

Creating Large Disturbances in the Power Grid: Methods of Attack After Cyber Infiltration

Senior Honors Thesis
June 2010

Loren Sands-Ramshaw
lsr@alum.dartmouth.org

Advisor: Sean W. Smith
sws@cs.dartmouth.edu

Abstract

Researchers are pursuing methods of securing the cyber aspect of the U.S. power grid, one of the country's most critical infrastructures. An attacker who is able to infiltrate an Energy Management System (EMS) can instruct elements of the grid to function improperly or can skew the state information received by the control programs or operators. In addition, a cyber attack can combine multiple attacks and affect many physical locations at once. A study of the possible adverse effects an attack could generate can underline the urgency of improving grid security, contribute to a roadmap and priority list for security researchers, and advise on how defending against cyber attacks can differ from defending against point failures and physical attacks. In this paper I discuss the physical and cyber systems that compose the power grid, and I explore ways in which a compromise of the cyber system can affect the physical system, with a particular emphasis on the best means of creating large disturbances. Further, I consider ways in which cyber attacks differ from physical attacks.

Contents

1	Introduction.....	5
2	Related Work	9
3	Physical Infrastructure	12
3.1	Generation	12
3.2	Transmission.....	13
3.3	Distribution	15
3.4	Circuit Breakers.....	16
3.5	Loads	18
3.6	Compensators	19
4	Cyber Infrastructure	21
5	Attack Scenarios.....	23
5.1	Direct Signaling	23
5.1.1	Generators	23
5.1.2	Circuit Breakers.....	25
5.1.3	Compensators	29
5.1.4	Transformers.....	30

5.1.5	Signaling Methods	31
5.2	Scope of Effect	34
5.2.1	Distance	34
5.2.2	Time	37
5.3	Human Operators	43
5.4	Intercompany Interactions.....	45
6	Conclusion	47
7	Acknowledgements	49
8	References	50

1 Introduction

The U.S. power grid is the national network of electricity on which we depend in order to function. It is the solution to the enormous problem of how to get power from six thousand power plants to three hundred million people [1, 2]. It includes power generators, such as nuclear plants or wind farms, transmission networks, which move power from the generators to where it is needed, and local distribution networks, which move power from the transmission networks to businesses and homes. Despite its scale and complexity, the power grid is extremely reliable and resilient, and remarkably so, considering that it is happening in real time – power is consumed less than a second after it is produced. While it rarely fails, its scale and speed means that when it does fail, the failure can be quick and widespread. Also, the resiliency currently built into the power grid was mainly designed to deal with non-malicious events, such as equipment malfunctions or quick changes in demand, as opposed to malicious events caused by people attacking the grid. Given our dependency on electricity, failures are a serious concern.

The power grid is operated in real time because we do not yet know how to efficiently store power in large quantities and because it needs to be able to respond to real-time failures. Continuously gathering and analyzing data from the grid as well as responding to failures and changes in demand requires a large amount of cyber infrastructure. The power grid is managed by control stations that are connected to substations, switches, and sensors in arrangements termed *SCADA* (supervisory control and data acquisition) systems. These systems are increasingly using open system architecture and becoming networked: for instance, using IP to communicate between the control center and remote equipment or among equipment in a substation. One example that demonstrates the dangers of such a move is a 2002 penetration test by a cyber security firm for a California power company that at the time had four million customers, in which the testers parked a van outside a substation, connected to the substation's open wireless local area network, and not only mapped the networked equipment in the control station but, "within 15 minutes, they mapped every piece of equipment in the operational control network. Within 20 minutes, they were talking with the business network and had pulled off several business reports" [3].

Also, instead of controllers depending only on sensors placed on power lines between the controllers and other stations, they have much larger areas of control, receiving data from many lines and unmanned substations. This has enabled *RTOs* (regional transmission organizations) to oversee areas that include many power companies, states, and control stations, increasing the reliability of the grid. These advanced SCADA systems allow power companies to address a number of challenges, including: 1) predicting and quickly reacting to changes in demand on a large scale; 2) keeping waste to a minimum by holding fewer generation plants online and less spinning electric reserve (unused capacity provided by generators that are synchronized to the grid); 3) maintaining this stability with the increasing prevalence of green energy sources such as wind and solar that have intermittent supply; 4) dealing with generation in the distribution grid, for instance, homes with solar panels transferring excess power to the grid; and 5) interacting with new, “smart” home appliances in order to help reduce peaks in demand.

The new smart technologies that are currently being integrated and will continue to be integrated into the U.S. power grid will introduce more vulnerability to attack. At a basic level, more sensors and controls will be networked together,

and there will be more connections between those networks and the power company's enterprise network, which in turn may be connected (although when connected, they are usually behind a firewall) to the Internet. The fact that some conceptualizations of the smart grid contain more networked devices than are on the Internet does not bode well for security in the smart grid, given that many Internet security problems are as yet unsolved.

One challenge we face is examining what a successful attack can do to the power grid and how. This study's importance is twofold: 1) to underline the importance of security for the power grid; and 2) to help security researchers determine and prioritize which parts of the system to protect.

2 Related Work

While there are many assessments of power-grid cyber security and studies of cyber security improvements, there are none that go into depth on what can be done after some level of infiltration has been achieved. The largest published attack-scenario list appears in a study from Idaho National Laboratory, one of the main government labs researching power grid cyber security, and is simply a short, very general list. The first two items are, for instance, “Take direct control of devices in substations and/or generation plants, shutting these facilities down,” and, “Plant malicious code or a ‘logic bomb’ that executes on a given event or at a preselected time to disrupt the system” [4]. The authors do not elaborate on how the scenarios can be carried out and what the range of effects might be. Such brief attention to the subject is commonplace. Another paper on protecting power systems against electronic intrusions has a similar list, containing, for instance, the following scenarios: “Shut down the substation or any portion of the subsystem controlled by the compromised device, either immediately or in a delayed manner,” and “Gather control and protection settings information that could be used in a subsequent attack” [5].

