compromised sensors required to achieve a desired outcome.

Ozay *et al.* [51] expand upon the notion of finding the least number of meters required for a succesful attack. They propose a method of developing sparse false data injection attacks so as to avoid detection more effectively.

Esmalifalak *et al.* [19] and Rahman *et al.* [53] lower attacker requirements significantly by eliminating the need for preliminary knowledge of the grid's configuration. Their methods involve passively obtaining as much information about the admittance of the grid as possible, and using this partial knowledge to exploit *subsections* of the grid for which admittance (and therefore H) can be determined with some degree of confidence.

5.1.3 Results & Potential Outcomes

False data injection attacks has numerous consequences:

State Estimation Error: Liu *et al.* [37] quantify the consequences of false data injection attacks on the state estimation control loop experimentally. Figures 5 and 6 show the number of meters an attacker needs to compromise to inject false data into a number of targeted state variables, given an attacker constrained to a set of vulnerable meters and unconstrained, respectively.

State estimation is responsible for protecting the grid's equipment from failure, ensuring efficient energy routing, and satisfying changing load. The consequences of its improper function could therefore be damage to overloaded power line, which can result in cascading failures, inefficient energy routing, which costs the utilities money, and failure to satisfy load, which can result in load shedding or generator disconnection. The manipulation of just a few state variables could lead to these outcomes, which is entirely feasible according to the experiments run by Liu *et al.*

Some authors have expanded upon false data injection attacks on state estimation. Kim *et al.* [30] demonstrate an attack on the topology of the grid. They formulate their attack such that, using false data injection, attackers can deceive state estimation into believing incorrectly that chosen power lines are connected or disconnected, as in Figure 7. This attack is able to convince state estimation that a line whose circuit breaker is open is carrying power. Such an attack could be used to overload and damage lines by hiding them from the operator, or to force inefficient routing of energy.

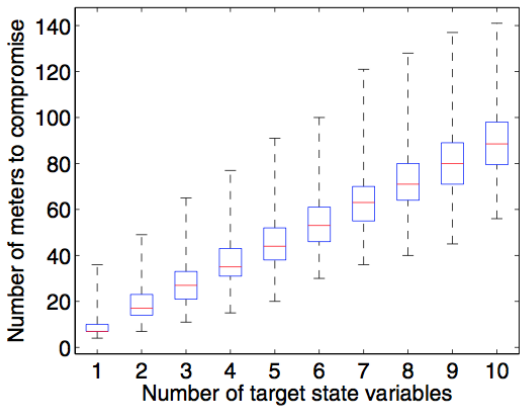

Figure 5: Number of meters to compromise to affect targeted state variables in IEEE 300-bus simulation - constrained case [37]

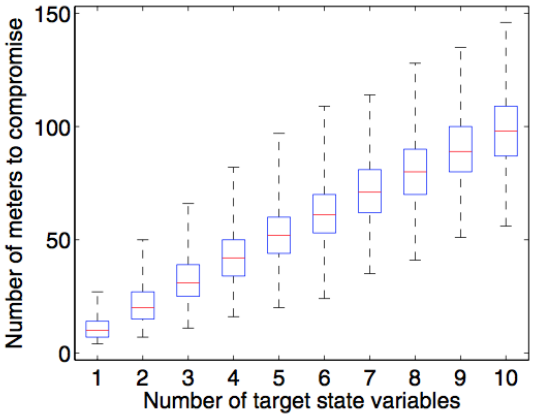

Figure 6: Number of meters to compromise to affect targeted state variables in IEEE 300-bus simulation - unconstrained case [37]

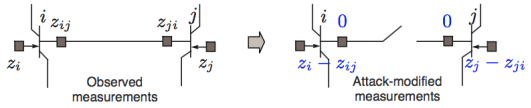

Figure 7: Illustration of false data injection attack against grid topology [30]

Attacks that directly give the attacker the ability to manipulate energy routing also exist. Lin *et al.* [35] describe the impact of false data injection attacks against state estimation on energy routing efficiency. Figures 8 and 9 show how the number of nodes compromised directly affects the energy lost in transmission and the total cost of energy transmission, respectively.

Automation Generation Control Error: Sridhar *et al.* [63] consider the effects of false data injection attacks against AGC. In AGC, sensor data regarding system frequency and power flow through tie-lines is used to determine if the load on any part of the system is increasing or decreasing. By manipulating this sensor data, an attacker is able to convince a generator that its neighboring load has increased, causing the generator to increase its power production inappropriately. Table 1 shows the generation-load imbalance this attack against AGC creates in a two-generator simulation.

Market manipulation: The data processed by state estimation is used in the real-time energy market. In certain ISOs, trading on virtual (day-ahead) power is a legitimate market. According to Xie [78], "A market participant [may] purchase/sell a certain amount of virtual power P at [a] location in [the] day-ahead forward market, and will be obliged to sell/purchase the exact same amount in the subsequent real-time market."

Jia *et al.* [27] and Xie *et al.* [78] both present attacks in which an adversary can successfully make a profit by defrauding the real-time energy market. The fraudster makes the decision to buy and sell power in the day-ahead market, and the next day manipulates the prices of the real-time (ex-post) market using false data injection attacks such that the power purchased the day before is now worth more, while the power sold is worth less.

5.1.4 Overview of Proposed Defenses

Manandhar *et al.* [38] propose the use of more complex models, known as Kalman filters, to detect bad data. Unfortunately, the processing cost of using models more complex than the current bad data detection schemes often makes them unusable in the grid's time-critical applications. To resolve this, Liu *et al.* [36] proposed partitioning the grid into sub-grids for bad data detection, and applying complex techniques like Kalman filtering to each sub-grid for expediancy. Though not implemented currently (to the author's knowledge), this has promise.

Bobba *et al.* [11] propose protecting a set of meters to ensure full network observability, which will guarantee bad data's detection. The shortcoming with this method is that some subset of meters must be absolutely protected, a warranty that is difficult to fulfill.

5.2 AMI Attacks on Grid Health

Advanced metering infrastructure [67] "is an integrated system of smart meters...that enables two-way communication between utilities and customers." AMI is the infrastructure that supports demand response, the control loop that provides load shedding at the individual customer, or even individual appliance, level of granularity.

