Title

Not Everything is Dark and Gloomy: Power Grid Protections Against IoT Demand Attacks

ISBN

9781939133069

Authors

Huang, Bing
Cardenas, Alvaro A
Baldick, Ross

Publication Date

2019

Peer reviewed

Not Everything is Dark and Gloomy: Power Grid Protections Against IoT Demand Attacks

Bing Huang, ***The University of Texas at Austin;*** **Alvaro A. Cardenas,**
University of California, Santa Cruz; **Ross Baldick,** ***The University of Texas at Austin***

https://www.usenix.org/conference/usenixsecurity19/presentation/huang

This paper is included in the Proceedings of the 28th USENIX Security Symposium.

August 14–16, 2019 • Santa Clara, CA, USA

978-1-939133-06-9

Open access to the Proceedings of the 28th USENIX Security Symposium is sponsored by USENIX.

Not Everything is Dark and Gloomy: Power Grid Protections Against IoT Demand Attacks

Bing Huang
The University of Texas at Austin
binghuang@utexas.edu

Alvaro A. Cardenas
University of California, Santa Cruz
alvaro.cardenas@ucsc.edu

Ross Baldick
The University of Texas at Austin
baldick@ece.utexas.edu

Abstract

Devices with high energy consumption such as air conditioners, water heaters, and electric vehicles are increasingly becoming Internet-connected. This new connectivity exposes the control of new electric loads to attackers in what is known as Manipulation of demand via IoT (MadIoT) attacks. In this paper we investigate the impact of MadIoT attacks on power transmission grids. Our analysis leverages a novel cascading outage analysis tool that focuses on how the protection equipment in the power grid as well as how protection algorithms react to cascading events that can lead to a power blackout. In particular, we apply our tool to a large North American regional transmission interconnection system consisting of more than 5,000 buses, and study how MadIoT attacks can affect this power system. To help assess the effects of such cyber attacks, we develop numerical experiments and define new and stronger types of IoT demand attacks to study cascading failures on transmission lines and their effects on the system frequency. Our results show that MadIoT attacks can cause a partition of the bulk power system, and can also result in controlled load shedding, but the protections embedded in the operation of the transmission grid can allow the system to withstand a large variety of MadIoT attacks and can avoid a system blackout.

1 Introduction

The vulnerability of Internet of Things (IoT) devices is a well-known problem [11, 25, 46]. Previous work has demonstrated that devices from cameras to door locks can be compromised directly or through their designated smart phone applications [29, 43]. A large-scale compromise of these devices can enable attackers to affect network infrastructures, as exemplified by the Distributed Denial of Service (DDoS) attacks by the Mirai botnet—which consisted of more than six hundred thousand IoT devices [13].

The collective effect of compromised IoT devices can go beyond traditional computer network infrastructures. Recent work proposed a novel form of attack called Manipulation of demand via IoT (MadIoT) [47], and showed that if an attacker compromised hundreds of thousands of high-energy IoT devices (such as water heaters and air conditioners), the attacker could cause various problems to the power grid, including (i) frequency instabilities, (ii) line failures, and (iii) increased operating costs. These attacks paint a dire picture of the security of the power grid as they show that a 30% increase in demand can trip all the generators in the US Western interconnection causing a complete system blackout, and a 1% increase of demand in the Polish grid results in a cascade of 263 transmission line failures, affecting 86% of the load in the system.

In this paper we re-evaluate the potential impact of MadIoT attacks by modeling in detail the protection equipment and the operational responses to sudden load changes in the power grid. Our analysis leverages a novel cascading outage analysis tool that focuses on how the protection equipment already embedded the power grid reacts during cascading events, where multiple protection equipment is activated one after the other.

Our analysis shows that while MadIoT attacks can create negative consequences on the power grid, the negative impact on the grid will not be as dire as originally thought. In particular, while the most powerful MadIoT attacks (assuming the attacker compromises more than 8 million air conditioners) might cause the power system to partition and operate as separate islands, or can also cause some controlled load shedding, our results show that creating a system blackout—which would require a black start period of several days to restart the grid— or even a blackout of a large percentage of the bulk power grid will be very difficult.

This paper is organized as follows. Section 2 introduces the background necessary to understand power systems and how our tool compares to state-of-the-art practices for cascading analysis. Section 3 presents the details of our simulations and models. Section 4 illustrates why our cascade analysis tool has advantages over competing alternatives in a simplified model used in previous work. Our main results focusing on the analysis of a large-scale North American interconnec-

tion undergoing MadIoT attacks are presented in Section 5. Section 7 summarizes related work and Section 8 provides conclusions, limitations, and future work.