More often than not in power grid security papers, the topic either is not examined at all or is given a sentence or two in the introduction in order to convey importance. One paper proposing a certain cyber-security-assessment approach simply lists the “affected components” as, “Electronic devices, IED’s [Intelligent Electronic Devices], Controllers or SCADA system” and “Data altered or destroyed, devices reset, communication blocked or re-routed” [6]. Another paper on cyber-security assessment contains the general warning, “Compromised cybersecurity of a SCADA system can cause serious impact to a power system if the attack is able to launch disruptive switching actions leading to a loss of load. This is particularly troublesome if the attack can penetrate the control center network that is connected to substations under the SCADA system” [7]. Instead, papers focus on other areas, such as methods of cyber infiltration, methods of cyber defense, and ways in which cyber security for the power grid differs from Internet security.

This trend continues with newer publications that deal with the smart grid. The Department of Energy’s 2009 smart-grid cyber-security study devoted these two sentences to the consequences of a successful attack: “Many compromised Smart Meters or data collector nodes could be programmed by the attacker to

simultaneously send messages that cause power demand to be reduced dramatically and then to be increased dramatically. These phony messages could cause grid instability and power outages” [8]. While it does give an actual example of something that an attacker could manipulate, it does not describe possible variations of the attack, the possible scopes of disturbance, or other attack scenarios. This paper presents a more comprehensive, in-depth study of the topic. Some of the ideas I present do not appear in other publications, and I analyze the possibilities from a number of different angles, including possible effects, locations of attack, and interaction in concert with other ideas, with a focus on creating large, cascading outages.

3 Physical Infrastructure

Electricity is produced by generators, which are connected to transmission lines.

The electricity flows from the generators over the transmission lines to where the electricity is needed. The network of power lines that fans out, carrying electricity from the end of the transmission line to the many different customers in the area, is called the distribution grid.

3.1 Generation

The great majority of generators connected to the power grid are three-phase, AC (alternating current) generators. The generators basically consist of a magnet inside a hollow cylinder called a *stator*. Inside the surface of the stator there are three coils of wire. When the magnet rotates inside the stator, it creates a voltage that is supplied to customers, who then apply loads, which create currents in the coils of wire. Current is a measure of electrons moving down a wire, and electron movement is the electricity that powers electrical equipment such as lights and computers. Since the magnet is constantly rotating, the current in the wire is constantly alternating between electrons flowing first one way and then the other along the wire (hence the term ‘alternating current generators’). Further, each of

the three coils of wire wind through different parts of the stator. Thus at any given point in time the north end of the magnet is pointing in a different direction relative to each coil. Since the orientation of the magnet relative to the wire determines how the current changes in the coil, each wire's coil's current is alternating at different times. Each of the three coils and its associated current graph is termed a phase, which is why the generators are called three-phase generators.

At the basic level, the rotating magnet makes electrons move back and forth across a wire, and those electrons power electrical equipment. The magnet is moved by the generator's power source. For instance, in a wind generator, the wind turns the wind tower's blades, which turn the magnet. In a hydropower plant, water turns the magnet, and in coal and nuclear power plants, water is heated to form steam, which turns the magnet.

3.2 Transmission

Two important measurements other than current are voltage and resistance.

Voltage is a measure of the electrical force that moves the current along the wire, and the resistance of an object is a measure of the object's opposition to current

flowing through it. All objects have resistance, including the wires carrying the electricity. The amount of power consumed by current passing through an object is determined by the amount of current and the object's resistance. Take for example a battery and a light bulb, which are connected by a wire. Current flows from the battery, through the wire, and then through the light bulb. Since the goal is to produce light, we want most of the power to be spent by the light bulb, so we pick a wire that has low resistance. Since the wire has low resistance, only a small amount of power is consumed when the current travels over the wire, and then the light bulb, which has high resistance, consumes the remainder.

In the power grid, the customers are often far away from the generators, so the power must be transmitted over power lines (which are simply thick wires) from the power plants to the customers. In order for this process to be efficient, we want only a small amount of power to be consumed by this transmission. We achieve this in part by choosing a type of wire that has low resistance (usually copper). However, recall that the amount of power used depends on both the amount of resistance and the amount of current. We could further decrease the power consumed by the power line by decreasing the current. One important equation governing the behavior of electric power is $P=VI$, where P is power, V

is voltage, and I is current. Given a certain amount of power coming from a generator, in order to lower the current, we must raise the voltage.

This is precisely what is done in the power grid. Pieces of equipment called *transformers* are placed between a power plant and a transmission line, and they *step up* the voltage to a higher level, thereby decreasing the current. *Substations* are groupings of equipment based around transformers that have one or more power lines that come in at a certain voltage level, go through transformers, and go out on one or more lines at a different voltage level. For example, two power plants could have short lines going to a substation, which connects the lines and steps up the voltage for power to leave on a single long-distance, high-voltage transmission line. The other end of this long-distance line could, for instance, be at a substation that *stepped down* (lowered) the voltage to go out on two or more lower-voltage shorter-distance transmission lines.