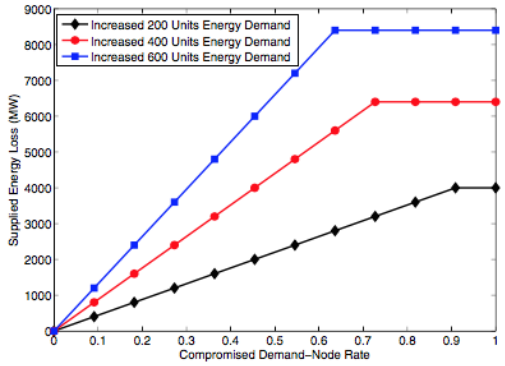

Figure 8: Energy lost in transmission per number of nodes compromised [35]

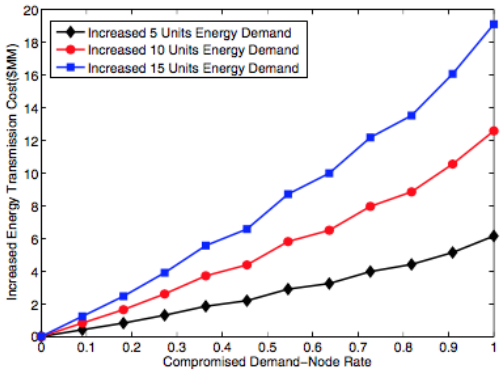

Figure 9: Cost of energy transmission per number of nodes compromised [35]

Parameter	Before Attack	After Attack
Frequency (Hz)	60	60.156
Tie-Line Flow from Area 1 (pu)	0.4	0.4049
Unit 1 Generation Change (pu)	0	0.01
Generation-Load Imbalance (pu)	0	0.01

Table 1: Generation-load imbalance in a two-generator system caused by false data injection attack against AGC [63]

The goal of demand response is to implement load shedding that does not affect customers as severely as planned blackouts, but that still provides a comparable amount of load reduction [12]. There are two current implementations of demand response: dynamic pricing, and direct load control [9].

In dynamic pricing, customers are motivated to conserve electricity at peak hours with price controls, such as increased rates at peak hours or rebates for lowered energy consumption. In direct load control, utilities have the ability to remotely disable customers' power appliances during times of high demand.[13] These power appliances would generally be non-essential, high-load devices like air conditioning units, such that load could be shed without antagonizing customers [21].

We now explore attacks against AMI that intend to cause actual damage to the power grid, beyond the scope of energy theft. We consider attacks against both forms of demand response.

5.2.1 Attack Vectors

The vectors of these attacks are smart meters. An attacker must either be able to issue commands to the smart meters and manipulate their responses (the Control and Data Acquisition Channel attack surfaces), or directly compromise the smart meters (the Physical Device attack surface).

Issuing commands to smart meters is often trivial, as is manipulating their responses. In the U.S., there is a lack of diversity in the smart meter market. As mass-produced, low-cost embedded systems, they are likely to have design flaws with regards to security [14] [15]. For example, Illera and Vidal [26] show that a brand of smart meter widely deployed in Spain uses the same asymmetric key for all of its communications. Every meter of that brand uses the same key, which the authors extracted from a sample meter with relative ease.

For the same reasons, compromising smart meters directly is also often trivial. C4 Security [12], Pollet [52], and McLaughlin *et al.* [42] experimentally explore various methods by which name-brand smart meters currently in the U.S. market can be compromised, from debugging backdoors and lack of authentication to firmware upgrade vulnerabilities.

As the attacks we are discussing require access to a large number of smart meters, we consider how widespread compromise may be achieved. Davis [17] presents a proof-of-concept smart meter worm that spreads between meters via their built-in radios, offering an attacker control over many devices having compromised just one.

[13] Direct load control is only installed with customer agreement.

5.2.2 Methodology

Barreto *et al.* [9] present a series of attacks against AMI. We focus on two scenarios: a malicious attacker who wished to harm the grid in a direct load control environment, and the same attacker in a dynamic pricing environment.

Direct Load Control: In this scenario, the attacker is able to automatically control the level of power provided to all connected appliances. With this capability, the attacker may choose any one of the following attacks:

1. Block or ignore all commands to victim smart meters, but confirm their receipt. During high demand, this will prevent requested load from being shed, while the demand response controller will make decisions in the short-term as if it had. Likewise, at times of low demand, demand response will make short-term decisions as if load had been restored.

2. Issue commands allowing appliances to run at their highest power during peak demand, directly causing an additional spike of demand.

3. Issue commands disabling appliances during periods of low demand, directly causing further loss of load.

Dynamic Pricing: In this scenario, the attacker is only able to offer pricing incentives to customers with regards to the real-time price of energy. However, customers in this scenario often employ automated appliances that will defer their operation until energy prices are below a certain threshold.[14] The ability to adjust the real-time price of energy therefore affords the following attacks:

1. Increase the price of energy during a period of low demand. Customers with automatic deference applications will further drop load from the grid.

2. Decrease the price of energy during a period of high demand. All customers' devices that were awaiting lowered prices will activate, resulting in a spike in demand.

3. Strategically keep the price of energy artificially high prior to a period of expected high demand, in order to force automatic deference appliances to deactivate. Once high demand does occur, lower the price of energy. The deferring appliances will all connect simultaneously, causing a large spike in demand.

Figure 10 shows how attacks 2 and 3 affect load in a simulation. The superiority of the strategy employed in attack 3 is apparent.

5.2.3 Results & Potential Outcomes

Sudden, unexpected spikes in demand are hard to account for, regardless of the monitoring and protection systems in the grid. An attacker employing one of these demand-spiking attacks may be able to force control loops like state estimation and wide area monitoring to load shed or disconnect generators from the grid as a protective measure. This could result in blackouts or even cascading failures depending on the generators' criticality. If this attack were coordinated with another event, like a natural disaster or a severely hot day, there is potential for significant damage.

[14] This deference option is particularly attractive to industrial-grade customers.

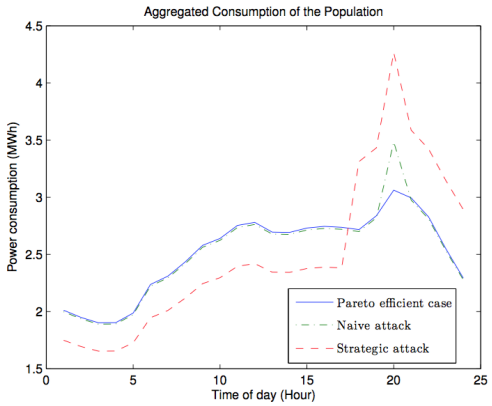

Figure 10: Demand under attack in dynamic pricing scenario [9]

Though perhaps not as severe, the attacks in which load is dropped during a period of low demand can still introduce frequency imbalance to the grid. Load on the grid is required for its proper function, an unexpected drop in which could potentially force a generator to disconnect from the grid and shut down to maintain balance. Were this to happen just before a period of high demand, the grid could be forced to load shed.

In terms of quantifying the impact of this attack, the spike in demand (or drop in load) is directly proportional to the number of smart meters compromised.

We also note that AMI is not only susceptible to attacks intended on damaging the grid. McLaughlin *et al.* [42] prove that smart meters can be modified to steal energy for customers, as do Barreto *et al.* [9]. Rouf *et al.* [58] demonstrate that smart meters' insecure broadcasts can be used by attackers to profile households' daily patterns, including when residents are present or what appliances might be running.

5.2.4 Overview of Proposed Defenses

The issue with protecting against these attacks is that they exploit legitimate features of demand response from the controller's point of view. Two proposed defenses exist: first, use redundant control loops like AGC, state estimation, or wide area monitoring to catch demand spikes and drops immediately (this is the defense currently implemented); or second, provide better perimeter security

to AMI, in the form of secured meters and communication channels. This is ultimately the better defense, but its cost is prohibitive to utilities at present time.

5.3 GPS Spoofing Attacks on PMUs

Phasor measurement units are the backbone of the wide area measurement control loop. Distributed networks of PMUs provide a real-time view of the state of the grid by synchronizing their clocks using the Global Positioning System (GPS) and using this to time stamp their voltage and current measurements. By spoofing GPS signals to PMUs, they can be desynchronized, resulting in misreported information to the wide area measurement controller.

5.3.1 Attack Vectors

PMUs are the attack vector of this attack (part of the Physical Devices attack surface); specifically, the signals into their GPS clocks' receivers.

GPS spoofing attacks have been studied in detail. Tippenhauer *et al.* [66] present a set of requirements for launching a successful GPS spoofing attack against individual and groups of GPS receivers. Nighswander *et al.* [50] elaborate on possible ways to launch GPS spoofing attacks (among others) and confirm their feasibility by attackers with few resources.

5.3.2 Methodology

Shepard *et al.* [59] propose a replay attack against the GPS signal received by a PMU. The goal of their attack is to induce a time offset of at least 26.5μs to the victim PMU's GPS clock. Such an offset will manifest itself in a a 0.573° phase angle difference between the victim PMU and other PMU's in its network, as shown in Figure 11. In order to induce this time offset, the authors acquire the real GPS signal[15] and produce a counterfeit signal identical to the captured signal. They then increase the power of their counterfeit signal until it is more powerful (to the victim) than the authentic signal, at which point their signal has taken control of the victim PMU's GPS receiver. They then introduce time delay incrementally to the counterfeit signal until they have successfully inducing a 26.5μs time offset on the PMU's clock.

Jiang *et al.* [29] expand upon the work done by Shepard *et al.* [59]. Rather than simply cause uncontrolled phase angle difference to a victim PMU's measurements, the authors attempt to determine the *maximum* error they can introduce through spoofed GPS signals, where error is the "difference between the spoofed clock offset and the pre-attack clock offset." They recognize that

[15] Specifically, the GPS L1 C/A and L2C signals.

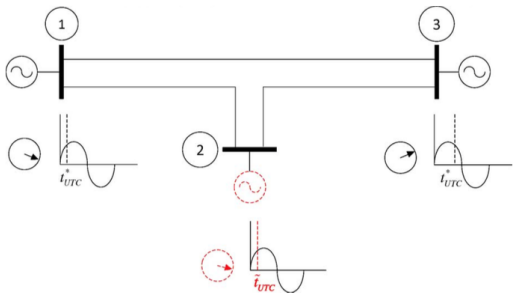

Figure 11: PMUs and an introduced phase angle difference [29]

GPS receivers use the following to determine their position and clock offset: ephemerides,[16] satellite positions (computed using the ephemerides), and pseudorange.[17]

The final variable, pseudorange, is the variable which Shepard *et al.* [59] affect by time-shifting the authentic GPS signal. The other variables, ephemerides and satellite positions, can be affected by modifying the GPS signal rather than simply delaying it. By manipulating all three of these variables in the spoofed GPS signal fed to the victim PMU, controlled phase error can be induced. During testing, Jiang *et al.* [29] were able to introduce a phase angle error of more than 50° to a victim PMU.

5.3.3 Results & Potential Outcomes

Wide area measurement is used to monitor the health of the grid and make automatic decisions to protect equipment from failure using remote circuit breakers. Specifically, the system uses the phase angle difference between two PMUs as an indicator of a fault, which can cause the system to automatically trip breakers to prevent equipment damage. By spoofing the GPS signals to a victim PMU and changing the phase angle difference between it and another PMU, two potential outcomes can occur:

False Alarm: The spoofing could increase the phase angle difference between the two fixed PMUs, resulting in an automatic breaker trip. The possible effects of disconnecting a line in the power grid are many-fold: a generator could be disconnected from the grid, forcing it to shut down; important HVDC lines between grids could be disconnected, unbalancing demand and forcing load-shedding; entire distribution networks could be blacked out. As an illustration of this attack's potential, the lines connecting a hydroelectric plant to its load in Mexico will trip if a phase angle difference above 10° is detected,

[16] Ephemerides are a set of constants of integration known as the Keplerian elements, as well as several additional variables, that can be used to determine the trajectory of a satellite.

[17] Pseudorange is the approximate distance between a satellite and a receiver.

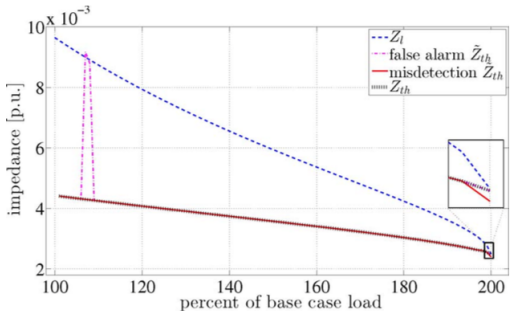

Figure 12: Real vs. forged Thevenin impedance under PMU GPS spoofing attack [29]

well within the tested capabilities of the author's attack [59].

Hidden Fault: The spoofing could decrease the phase angle difference between the two fixed PMUs, hiding faulty conditions from the wide area management system and inhibiting its automatic protection [29]. If other monitoring control loops like state estimation do not detect the situation in time, this attack could lead to failure of important power lines or damage to generators which have been disconnected from their load unknowingly.

Figure 12 shows both a false alarm and a hidden fault attack induced by PMU GPS spoofing in a slightly different scenario, wherein alarms are triggered if Thevenin impedance surpasses load impedance – phase angle difference between PMUs in the system directly affects perceived Thevenin impedance At 110% of base case load, we see an artifact corresponding to a false alarm being triggered. This was caused by a shift of -10°. At 200% of the base case load, we see a second artifact corresponding to a hidden period of instability. This was caused by a shift of 2°.

5.3.4 Overview of Proposed Defenses

There are two main defenses proposed against such attacks:

First, wide area monitoring bad data detection could be modified to validate the timestamps reported by the PMUs, in addition to the measurement values. As its current implementation is directly adopted from state estimation, according to Choi *et al.* [13] and Terzija *et al.* [65], wide area monitoring bad data detection does not currently appear to validate PMUs' GPS timestamps for validity.

Second, there has been substantial work done in the area of GPS spoofing countermeasures. Tippenhauer *et al.* [66] and Jiang *et al.* [29] both suggest that GPS spoofing countermeasures could defeat these attacks; the only inhibition is cost to the utilities to add antennas/receivers to current GPS units, or to implement software that provides redundancy checking with fixed, trusted sources.

5.4 Malware Attack on Unknown PLCs

In the grid's power generators, programmable logic controllers (PLCs) are used to automatically control and maintain the power generation process. In addition to implementing automatic voltage regulation and governor control, which ensure that the output of the generator meets grid voltage and frequency requirements, a generators' PLCs also automatically maintain safe operation of the generator itself. For example, the PLCs in a nuclear plant are responsible for monitoring the internal temperature of the plant and increasing the supply of cooling water as needed. As such, PLCs are an essential part of power generation.[18]

Comprehensive attacks against PLCs have been studied for some time. Milinkovic *et al.* [45] present the results of basic reverse engineering efforts against five common industrial PLCs, finding at least one major vulnerability in each. Beresford [10] performs a more comprehensive reverse engineering effort against the Siemens Simatic S7 line of PLCs (a popular line in the power industry), and finds it vulnerable to replay attacks, memory dumps, and a host of other exploits. These findings are serious, as the implications of a successful attack on the internal components of a power generator can be severe; an attack which prevents a nuclear plant from cooling properly could cause an environmental disaster.