3.3 Distribution

Transmission lines run to distribution substations, which step down the voltage from one or more lines so that the power can leave on a number of low-voltage, short-distance power lines to serve a certain group of customers, for instance, a

town. There will be further, smaller transformers located on power line poles that lower the voltage even more. For example, there may be a transformer on the pole in a house's back yard that has a line going to the house and a line going to its neighbor and that lowers the voltage to a level that the house's appliances can handle.

3.4 Circuit Breakers

Circuit breakers are devices that break the connection between two power lines, thereby stopping the flow of electricity between the lines. For example, at a substation A that had a high-voltage transmission line coming in and a low-voltage distribution line going out, the transmission line would enter a transformer that stepped down the voltage, and then a short line would run between the transformer and a circuit breaker, and then the distribution line would run from the circuit breaker out of the substation. If the circuit breaker was triggered – for instance, automatically by a nearby relay or the control station due to the current in the line being too high or manually in order to perform line maintenance – the short line would be disconnected from the distribution line, and power would not be able to flow from the transmission line to the distribution line, so the power would have to find some other way to flow.

If substation A's incoming transmission line originated in a substation B that had a second outgoing transmission line, all the power that would normally flow through the first transmission line would now have to flow through the second line.

This large power flow change is the mechanism that enables *cascading failures*.

Each power line has a certain maximum current capacity. When a line carries too much current, as the second transmission line may carry at the end of the above scenario, a circuit breaker on the line is automatically triggered, thus causing the power to flow on different line(s), some of which may consequently be overburdened, and so on. I will discuss cascading failures further later on in the paper.

Circuit breakers are commonly located between generators and substations, at substations on both ends of transmission lines, and on various parts of distribution lines. They can be controlled locally by a *relay* or remotely by a control system. A relay is a device on a power line that takes a measurement, for instance, the current level, and in the case of a bad state sends one or more messages. For instance, if the current was too high the relay could send a

command to a nearby circuit breaker to *open* (disconnect the wires) and a message to the control station to notify it of the opened line. Relays and circuit breakers are an important protective measure in the power grid because high currents can permanently damage power equipment such as transformers and generators.

A *fuse* is like a circuit breaker in that it disconnects two lines, but unlike a circuit breaker, it cannot automatically *close* (reconnect the lines together). A fuse is meant to protect equipment from higher currents than relay-controlled circuit breakers are meant to protect it from. It is simply a small section of the wire that is made of a different material – one that burns out when a high enough current passes through it, burns out, thereby creating a break in the line and interrupting the flow of power.

3.5 Loads

In AC power systems, not only the current, but also the voltage alternates back and forth, and they together produce two different types of power: real and reactive. *Real power* is produced when the current and voltage alternate at the same time, *reactive power* is produced when they alternate at completely

opposite times, and different amounts of both are produced when the current and voltage alternate at somewhat-overlapping times. On the side of the power system opposite generation, there are three types of devices, or loads, that use electricity. *Resistive* loads, such as light bulbs or electric ovens, run electricity through a high-resistance wire, which produces light and heat and consumes real power. *Inductive loads*, such as fans and vacuum cleaners, contain motors that generate magnetic fields and consume both real and reactive power. *Capacitive loads*, such as capacitors in the power supply of personal computers, are commonly said to produce reactive power.

3.6 Compensators

We use more inductive load than we do capacitive load, so the power grid has to compensate for consumers using more reactive power than the consumers produce. The two main types of compensators are *capacitor banks*, which produce reactive power, and *reactors*, which consume reactive power. Given the disparity in type of consumer loads, capacitor banks are more common. Both types of compensators are placed close to where they are needed, in the distribution network. Compensators can be controlled locally by a relay – for instance one on a distribution line that measures the reactive power level – and/or remotely

through a *SCADA* system. Reactive power is said to ‘support’ voltage:

insufficient reactive power leads to *voltage collapse* – the loss of voltage in a large

part of the system, which causes an outage.

4 Cyber Infrastructure

A SCADA system, as the name states, involves both centralized control and data gathering. The main controls include a surprisingly small set of possible actions: changing generator output levels, opening and closing circuit breakers, changing transformer *tap changers* (devices in transformers or neighboring voltage regulators that alter the voltage level), and turning compensators on and off. The data include generation power output as well as a number of readings along power lines, such as real and reactive power levels, current levels, and voltage levels. *RTU* (Remote Terminal Unit) is a broad term that encompasses all devices that feed data to the SCADA system and can alter some part of the grid: for instance, a modern relay that both sends updates to the control center and can close and open a circuit breaker. RTUs can send their data to the control center in a number of ways, such as via fiber optics or microwave radio.

An *EMS* (Energy Management System) encompasses a SCADA system as well as a number of different types of computer programs running on real-time data: automatic generation control, which changes generator output based on customer demand and the prices of buying power from or selling power to neighboring

companies; state estimation, which gathers data from all the RTUs to form a wide-area model of the state of the grid; contingency analysis, which continuously simulates what would happen if, for instance, an arbitrary line were to fail; automatic emergency load shedding, which cuts off load in order to preserve grid stability; and forecasting, which uses past load data and current weather data in order to predict future load.

A control station that contains an EMS can cover as small an area as a city or county or as large an area as multiple states. A number of station computers usually perform the EMS tasks on one or two private local area networks (the second network put in place for redundancy), which in many cases are connected to the corporate network via a firewall. In addition to a second private network, there can also be backup computers ready to perform the necessary functions in case the main computers fail. Further, especially for control stations that serve large areas, there can be entire backup control stations that have duplicate EMS systems and that receive duplicated messages from each RTU.

5 Attack Scenarios

5.1 Direct Signaling

When one typically imagines an attack on the power grid, what typically comes to mind is blowing up generators or substations, or perhaps the easier task of interrupting the transmission at some point along the miles of unguarded power lines (by, for example, cutting the line or creating a short circuit), and perhaps performing a number of these actions at the same time in a number of key places (although such a task is much easier to achieve with a cyber attack) in order to cause a large, long-lasting outage. The first two examples of direct signaling, sending commands to generators and circuit breakers, are the cyber parallels to the above physical attacks, albeit usually not as long-lasting.

5.1.1 Generators

An EMS has to constantly keep a balance between the amount of power generated and the demand for power. Customers do not notify their power company in advance how much electricity they will be using and when. While future load amounts can be approximated to some degree, they can never be

Action	Mechanism	Effect	Typical Reason
Raise or lower generator output	Varied	Raises or lowers amount of power entering the system	Match increasing or decreasing load (which is indicated by decreasing or increasing voltage readings)
Open or close circuit breaker	Physically separates or connects two sections of a line	Prevents or enables power from flowing through a line	Begin or end scheduled maintenance Drop or add load Isolate disturbance (open only; could involve a number of different readings)
Switch compensator on or off	Varied	Alters amount of reactive load or generation	Maintain reactive generation:load balance (based on reactive power measurements) Voltage support
Alter tap changer in transformer or voltage regulator	Sets new transforming ratio	Changes voltage level of power leaving the transformer	Maintain desired voltage range (when voltage readings are approaching normal operating limits)

Figure 1: The main available actions in a power system.

predicted perfectly. An EMS must continuously calculate the load on the system using data gathered at different points on the power lines. When there is more demand than power generated, the voltage sags (dips below the desired level) and there is a time window in which more power must be generated, or else a voltage collapse occurs. If attackers are able to make one or more generators either lower their output or shut down, they can create such an imbalance and perhaps a collapse.

There are a number of ways in which attackers could trigger changes in generation. They could infiltrate the generator's local electronic control system in

order to send the “turn off” or “lower output” directives. Also, most generation control systems accept remote commands (this is what enables the automatic generation control that modern EMSs employ). If attackers could access the data link between an EMS and a generator or compromise the EMS itself, they may be able to send the generator commands.

5.1.2 Circuit Breakers

Each EMS controls many circuit breakers throughout the section of the grid under its purview. Opening a breaker on a line that carries power away from a generator (the top-left circuit breaker in Figure 2) removes that supply of power from the grid, just as shutting down the generator does. However, there are two new situations that opening breakers can create.

The first situation is in a section of a distribution grid that is radially oriented, with a single substation supplying an area through lines that fan out from the substation, splitting at various points in order to reach each customer. When a breaker is opened in this situation, the breaker’s line loses power, as well as the lines that split away from that line and the customers on those lines. For




Figure 2: Example power grid diagram.

instance, consider the town in Figure 2, but without Line B or the Capacitor Bank. The town's power is supplied from a single substation (Substation C) that has five lines – one transmission line (Line A) that carries power to the substation and four distribution lines that carry power out of the substation and fan out to cover equally sized sections of the town. Opening a breaker on one of the four main distribution lines at a point close to the substation would cause a quarter of the town to be blacked out. This situation can be generalized to apply to any case in which the sole power source for a certain area is cut off. In the

example, the only power source to the one substation is Line A, so opening a breaker on that line would cut power off from the whole town. While almost every individual customer is served by a single load, large enough groups of customers, such as cities, have multiple transmission lines feeding their substations.

The second situation is when the breaker opens a line that is not the sole supplier of power to anyone. Unlike the first case, customers do not immediately lose power, so the system load is unchanged. Thus power must be rerouted to reach the customers. However, in the majority of the grid, power cannot be intelligently routed or even at all actively routed. Given, for instance, a simple substation connected to a generator and two outgoing transmission lines (Substation A in Figure 2), the substation cannot alter how much power goes out on one line versus the other. The two transmission lines and the supply line coming from the generator are connected to each other via a bus, which is simply a thick rod of copper or aluminum. The distribution of power on the outgoing lines is determined by aspects of physical rules, such as the resistance of the wire and the load on the other end. There are new devices called FACTS controllers that

would enable some degree of active routing, but they are not in widespread use [9].

For an example of the second situation, consider a substation that supplies a town and has two incoming transmission lines (Substation C). Line A can carry a maximum 1000 MW (megawatt, equal to one thousand watts; a watt is a unit of power), and Line B can carry a maximum of 500 MW. The town at this point in time has a demand of 600 MW, and each line is supplying 300 MW. The attackers obtain the model of this section of the system and intelligently decide to open a breaker on Line A. This causes the full 600 MW needed by the town to attempt to flow through Line B, which is only rated to carry 500 MW. If the 600 MW is permitted to flow through line B for a certain amount of time (which varies based on the type of cable and its safety margin), the line will be damaged. One possible scenario is that a power level higher than a line's rating would not be permitted to flow through the line; a relay along the line would sense the power level jump and would immediately open a circuit breaker, preventing damage to the line. Another possible scenario is that the automatic line opening would be delayed, giving time for a computer program or human operator to selectively shed load. In this case, 100 MW or more of load in the town would be

shed by opening one or more distribution lines, disconnecting a portion of customers in order to keep line B providing power to the remainder of the town.

This situation can also create a large cascading outage. The last example ends with the possibility of just the town out of power, but in the case of major lines in the middle of the transmission system, opening one line can first cause neighboring lines to overload, but then the power surges in another direction to try to get where it was going in a more roundabout way, but overloads those lines as well, etc.