However, these attacks have shortcomings. First, they require significant manual effort against individual PLCs, which does not scale well in an attack scenario where many different models of PLC could be targets (for example, against a power plant). And second, these attacks do not consider the specific logic running on each PLC, and cannot inflict damage more fine-grained than simply disabling the PLC.

Therefore, McLaughlin [43] posits a different attack on PLCs. He suggests the idea of a malware-based attack on a PLC's master terminal unit (MTU), or control system, which would map the functionality of an unknown PLC with no *a priori* knowledge, and subsequently construct a malicious payload for the PLC based on a preconfigured goal. This idea is conceptualized in Figure 13. We now examine a proof-of-concept variant of this attack.

[18] Synchronous condensers are also managed by PLCs. Due to their similarities to generators, we do not discuss them separately.

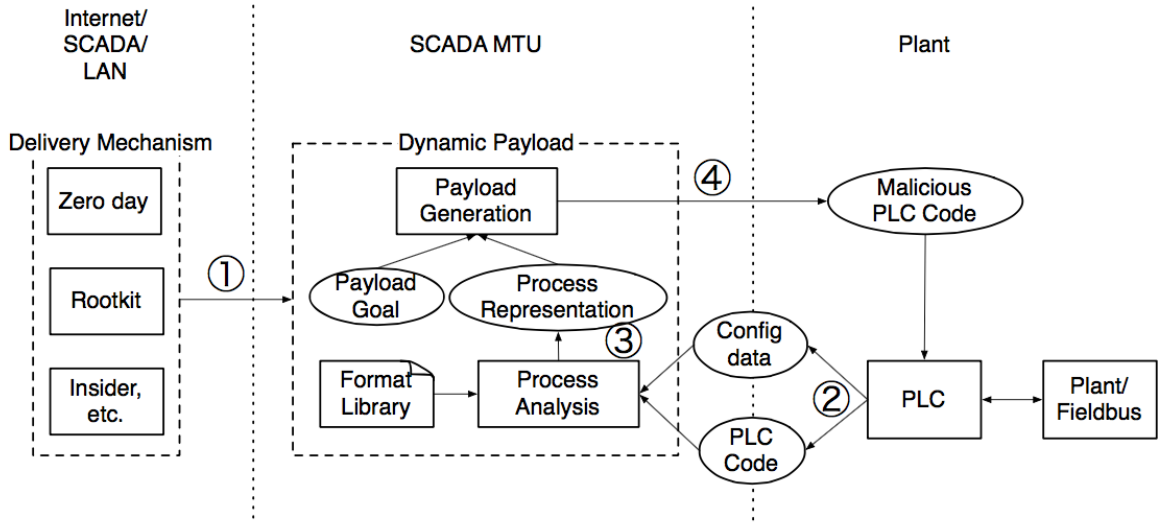

Figure 13: PLC malware illustration [43]

5.4.1 Attack Vectors

The initial vector of this malware-based attack is the MTU (part of the Control Systems attack surface), connected to one or more PLCs inside a generation plant. Control system computers can be infected by malware in a number of ways. As they are often behind network firewalls and sometimes even air-gapped, a promising method of attack is through a human operator. Stuxnet, a piece of malware which targeted PLCs behind an air-gapped control system computer, was able to reach its target through an operator's infected USB stick [34]. Commercial security researchers have also attempted remote intrusion into control system computers with success. Maynor and Graham [40] offer pen-testing anecdotes wherein they gain complete access to the control system computers in power plants through insecure company wireless networks.

Once the malware has infected the MTU, its secondary vectors of attack are the PLCs themselves; specifically, their control ladder logic (PLCs are considered part of the Physical Devices attack surface). PLCs by design accept logic updates from their trusted controllers [45], so this second vector is a feature rather than an exploit.

5.4.2 Methodology

McLaughlin and McDaniel [41] propose an attack on unknown PLCs in an industrial control system where the attacker's only *a priori* knowledge is how the system works and what equipment it is expected to have. We will define these requirements more concretely shortly. The major divergence from the attack conceptualized by McLaughlin [43] is thus that the attacker must be familiar with the behavior of the target system; this attack will *not* work blindly on any industrial control system. As the authors elaborate, "SABOT is not for adversaries that do not understand the behavior of the victim plant. In such cases, an adversary can erase the PLC's memory, upload random instructions, or attempt to bypass safety properties of the control logic."

Some familiarity with PLC operation is required to understand how the attack works [43]. PLC logic, the program it runs, is essentially a set of Boolean expressions evaluated in a continuous loop. In each iteration of the loop, or cycle, a set of input variables is read from the system's sensors and used to evaluate the Boolean expressions in order. These expressions may also be dependent on internal state and timer variables. At the end of each cycle, a set of output variables will have been produced, which are used to make decisions and issue commands to the system's actuators and monitoring apparati. Each of these variables is stored in a static memory location on the PLC.

We now examine the authors' three-part attack, known as SABOT:[19]

1. Behavior Encoding:

The attacker creates specifications containing all knowledge of the system's behavior. This includes all suspected devices in the system, as well as how they interact. For example [41], if the system has a button and a valve, and pushing the button opens the valve, this information will be fed into the specifications.

[19] SABOT: Specification-based Attacks against Boolean Operations and Timers

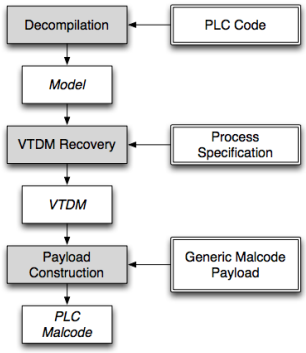

Figure 14: PLC attack illustration [41]

2. PLC Decompilation:

Once on the infected MTU, SABOT downloads the logic bytecode from the PLC. It symbolically executes the bytecode to obtain a control flow graph and modifies the constraints obtained from the CFG to fit a NuSMV model. Using this model and the encoded specifications from Step 1, a Variable To Device Mapping is determined, wherein the variables in the attacker's specifications are mapped to actual devices in the PLC in accordance with the rules found in the CFG and assumptions made in the specifications. *The attacker assumptions may not agree with the results of decompilation.*

3. Payload Construction:

If a fitting mapping is found, SABOT constructs a custom malicious payload for the PLC based on the goals the attacker also encoded in the specifications. For example, if the goal of the attack was to permanently seal the valve, then the payload will include Boolean logic to disregard button input. The payload is then uploaded to the victim PLC, completing the attack.

5.4.3 Results & Potential Outcomes

SABOT is not perfect. Attacker assumptions about the plant may be wrong, and the mapping stage may incorrectly map devices to variables. The specifications encoded by the attacker may also have multiple fitting mappings to the devices connected to the PLC, only one of which executes the desired attack. An incorrectly executed attack will likely be caught by alarms that detect improper operation, so this attack is not stealthy when it fails.