5.1.3 Compensators

Capacitor banks and reactors can be switched off and on remotely and are automatically controlled by EMSs. Both also can be used maliciously to cause instability. The most common problem that compensators address is a dearth of reactive power. Switching off capacitor banks and switching on reactors can sharply diminish the supply of reactive power in the distribution network. Firstly, the lack of reactive power leads to voltage collapse. Secondly, reactive power must then be brought to where the demand is from generators via transmission lines, instead of it being locally produced by capacitors. This extra reactive power

traveling over the transmission lines contributes to the line's total power, which may then approach or exceed the line's power rating.

5.1.4 Transformers

Transformer tap changers can be remotely set to a different current-to-voltage ratio so that the outgoing voltage is outside of the bounds of normal levels. For instance, if the change is made in a location for which there are no other incoming power sources between it and loads, as in the first circuit breaker situation discussed in section 5.1.2 above, then a lowering of the voltage would produce a voltage sag everywhere down the line (or lines if the line splits).

Voltage sags can adversely affect electrical devices, including causing them to reset. The more sensitive types of devices, such as computers and process control machinery, can be affected – for instance can be reset – when the voltage sags as little as ten percent. Since such devices are the backbone of the nation's critical infrastructures, voltage sags are a serious concern. Of even more concern would be a power outage, and a large enough voltage sag could lead to a voltage collapse.

The situation changes if the alteration is made in a location that is not along the line of a sole power supply to a group of customers. If a single change is made, it is likely that at some point down the line when the power is combined with another incoming source that the voltage dip will be absorbed to a large degree. However, if most or all voltages are lowered on lines entering a certain area, then the instability could persist. Further, if the voltages can be lowered on a group of long-distance, high-powered transmission lines, then a much larger region could experience instability and even a cascading outage.

5.1.5 Signaling Methods

The four parts of the power grid discussed above are the main items that an EMS controls. There are various points in the system from which each of these parts can be given commands to change the state. Gaining access to a substation's communications network would allow attackers to send messages directly to the substation's transformers, circuit breakers, and compensators. Another point of entry is the line of communication between the equipment and the control station. If attackers can gain access to a line of communication, they can inject false commands to the equipment. For some companies – in particular, those that have one control center overseeing a large area – equipment

communication lines are consolidated by area before being sent (for example, via a dedicated T1 line) to the control center. These communication hubs can also be compromised.

The last point of entry is the control center. If attackers were able to access the control center's communications hub, they could send the commands directly out on the wire. If not, they might be able to alter the behavior of the manager program – for instance, through code injection or configuration file modification – which extrapolates the current grid power flow model from the incoming data and sends out automatic corrections, so that the program sends the attackers' desired commands via the communications hub. Also, instead of altering the behavior of the manager directly, they could alter values in the database the program uses to form its model of the grid's current state. For instance, artificially raising the datum corresponding to the measured voltage on a line may incite the manager to automatically correct the perceived problem by signaling the transformer to lower the voltage.

This same effect can be reached through modifying the data that RTUs send to the control station rather than the modifying the data after it reaches the control

station. For instance, in Figure 2 the attackers could compromise the power plant and substations and change the data that is sent from those locations to the control center. This method may be preferred for ease of entry – the substations may be easier to electronically break into than the control center – or for persistence – compromising the substations may be harder to detect and fix than compromising the control center. Such attacks were formerly considered to be largely or even fully impeded by algorithms used by state estimators to account for erroneous sensor measurements. However, a recent paper by Liu et al. [10] reveals how to create attacks that “successfully introduce arbitrary errors into certain state variables while bypassing existing techniques for bad measurement detection.” While it may be difficult to obtain the degree of system configuration knowledge necessary for the proposed attacks without having compromised the control center, such information could be acquired through means other than cyber attack. One non-malicious example of an EMS performing a detrimental action based on bad data can be found in NERC’s (North American Electric Reliability Corporation) archive of system disturbances. NERC publishes a yearly summary of disturbances that power companies are required to report to it, and the 2007 report describes an incident where an EMS disconnected 98,700 customers, saying that it was “triggered by the status change of several dynamic

transfer signals to telemetry error which biased the calculated load higher than the target load set in the load shed application” [11]

5.2 Scope of Effect

5.2.1 Distance

Out of the different parts of the power grid, attacking the transmission section of the grid would create the largest effect. Many of the original long, high-voltage transmission lines that connect different regions of the country were put in place for safety reasons, to help with emergencies. They were not meant for the continuous high usage that they experience today, with many regions getting all or most of their electricity from remote generators. For instance, a large portion of the electricity used by the Northeastern United States comes in around the clock on long transmissions lines from hydroelectric plants.

Since many of the effects of physical attacks can also be created by cyber attacks, methods of determining the best places for physical attack, such as the method put forth by Salmeron et al. [12] [13], can also be used to plan cyber attacks.

There have been a few studies that conclude that transmission is the most vulnerable section of the grid. Albert et al. [14] create and analyze a model of the current U.S. grid from a graph theory perspective and find that “disturbances

affecting key transmission substations greatly reduce [the power grid's] ability to function.” The insufficiency in transmission capability has long been recognized, and in 2005 Congress authorized the Department of Energy to create National Interest Electric Transmission Corridors – regions with high transmission congestion in which the Federal Energy Regulatory Commission has the power to make compulsory purchases of land in order to construct new transmission lines. While two National Corridors were created in 2007 – one in the Mid-Atlantic and one in southern California and western Arizona – routes for new lines have not yet been chosen. The Department of Energy’s 2009 National Electric Transmission Congestion Study [15] says of the Mid-Atlantic corridor, “little new transmission has been built in the region in the past three years.” Not only has the situation not improved, but the study also states that the region “has added new generation since 2006,” exacerbating the congestion problem.