However, SABOT was able to successfully compile attacks against plants similar to power generators in the authors' testing at least some portion of the time. Figure 15 shows the results of SABOT against a simulated pH neutralization plant[20] and traffic light. All blank cells indicate successful attacks; cells with a "P" indicate a false positive; and gray cells indicate an attack not attempted. SABOT attacks this plant with a high degree of success.

A successful attack against a power generator, or against multiple generators simultaneously (as SABOT scales well, requiring low manual effort beyond initial reconnaissance), could be catastrophic. Malfunctioning PLCs could be caught by redundant control loops like state estimation and AGC before damaging any equipment; however, the plants would still need to be taken offline for repair, forcing short-term blackouts. In a worse scenario, the corrupted PLCs would succeed in causing damage to the generators, resulting in persistent blackouts while the generators are replaced. The Aurora Generator Test [3] demonstrated that an attacker with control of a generator's electronics (MTU and PLCs) can cause it to physically destroy itself.

5.4.4 Overview of Proposed Defenses

As PLCs by design accept logic ladder updates from their MTUs, and we are assuming an attack model where the MTU is compromised by malware, we must consider defenses with these limitations. McLaughlin and McDaniel [41] posit two promising defenses: first, the use of safety PLCs, additional PLCs in the system with the sole purpose of ensuring proper operation of subjugate PLCs; and second, obfuscated control logic in PLCs to prevent successful decompilation of downloaded PLC control logic.

6 Assessment

We now discuss the implications of the demonstrated attacks, in terms of grid security as well as research direction.

6.1 Summary of Attacks

Table 2 summarizes the attacks we have discussed. From this summary, we make the following observations:

1. Every surface exposed by the grid's EMS is vulnerable to attack.

2. The barriers for attack are not high for a capable attacker.

3. Most important control loops that govern the grid can be attacked and disrupted.

4. The outcomes of these attacks can entail high socioeconomic cost.

[20] A pH neutralization plant is equivalent in complexity to a power generation plant.

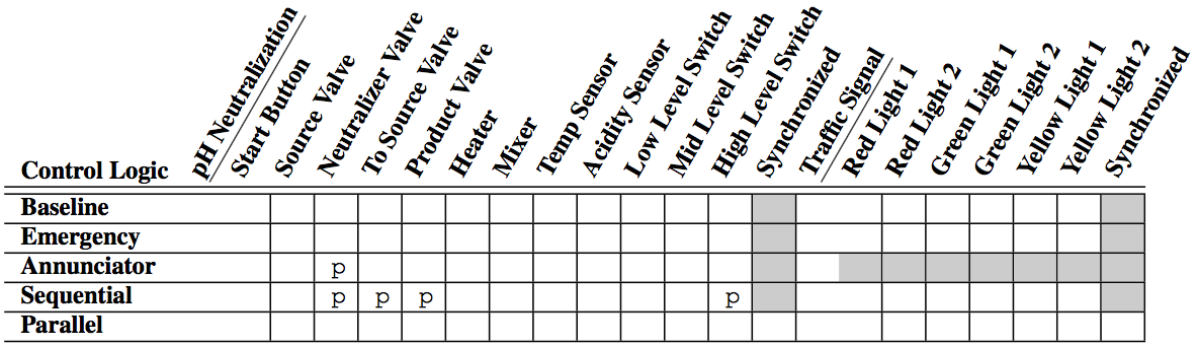

Figure 15: SABOT pH neutralization plant and traffic signal attack results [41]

Attack	Attack Surfaces	Control Loops Affected	Potential Outcomes
False Data Injection	Data Acquisition Channel Physical Devices	State Estimation Automatic Generation Control	Stress/disable power line Stress/disable generator Inefficient routing Load shedding Market fraud
AMI Load Alteration	Control Channel Data Acquisition Channel Physical Devices	Demand Response	Load shedding Disconnect generator
PMU GPS Spoofing	Physical Devices	Wide Area Monitoring	Stress/disable power line Stress/disable generator Load shedding
PLC Malware	Control Channel Control Systems Physical Devices	Automatic Voltage Regulation Governor Control	Stress/disable generator Violate generator safety

Table 2: Summary of observed attacks

6.2 Impact of These Findings

The trend in historical power grid failures is that seemingly insignificant events can lead to significant damage. The complex dependencies between its control loops, the scale of its operation, and the influence over it by uncontrollable factors like weather all imbue the grid with the property of escalation; small problems have the potential to intensify rapidly.

The 2003 Northeast blackout was triggered by a control system software bug and a single transmission line's failure. When the transmission line hit a tree branch and shorted (a common occurrence), the software bug in the control system software delayed to inform the operators of the lost line for an hour. During this time, the single line failure caused a cascading series of line overloads and failures, resulting in the loss of power to 55 million people.

The 2011 Southwest blackout was triggered by the accidental cutting of a single transmission line. Less than a minute after the line's failure, the resulting phase shift forced the rapid disconnection of all overloaded generators in the under-supplied area, leading to blackouts across Southern California, Arizona, and Mexico. Despite the rapid response of a fully-functional EMS, roughly 3 million people lost power.

These blackouts were caused by trivial, unintentional errors: a bug that induced deadlock in a single control center's monitoring services; a single failed transmission line. This highlights how fragile parts of the U.S. power grid are to mistake alone, much less intentional attack.

According to Miller and Rowe [46], the U.S. has yet to experience a significant, targeted electronic attack against the power grid.[21] But given the fragility of the system, the attacks we have examined in this paper would theoretically be able to induce blackouts similar to, or even worse than, these historical examples. A false data injection attack against the state estimator, if used to hide the state of a transmission line under duress, would produce a scenario equivalent to the 2003 blackout, as operators would be unable to detect the overloading and eventual failure of the targeted line. The attackers could even actively disable the line by deceiving operators into believing the line was under-utilized, resulting in a damaging amount of power being fed into the line. False data injection is not the only attack that could trigger these scenarios. For utilities already using wide area monitoring, GPS spoofing attacks against the PMUs monitoring a critical line could induce the same levels of deception.

If minor, unintentional mistakes can cause severe damage to the grid, it is probable that a large-scale, coordinated attack would be effective and catastrophic.

[21] There have been unintentional attacks against the grid. In 2003, the Davis-Besse Nuclear Power Plant was infected by the SQL Slammer worm and its alarm system was disabled for several hours.

6.3 Current Research Direction

Academia has largely been focused on false data injection attacks. Though a relatively young attack (introduced in 2009 by Liu *et al.* [37]), it has gained traction in the research community for a variety of reasons:

Generalization: False data injection attacks are equipment-independent. They are effective attacks regardless of what brand of sensor is deployed, what communication protocol is chosen, or what control system software suite is running; as long as there exists some attack vector through which the data can be injected, the control system's bad data detection models can be exploited. False data injection attacks are not even specific to the power grid; they can be tailored to any control system which uses a similar bad data detection scheme.

Difficulty: It is non-trivial to defend against a false data injection attack. By design, they evade the current bad data detection method used by utilities' control systems. Bad data detection methods more complex than using measurement residuals have been developed, but they are handicapped by their complexity. The grid requires expedited decision-making based on the results of state estimation; complex bad data detection schemes historically introduce more delay than is acceptable. An alternative method of defending against false data injection attacks is to secure the communication channels, physical devices, and control systems against compromise (perimeter security). Unfortunately, there is little economic incentive for utilities to replace working equipment and software in the grid, even though some of which is known to be insecure.

Longevity: The grid is increasingly using data collected from its sensors to make automated, impactful decisions. False data injection attacks will not only remain viable; their potential is in fact growing.

As such, the focus of formal academic research on false data injection attacks is objectively positive. But formal academic research is not the only input to the study of grid security; independent and commercial security researchers have also made significant contributions.

The primary focus of independent and commercial security research is control system and physical device security. Specifically, this community focuses on reverse engineering and penetration testing efforts against specific embedded devices and control system software suites used by the grid's EMS. One focus of this community has been the compromise of smart meters. Widely deployed, accessible to customers, and responsible for both billing and actual power provision, AMI is a likely target for criminals and thieves, making valuable a thorough investigation of its vulnerabilities. Another re-

search focus, mainly of the commercial pen-testing community, has been control system software. Due to the nature of commercial relationships, details about control system software vulnerabilities are often concealed or generalized; however, commercial research into control system software vulnerabilities is still valuable, as it provides insight into the vulnerabilities of systems which are typically (erroneously) considered out of an attacker's reach.

As we see, there currently exists a mild dichotomy between academic and independent/commercial security researchers in terms of studying power grid attacks. Academia tends to approach generalized attacks at the system level, while the independent/commercial community tends to focus on low-level attacks against individual components of the grid.

6.4 Future Work

We propose several directions for future work in the field.

The future of the power grid: The U.S. power grid is experiencing growing pains due to the advent of the electric vehicle industry and the advancement of renewable energy technology. The grid is struggling to account for electric vehicles, as they entail considerable load, do not draw power from a set location, and do not necessarily follow standard demand/time models. Likewise, the grid is having trouble handling the influx of power from renewable energy sources (including household solar power). Renewable power generation is typically erratic, depending on factors like wind speed or cloud positioning, and the modern grid is ill-equipped to handle erratic injections of power into its distribution networks by customers expected to be constant loads.

These issues are being combated by the proliferation of AMI and PMUs.[22] AMI will allow for the close monitoring of household power generation and the shedding of electric vehicle charging load using demand response, while PMUs' rapid measurements will allow the Wide Area Measurement system to automatically adjust to these fluctuating loads and generators. The days of state estimation, where human operators make decisions manually for the grid, are coming to an end.

A strong future direction of power grid security research would be studying attacks against the wide area monitoring and demand response control loops. We inherently expect poor perimeter security in smart meters and PMUs; as such, a better focus for research is bad data detection in both wide area monitoring and demand response. According to Choi and Xie [13], the state of the art in wide area monitoring bad data detection is at the same point as state estimation. And Jiang *et al.* [28] observe that demand response and AMI security is focused on energy theft rather than injurious attacks.

Security through obscurity:

In each of the attacks we study in this paper, we see the influence of the power grid's adherence to security through obscurity.

False data injection attacks are explicitly hindered by a lack of knowledge of the victim grid's topology (which is necessary to construct the DC power flow matrix H); it is a significant enough obstacle that several authors even modified the false data injection attack to bypass the need for complete topological knowledge at the cost of a shrunken attack surface. GPS spoofing attacks against PMUs will soon encounter these same issues as wide area monitoring is incorporated into grid decisions and bad data detection techniques are applied.

Malware-based attacks against PLC's are distinctly useful in situations where a target PLC and its precise interaction with the plant it controls are hidden from the attacker. The attack will fail without some *a priori* knowledge of the target system's operation, but is designed to account for secrecy of target design as much as possible.

AMI attacks against demand response are the one attack we study that does not seem to be hindered by a lack of knowledge of the grid; this is because the protocols by which control systems communicate to AMI are open sourced and standardized, and no additional topological knowledge is needed to attack AMI via demand response.

Therefore, the grid's practice of design and implementation secrecy clearly imposes some cost on attackers in terms of reconnaissance and potential for failure and/or detection.

An interesting direction for future research would be to attempt to quantify or measure the cost security through obscurity imposes on grid attackers, or the benefit gained by grid defenders through this practice. It is a compelling question, largely because of the conflicting interests involved: security engineers tend to subscribe to Kerckhoff's principle, wherein full system disclosure to the attacker (minus a "key") is assumed; industrial system engineers trend oppositely, favoring minimal design and implementation disclosure as a "defense in depth" strategy.[23]

Determining the actual value of practicing secrecy at the system level would be valuable, and difficult.

Modeling cyber-physical events:

In the Potential Outcomes of each attack we study, we see that there is leeway in interpreting the physical outcomes (overloaded power lines, stressed generators, etc.)

[22] The widespread installation of AMI and PMUs is an essential part of the "Smart Grid," according to Amin [5].

[23] Security through obscurity is often considered a part of defense in depth strategies. Unfortunately, it is sometimes the only level of security employed.

that attacks can cause. Few grid attack authors attempt to model the grid in such a way that physical results can be accurately linked to cyber-attacks.

This is understandable. Modeling the grid in such a way that cyber-events are linked to physical outcomes is a hard problem. A cyber-physical system as large as the power grid does not exist in a vacuum; climate and weather, for example, directly affect the grid's susceptibility to physical damage and its severity.

Modeling attacks on (or changes to) the grid based only on their effect on the power grid's routing efficiency, equipment tolerances, or load capacity is fallacious; outside factors affect both the operation of the grid and the results of its failure. A nuclear generator's failure has environmental effects in addition to generational loss. A power line's failure can cause fire damage disproportionate to its effect on the grid's routing efficiency. Factoring weather, market fluctuation, social conditions, and other outside variables into a model of the grid could be essential to ensure the accuracy of its predicted physical outcomes.

As such, another interesting direction for future research in power grid security, as well as cyber-physical system security as a whole, is to develop models to accurately predict the physical outcomes of cyber-attacks. Such models would allow for more informed discussion of the potential damages caused by attacks on industrial-scale cyber-physical systems, in terms of cost, property damage, and socioeconomic outcome. An issue in the security community is conveying the implications of security breaches to parties interested strictly in economic consequence; work in this area would have value, and would also entail a high degree of difficulty.

7 Conclusion

We have shown that the power grid, being a cyber-physical system, has exposed numerous attack surfaces to malicious electronic attack. The control loops which maintain the grid's functionality and safety are inherently vulnerable to disruption by these attacks, disruption which can potentially lead to damage to the grid and its generators, and denial of electricity to the grid's customers. Blackouts entail serious socioeconomic ramifications, and as malicious electronic attack on the grid can directly cause them, this area of research should continue to be investigated thoroughly.

8 Acknowledgments

I would like to thank my advisors, Kirill Levchenko, for providing valuable feedback during the initial research and writing processes, and Stefan Savage, for helping me select my topic and find relevant research. I would like to thank Geoff Voelker and Mihir Bellare for serving on my research exam committee. And I would especially like to thank Alex Snoeren for serving as my chair and for guiding me through the research exam process.