An excellent example of a large blackout caused by a cascade over transmission lines is the Northeast blackout of 2003, which put the majority of eight states as well as Ontario, Canada out of power. It began with the unplanned shutdown of a power plant in Ohio. Over the next two hours, the circuit breakers on a few high-voltage transmission lines opened, mostly due an unusually high current

causing the lines to sag and come into contact with trees. This contact created a *short circuit*, which caused a very high current in the power line, which was detected by a relay, which automatically opens a breaker on the line. After approximately two hours, the overcurrent and undervoltage conditions were so severe that an entire swath of transmission lines quickly tripped out. This separated a region of power plants to the west from eastern demand, causing the power plants to go offline. The eastern region then experienced a surge of power from the far east, further tripping transmission lines and taking far east generators offline.

Current EMSs do not adequately prevent maliciously caused cascading outages. This inadequacy is due in part to their *N-1 contingency analysis* systems. An N-1 contingency analysis program continuously runs scenarios in which one piece of the grid goes down – for instance, a transmission line or a generator – to make sure that the grid is able to return to a stable state. A system designed to respond to single-point failures is insufficient, however, since attacks can easily create more than one contingency at once. One physical phenomenon that similarly creates multiple failures at once is a storm that causes geomagnetic field disturbances. One example is the 1989 collapse of the Quebec hydro grid, which

experienced seven contingencies in 57 seconds [16]. While N-2 or greater systems do exist, they are rarely used. Instead, sometimes if-then statements are added to the program, such as, “if transmission line X fails, generator Y usually does as well,” based on the prior experience of the operators, but these small additions to the program fail to address many of the multiple contingencies attackers could create [17].

5.2.2 Time

The magnitude of an event can be measured in time as well as distance. In the first circuit breaker example of section 5.1.2, opening one distribution line leaves a fourth of the town without power. However, it would be a relatively small outage if the breaker was automatically reclosed a second later. Two possible methods of prolonging a malicious change made to the system are to continue re-sending the original message, and to prevent further commands from being sent to the altered device, whether it be by crashing the control program or blocking the communications line.

When a generator is shut down, its downtime is longer than that of a line. A line in many cases can simply be switched on, whereas a generator must be started,

which can take a few hours for coal power plants or an entire day for nuclear plants. Also, many power plants need power in order to start. If the disturbance affects a wide area, including the power lines connected to the shut-down generators, then the generators must be *black started* in order to bootstrap the power grid back into working order. The process begins with those generators that have diesel generators that can be connected to the plant for this very purpose, and then incrementally with the other generators, using the power generated by the first set, all the while connecting the exact amount of load to the system in order to balance out the amount of power being generated. When the area is large and lies across company boundaries, the process of coordinating the startup can be an enormous task. Restoring power following the 2003 Northeast blackout took over forty-eight hours.

The delicacy of this bootstrapping process suggests a method of prolonging an outage even further. If the control center can be compromised and remain so (the attacker could, for instance, implant intelligent code that could continue to attack even after the attacker's connection to the control center has been cut), then the bootstrapping process can be sabotaged more easily due to the simpler model of the emerging system. At the beginning, for instance, it is only one

black-started generator supplying a small load. Taking out, say, the one transmission line between that generator and its load returns the grid to a complete blackout.

One advantage that a cyber attack has over a physical attack is its ability to continue and hide past the point a physical attack could. The physical attackers of a control station or substation cannot remain and wait to perform the same attack – police officers or members of the armed forces will come to investigate and guard the area. A cyber attack may not be recognized as an attack immediately – it could seem to be a software error – and may take a substantial effort to stop. Breaking the line of cyber entry, for instance, would not prevent malicious programs from being left behind, scanning the machine for malicious code may not find it, and wiping and reinstalling the machine performing power flow analysis, for instance, would not prevent malicious programs on other machines on the local network from spreading back to the power flow machine. Large power companies often have a redundant backup control center, and they may switch over to that facility in the event of a cyber attack. However, if the cyber attack was distributed – for example, via a number of substations – this

control center switch would not help, and it may take quite a while to identify the compromised substations and vet their equipment.

Another way in which time plays into the size of the event is the state of the power grid at the instant in which the attack is carried out. An attack can be more effective – more likely to cause a cascading outage, for instance – if the system is under stress to begin with. The peak load times are always during extreme hot or cold weather due to the heightened use of air conditioning and electric heating, respectively. Of the two, extreme heat creates the greater load, because the common methods of heating buildings depend less on electricity.

Thus attackers can simply look at the weather forecast to choose a time that will increase their chance of success.

An even more accurate method of determining the current system load is by looking it up online. RTOs that manage the deregulated open-market power system must provide availability and prices to the companies in the area. Of those that provide a web interface, most require registration; however, at least one major wholesale power provider does not require registration. CAISO, the California Independent System Operator, has a real-time, publicly-available online

system named OASIS (Open Access Same-time Information System), which provides real-time transmission loads [18].

Weather can be used not only to estimate the system load, but also to predict when there may be natural damage to the system that could complement the attacker's efforts: lightning striking power lines can create voltage surges; lightning or wind can cause trees or branches to knock lines and line poles over and create short circuits; and wind blowing lines against nearby trees can cause short circuits. Weather is by far the largest cause of damage to the grid, and large storms can cause a lot of damage. Both looking up regional loads, in cases when it is accessible, and paying attention to weather are helpful to choosing the timing in both physical and cyber attacks. However, the most accurate view of the state of the grid is obtained by looking at the real-time data gathered by the target region's EMS, which is only possible in a cyber attack.

The most long-lasting effect, however, is attained through permanently damaging equipment. While damaging equipment is much more likely to occur in a physical attack, it is still possible in a cyber attack. There is the commonly-cited possibility of taking control of a generator's control system and, say, cranking up

generation above the capacity of the machinery, raising heat and pressure to the blowing point. An attacker could also maintain a short circuit through a combination of cyber and physical means: if a short circuit is created, either by an attacker or a storm, circuit breakers on the line would usually trip. However, an increasing percentage of breakers can be remotely controlled. If an attacker could send commands to breakers to stay closed while the line is, for instance, in contact with the ground, then the extremely high current that comes with a short circuit could damage equipment along the line. Transformers – in particular, large ones on high-voltage transmission lines – would be a prime target, although many have fuses that would prevent the high current from continuing through the device for more than an instant. A less likely target would be the heart of generators, since there are almost always protective elements between the line and the stator. However, transformers are a large enough target themselves. While companies have modest stockpiles of extra smaller transformers, such as the ones used in the distribution grid, there are not many spare large transformers. If enough were damaged, new ones would need to be manufactured, which could take half a year.

5.3 Human Operators

While in contemporary times the everyday running of the power grid is done automatically by EMSs, there are still human operators sitting all day and night in the main control room in front of a large *mapboard* – a display that represents the current state of the region of the grid for which the control station is responsible. Although their usual responsibilities involve non-emergency situations, such as coordinating scheduled line outages for maintenance, they are trained to respond to emergency situations.

One element behind the magnitude of the 2003 Northeast blackout was an error in the EMS run by the company that oversaw the area where the cascade began. In particular, this error delayed the alarm from going off for over an hour. During this delay the mapboard showed the system in fine working order. This caused the operators to not only not notice the problems and thus not take any corrective action, but also to ignore calls from neighboring control centers asking them about the instability emanating from the faulty control area. The final report on the blackout from the U.S.-Canada Power System Outage Task Force [19] determined that there was a large window of time in which corrective actions could have been taken. If the control center's EMS had been working

correctly, the control center operatives would have been notified of the problems with the grid and would most likely have taken corrective actions that would have prevented the massive cascading outage.

The 2003 Northeast blackout example clearly suggests a factor that would greatly improve the chances of an attack's widespread success: preventing control center operatives from taking corrective measures. Many in the industry believe that a cyber attack at a given control station is no worse than a physical attack there. However, imagine a control station being subject to a silent cyber attack, in which the station managers do not know that their station is compromised. The attacker creates a dangerous imbalance in the grid but makes sure that the control system software continues to give normal readings. When neighboring controllers call the compromised station, they are told that everything is under control, and the problem has time to grow, perhaps enough to create a larger cascading failure.

The attack method of changing what control center operators see suggests another idea: create imaginary problems in the system designed to prompt the operator to perform a 'corrective' action that the attacker desires. The same

reason/action pairings listed in Figure 1 that can be used when sending imaginary bad data to an EMS can be used here. While most of the changes in the grid are made by the EMS control software, sometimes the decision is left up to the operator. For instance, in one EMS, some problems are presented to the operator along with possible fixes and the time by which a corrective action must be taken to avoid damaging equipment. There must be some point at which this method will not work, when the operators realize that, for instance, the data they see is too illogical to be real. However, as one control center manager states, “Status of devices is gospel” [17]. It may take quite some time or strange circumstances for an operator to begin questioning the validity of the data coming from the grid’s trusted devices.

5.4 Intercompany Interactions

Often neighboring companies share their real-time data amongst themselves, via, for instance, ICCP (Inter-Control Center Communication Protocol). This suggests other data-obscuring attack vectors similar to those described for within a company. Data sent from an attacked control center to the neighboring centers could be modified to hide the instability, thereby giving it more time to exacerbate. On the other hand, attackers could send imaginary bad data from

one control center to another in order to evoke an action in the latter's center. If, for instance, it is easier to infiltrate control station A than it is B, the target station, then the attackers could send incorrect instable data from A to B so that B cuts its section of the grid off from A's section as a protective measure, which would cause some degree of instability in B's section. For example, if there was power flowing to B's region from A's region, then separating the two regions would lower the power supply in B's region, which would lower the voltage until load was shed or local generation was increased.

Presenting false normal data in order to prevent corrective measures is important not only within a company and between companies, but also between the company and the RTO, which often also has a station that continually receives, analyzes, and monitors real-time data from the companies in its purview. In the 2003 Northeast blackout there was a problem not only with the original control room's EMS alarm, but also with the RTO's EMS: the RTO's state estimator had not been restarted after maintenance that day. If the RTO had had a working state estimator, its operator would have noticed the irregular state of the originating area and made sure the control station took corrective action.

6 Conclusion

In this paper I explain the basics of the power grid's electrical and cyber infrastructures and discuss ways that attackers could subvert the grid's normal operations and controls either directly (via electronic commands) or indirectly (via injecting false data to get the EMS control software or human operator to take harmful seemingly-corrective action). I also discuss the following items: the points in the system that can be compromised; using cyber attacks to create permanent physical damage; affecting larger areas through cascading outages; choosing an initial attack time based on load derived from weather, RTO data, and/or EMS data; and creating longer-lasting disturbances by crippling the control systems or the bootstrapping process or by manipulating the state information received by the control programs or operators located at the control center(s) of the companies in the attack area, their neighbors, and their RTO. The latter two topics are unique to cyber attacks; further, the multiple simultaneous contingencies required to incite a cascading outage are much easier to create via a cyber attack than via a physical one.