Acronyms

AGC: automatic generation control

AMI: advanced metering infrastructure

EMS: energy management system

HMI: human-machine interface

IED: intelligent electronic device

ISO: independent system operator

MTU: master terminal unit

PLC: programmable logic controller

PMU: phasor measurement unit

RTU: remote terminal unit

SCADA: supervisory control and data acquisition

VAR: volt-amperes reactive

References

[1] National security telecommunications advisory committee information assurance task force electric power risk assessment. NSTAC & IATF.

[2] Blackout 2003: How did it happen and why? In *Hearings Before the Committee on Energy and Commerce, House of Representatives, One Hundred Eighth Congress, First Session* (2003).

[3] Cyber war - the aurora project. *60 Minutes* (2009).

[4] Nist framework and roadmap for smart grid interoperability standards, release 1.0. In *NIST Special Publication 1108* (2010), NIST.

[5] AMIN, M. Urban energy & green design - smart grid cities. In *Proceedings of the 2010 Tufts Energy Conference on the Evolution of Energy* (2010), University of Minnesota.

[6] AMIN, M., AND STRINGER, J. The electric power grid: Today and tomorrow. In *MRS Bulletin: Carriers, Storage, & Transformation* (2008).

[7] ANONYMOUS. personal communication, 2015.

[8] ASHOK, A., HAHN, A., AND GOVINDARASU, M. Cyber-physical security of wide-area monitoring, protection and control in a smart grid environment. *Journal of advanced research 5*, 4 (2014), 481–489.

[9] BARRETO, C., CÁRDENAS, A. A., QUIJANO, N., AND MOJICA-NAVA, E. Cps: market analysis of attacks against demand response in the smart grid. In *Proceedings of the 30th Annual Computer Security Applications Conference* (2014), ACM, pp. 136–145.

[10] BERESFORD, D. Exploiting siemens simatic s7 plcs.

[11] BOBBA, R. B., ROGERS, K. M., WANG, Q., KHURANA, H., NAHRSTEDT, K., AND OVERBYE, T. J. Detecting false data injection attacks on dc state estimation.

[12] C4-SECURITY. The dark side of the smart grid - smart meters (in)security.

[13] CHOI, D.-H., AND XIE, L. Fully distributed bad data processing for wide area state estimation. In *Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on* (2011), IEEE, pp. 546–551.

[14] COSTIN, A., AND ZADDACH, J. Embedded devices security and firmware reverse engineering. In *BH13US Workshop* (2013).

[15] COSTIN, A., ZADDACH, J., FRANCILLON, A., BALZAROTTI, D., AND ANTIPOLIS, S. A large-scale analysis of the security of embedded firmwares. In *USENIX Security Symposium* (2014).

[16] CSANYI, E. Preparing to synchronize a generator to the grid. *Energy & Power, EEP* (2013).

[17] DAVIS, M. Smartgrid device security - adventures in a new medium.

[18] ERICSSON, G. Cyber security and power system communication 2014; essential parts of a smart grid infrastructure. *Power Delivery, IEEE Transactions on 25*, 3 (July 2010), 1501–1507.

[19] ESMALIFALAK, M., NGUYEN, H., ZHENG, R., AND HAN, Z. Stealth false data injection using independent component analysis in smart grid. In *Smart Grid Communications (SmartGridComm), 2011 IEEE International Conference on* (2011), IEEE, pp. 244–248.

[20] FADEL, E., GUNGOR, V., NASSEF, L., AKKARI, N., MAIK, M. A., ALMASRI, S., AND AKYILDIZ, I. F. A survey on wireless sensor networks for smart grid. *Computer Communications* (2015).

[21] FERC. 2014 assessment of demand response and advanced metering - staff report.

[22] FERC/NERC. Staff report on the september 8, 2011 blackout.

[23] GAO, W., AND MORRIS, T. H. On cyber attacks and signature based intrusion detection for modbus based industrial control systems. *Journal of Digital Forensics, Security and Law 9*, 1 (2014), 37–56.

[24] GRAHAM, S. Disrupted cities: Infrastructure disruptions as the achilles heel of urbanized societies. In *Disaster, Infrastructure, and Society: Learning from the 2011 Earthquake in Japan* (2012).

[25] HISKENS, I. A. Introduction to power grid operation. *Tutorial on Ancillary Services from Flexible Loads* (2013).

[26] ILLERA, A. G., AND VIDAL, J. V. Lights off! the darkness of the smart meters.

[27] JIA, L., THOMAS, R. J., AND TONG, L. Impacts of malicious data on real-time price of electricity market operations. In *2012 45th Hawaii International Conference on System Sciences* (2012), IEEE, pp. 1907–1914.

[28] JIANG, R., LU, R., WANG, Y., LUO, J., SHEN, C., AND SHEN, X. S. Energy-theft detection issues for advanced metering infrastructure in smart grid. *Tsinghua Science and Technology 19*, 2 (2014), 105–120.

[29] JIANG, X., ZHANG, J., HARDING, B. J., MAKELA, J. J., AND DOMINGUEZ-GARCIA, A. D. Spoofing gps receiver clock offset of phasor measurement units. *Power Systems, IEEE Transactions on 28*, 3 (2013), 3253–3262.

[30] KIM, J., AND TONG, L. On topology attack of a smart grid: Undetectable attacks and countermeasures. *Selected Areas in Communications, IEEE Journal on 31*, 7 (2013), 1294–1305.

[31] KOSUT, O., JIA, L., THOMAS, R. J., AND TONG, L. Malicious data attacks on smart grid state estimation: Attack strategies and countermeasures. In *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on* (2010), IEEE, pp. 220–225.

[32] KOSUT, O., JIA, L., THOMAS, R. J., AND TONG, L. On malicious data attacks on power system state estimation. In *Universities Power Engineering Conference (UPEC), 2010 45th International* (2010), IEEE, pp. 1–6.

[33] KOSUT, O., JIA, L., THOMAS, R. J., AND TONG, L. Malicious data attacks on the smart grid. *Smart Grid, IEEE Transactions on 2*, 4 (2011), 645–658.

[34] LANGNER, R. Stuxnet: Dissecting a cyberwarfare weapon. *Security & Privacy, IEEE 9*, 3 (2011), 49–51.

[35] LIN, J., YU, W., YANG, X., XU, G., AND ZHAO, W. On false data injection attacks against distributed energy routing in smart grid. In *Cyber-Physical Systems (ICCPS), 2012 IEEE/ACM Third International Conference on* (2012), IEEE, pp. 183–192.

[36] LIU, T., GU, Y., WANG, D., GUI, Y., AND GUAN, X. A novel method to detect bad data injection attack in smart grid. In *INFOCOM, 2013 Proceedings IEEE* (2013), IEEE, pp. 3423–3428.

[37] LIU, Y., NING, P., AND REITER, M. K. False data injection attacks against state estimation in electric power grids. *ACM Transactions on Information and System Security (TISSEC) 14*, 1 (2011), 13.

[38] MANANDHAR, K., CAO, X., HU, F., AND LIU, Y. Combating false data injection attacks in smart grid using kalman filter. In *Computing, Networking and Communications (ICNC), 2014 International Conference on* (2014), IEEE, pp. 16–20.