Much needs to be done to improve the security of the power grid. Important

initial steps would include: 1) reduce the possibility of cascading outages through increasing the number of transmission lines, further study of such outages from a systems and power engineering perspectives (such as Dobson et al. [20]), introducing FACTS controllers to substations, and improving contingency algorithms (such as in Mili et al. [21] and Motter [22]); and 2) prevent false data injection through adapting some of the many mechanisms that are used in computer security when certain entities in a system are not trusted, such as encryption and signatures.

7 Acknowledgements

I wish to thank first and foremost my advisor, Professor Sean Smith, as well as the Dartmouth College Security Lab and Professor Peter Sauer of the University of Illinois. This material is based upon work supported by the National Science Foundation under Grant No. CNS-0524695 and by the Department of Energy under Award No. DE-OE0000097. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation or the Department of Energy.

8 References

- [1] U.S. Energy Information Administration. (2010, Jan. 21). *Count of Electric Power Industry Power Plants, by Sector, by Predominant Energy Sources within Plant*. [Online]. Available: <http://www.eia.doe.gov/cneaf/electricity/epa/epat5p1.html>
- [2] U.S. Census Bureau. (2010, May 8). *2009 Population Estimates*. [Online]. Available: http://factfinder.census.gov/servlet/DTable?_bm=y&-geo_id=01000US&-ds_name=PEP_2009_EST&-mt_name=PEP_2009_EST_G2009_T001
- [3] A. S. Brown, "SCADA vs. the Hackers - Can Freebie Software and a Can of Pringles Bring Down the U.S. Power Grid?," *Mechanical Engineering*, vol. 124, no. 12, Dec. 2002.
- [4] K Barnes et al., "Review Of Supervisory Control And Data Acquisition (SCADA) Systems," *Idaho National Engineering and Environmental Laboratory*, INEEL/EXT-04-01517, Jan. 2004.
- [5] P. Oman et al., "Safeguarding IEDs, Substations, and SCADA Systems against Electronic Intrusions," *Proceedings of the 2001 western power delivery automation*, 2001.
- [6] C. Taylor et al., "Risk Analysis and Probabilistic Survivability Assessment (RAPSA): An Assessment Approach for Power Substation Hardening," presented at the ACM Workshop on Scientific Aspects of Cyber Terrorism, Washington, DC, 2002.
- [7] C.-W. Ten et al., "Vulnerability Assessment of Cyber-Security for SCADA Systems," *IEEE Transactions on Power Systems*, vol. 23, no. 4, pp. 1836-1846, 2008.
- [8] Office of Electricity Delivery and Energy Reliability, "Study of Security Attributes of Smart Grid Systems – Current Cyber Security Issues," *U.S. Department of Energy* INL/EXT-09-15500, Apr. 2009.
- [9] L. Grigsby et al., Ed. *The Electric Power Engineering Handbook*, Second ed., Boca Raton, FL: CRC Press, 2001.
- [10] Y. Liu et al., "False Data Injection Attacks against State Estimation in Electric Power Grids," *Proceedings of the 16th ACM conference on Computer and communications security*, pp. 21-32, 2009.
- [11] North American Electric Reliability Corporation. (2010, Mar. 14). *"2007 - Disturbance Index - Public"*. [Online]. Available:

- <http://www.nerc.com/files/NERC%202009%20Jan-Dec%20DAWG%20Disturbance%20Reports.pdf>
- [12] J. Salmeron et al., "Analysis of Electric Grid Security under Terrorist Threat," presented at the INFORMS Annual Meeting, Denver, CO, 2004.
 - [13] J. Salmeron et al., "Worst-Case Interdiction Analysis of Large-Scale Electric Power Grids," *IEEE Transactions on Power Systems*, vol. 24, pp. 96-104, 2009.
 - [14] R. Albert et al., "Structural Vulnerability in the North America Power Grid," *Physical Review E*, vol. 69, no. 2, Mar. 2004.
 - [15] U.S. Department of Energy. (2009, Dec.). *National Electric Transmission Congestion Study*. [Online].
 - [16] J. Kappenman, "Geomagnetic Disturbances and Impacts upon Power System Operation," in *Electric Power Generation, Transmission, and Distribution*, L. Grigsby, Ed., ed Boca Raton, FL: CRC Press, 2007, pp. 16-1 to 16-22.
 - [17] V. Tsolias, Manager, Energy Management Group, National Grid. Northboro, MA: Apr. 28, 2010. [Interview].
 - [18] California ISO. (2010, Apr. 20). *OASIS - California ISO*. [Online]. Available: <http://oasis.caiso.com/>
 - [19] U.S.-Canada Power System Outage Task Force, "Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations," Apr. 5, 2004.
 - [20] Dobson et al., "Complex Systems Analysis of Series of Blackouts: Cascading Failure, Critical Points, and Self-Organization," *Chaos*, vol. 17, 2007.
 - [21] L. Mili and Q. Qiu, "Risk Assessment of Catastrophic Failures in Electric Power Systems," *Int. J. Critical Infrastructures*, vol. 1, no. 1, pp. 38-63, 2004.
 - [22] A. E. Motter, "Cascade Control and Defense in Complex Networks," *Physical Review Letters*, vol. 93, no. 9, Aug. 27 2004.