[39] MAYNARD, T. Business blackout. In *Lloyds Emerging Risk Report - 2015* (2015).

[40] MAYNOR, D., AND GRAHAM, R. Scada security and terrorism. In *Proceedings of the 2006 Blackhat Federal Conference* (2006), Errata Security.

[41] MCLAUGHLIN, S., AND MCDANIEL, P. Sabot: specification-based payload generation for programmable logic controllers. In *Proceedings of the 2012 ACM conference on Computer and communications security* (2012), ACM, pp. 439–449.

[42] MCLAUGHLIN, S., PODKUIKO, D., AND MCDANIEL, P. Energy theft in the advanced metering infrastructure. In *Critical Information Infrastructures Security* (2010), Springer, pp. 176–187.

[43] MCLAUGHLIN, S. E. On dynamic malware payloads aimed at programmable logic controllers.

[44] MILINKOVIC, S., AND LAZIC, L. Industrial plc security issues. In *Telecommunications Forum (TELFOR), 2012 20th* (2012), IEEE, pp. 1536–1539.

[45] MILINKOVIC, S., AND LAZIC, L. Industrial plc security issues. In *Telecommunications Forum (TELFOR), 2012 20th* (2012), IEEE, pp. 1536–1539.

[46] MILLER, B., AND ROWE, D. A survey scada of and critical infrastructure incidents. In *Proceedings of the 1st Annual conference on Research in information technology* (2012), ACM, pp. 51–56.

[47] MOHAGHEGHI, S., STOUPIS, J., AND WANG, Z. Communication protocols and networks for power systems-current status and future trends. In *Power Systems Conference and Exposition, 2009. PSCE'09. IEEE/PES* (2009), IEEE, pp. 1–9.

[48] MONTICELLI, A. Electric power system state estimation. *Proceedings of the IEEE 88*, 2 (2000), 262–282.

[49] NICHOLSON, A., WEBBER, S., DYER, S., PATEL, T., AND JANICKE, H. Scada security in the light of cyber-warfare. *Computers & Security 31*, 4 (2012), 418–436.

[50] NIGHSWANDER, T., LEDVINA, B., DIAMOND, J., BRUMLEY, R., AND BRUMLEY, D. Gps software attacks. In *Proceedings of the 2012 ACM conference on Computer and communications security* (2012), ACM, pp. 450–461.

[51] OZAY, M., ESNAOLA, I., VURAL, F., KULKARNI, S. R., AND POOR, H. V. Sparse attack construction and state estimation in the smart grid: Centralized and distributed models. *Selected Areas in Communications, IEEE Journal on 31*, 7 (2013), 1306–1318.

[52] POLLET, J. Electricity for free? the dirty underbelly of scada and smart meters.

[53] RAHMAN, M. A., AND MOHSENIAN-RAD, H. False data injection attacks with incomplete information against smart power grids. In *Global Communications Conference (GLOBECOM), 2012 IEEE* (2012), IEEE, pp. 3153–3158.

[54] REEVE, T. 4sics: What hackers do when they access a power grid honeypot. *SC Magazine UK* (2015).

[55] RICHTER, A., VAN DER LAAN, E., KETTER, W., AND VALOGIANNI, K. Transitioning from the traditional to the smart grid: Lessons learned from closed-loop supply chains. In *Smart Grid Technology, Economics and Policies (SG-TEP), 2012 International Conference on* (2012), IEEE, pp. 1–7.

[56] ROBERTS, M., AND NUNNELLEY, M. E. How digital governors boost operation of multiple needle impulse turbines, 2011.

[57] ROSATELLI, F. Maintaining equilibrium in the european power grid. *Introduction to POWER-GEN Europe 2014* (2014).

[58] ROUF, I., MUSTAFA, H., XU, M., XU, W., MILLER, R., AND GRUTESER, M. Neighborhood watch: security and privacy analysis of automatic meter reading systems. In *Proceedings of the 2012 ACM conference on Computer and communications security* (2012), ACM, pp. 462–473.

[59] SHEPARD, D. P., HUMPHREYS, T. E., AND FANSLER, A. A. Evaluation of the vulnerability of phasor measurement units to gps spoofing attacks. *International Journal of Critical Infrastructure Protection 5*, 3 (2012), 146–153.

[60] SIEMENS. Communications network solutions for smart grids.

[61] SMITH, H. L. A brief history of electric utility automation systems. In *Electric Energy Online* (2010).

[62] SRIDHAR, S., AND GOVINDARASU, M. Cyber–physical system security for the electric power grid. *Proceedings of the IEEE 100*, 1 (2012).

[63] SRIDHAR, S., AND MANIMARAN, G. Data integrity attacks and their impacts on scada control system. In *Power and Energy Society General Meeting, 2010 IEEE* (2010), IEEE, pp. 1–6.

[64] STOUFFER, K., FALCO, J., AND SCARFONE, K. Guide to industrial control systems (ics) security. In *Recommendations of the National Institute of Standards and Technology* (2011), NIST.

[65] TERZIJA, V., VALVERDE, G., CAI, D., REGULSKI, P., MADANI, V., FITCH, J., SKOK, S., BEGOVIC, M. M., AND PHADKE, A. Wide-area monitoring, protection, and control of future electric power networks. *Proceedings of the IEEE 99*, 1 (2011), 80–93.

[66] TIPPENHAUER, N. O., PÖPPER, C., RASMUSSEN, K. B., AND CAPKUN, S. On the requirements for successful gps spoofing attacks. In *Proceedings of the 18th ACM conference on Computer and communications security* (2011), ACM, pp. 75–86.

[67] USDOE. Advanced metering infrastructure and customer systems, 2015.

[68] WALKER, M. 'soft bomb' knocks out power plants. *The Guardian* (1999).

[69] WIKIBOOKS. Control systems/feedback loops — wikibooks, the free textbook project, 2015. [Online; accessed 24-November-2015].

[70] WIKIPEDIA. Attack surface — wikipedia, the free encyclopedia, 2015. [Online; accessed 24-November-2015].

[71] WIKIPEDIA. Automatic generation control — wikipedia, the free encyclopedia, 2015. [Online; accessed 20-November-2015].

[72] WIKIPEDIA. Cyber-physical system, 2015. [Online; accessed 22-Nov-2015].

[73] WIKIPEDIA. Demand response — wikipedia, the free encyclopedia, 2015. [Online; accessed 20-November-2015].

[74] WIKIPEDIA. Phasor measurement unit — wikipedia, the free encyclopedia, 2015. [Online; accessed 21-November-2015].

[75] WIKIPEDIA. Programmable logic controller — wikipedia, the free encyclopedia, 2015. [Online; accessed 20-November-2015].

[76] WIKIPEDIA. Static var compensator — wikipedia, the free encyclopedia, 2015. [Online; accessed 20-November-2015].

[77] WIKIPEDIA. Voltage regulator — wikipedia, the free encyclopedia, 2015. [Online; accessed 20-November-2015].

[78] XIE, L., MO, Y., AND SINOPOLI, B. Integrity data attacks in power market operations. *Smart Grid, IEEE Transactions on 2*, 4 (2011), 659–666.