2 Power Systems Background

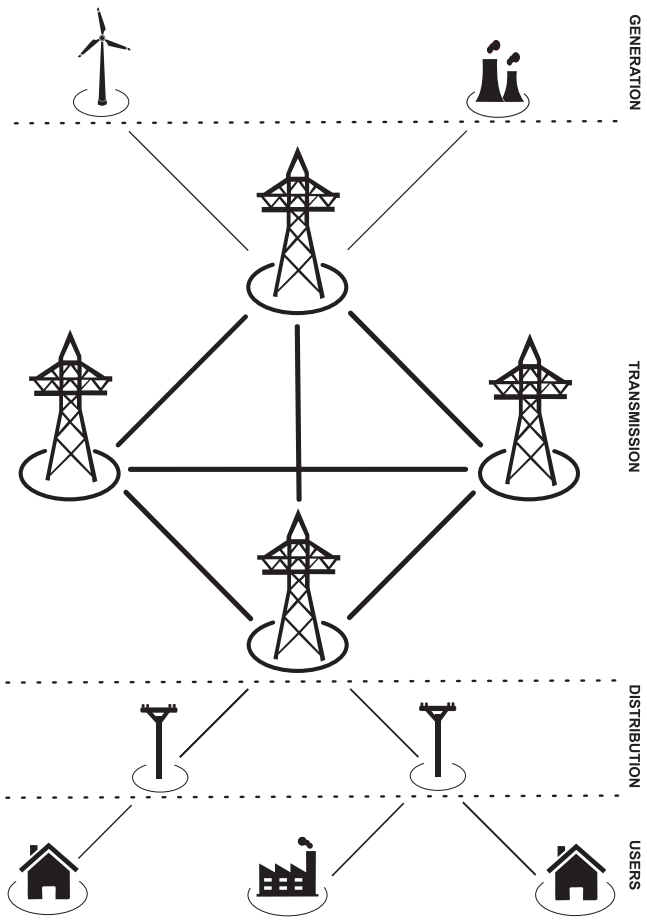

Figure 1: Generation and Transmission form the Bulk of the Power Grid. Transmission systems are redundant and have to satisfy the N-1 operation criterion, while Distribution systems are radial systems (non redundant) and affect a very small percentage of the system.

The objective of engineers and researchers in the power system industry is to deliver increasing amounts of electrical energy in a safe, clean, and economical manner [31]. The power grid has three major parts: (1) generation, (2) transmission, and (3) distribution. Electric power is generated wherever it is convenient and economical, and then it is transmitted at high voltages (100kV-500kV) in order to minimize energy losses—electrical power is equal to voltage times electrical current ($P = VI$), and given a constant power, high voltage lines have less electrical current, and therefore there is less energy lost as heat as the current moves through the transmission lines. Geographically, a distribution system is located in a smaller region thereby energy losses are less of a concern while safety (preventing accidents, fires, electrocutions, etc.) is more important, therefore they are operated at lower voltages. Figure 1 illustrates these three main parts of the grid. A distribution system is connected to a transmission system in a substation and the conductor that completes the connections is usually represented in electrical diagrams by nodes called **buses**.

Operators have to keep the nominal frequency (e.g., 60Hz in the Americas) and the transmission lines at their operating range (at a fixed voltage like 500kV, and with currents below a safety threshold) in order to ensure reliable operation of the grid. If there is a sudden increase in the demand of electricity, the frequency of the power grid tends to slow down, and automatic controls ramp up generation of electricity to take the frequency back to 60Hz. If there is a sudden decrease in the demand of electrical power, then the frequency of the grid tends to increase, and automatic controls then decrease generation of electrical power to reduce the frequency to the nominal level. Similarly sudden changes in electricity consumption might overload transmission lines and activate protection equipment (relays that prevent the flow of electricity through the line), and if this happens, the power is then distributed to other transmission lines.

2.1 Transmission vs. Distribution Outages

Large generation plants and the transmission network are usually referred to as the **Bulk Power System**, and this bulk power system is responsible for the reliable delivery of electricity to large areas. The bulk power system is an interconnected, redundant network that spans large regions—usually one country, but in North America there are three bulk systems: the Eastern Interconnection, the Western Interconnection, and Texas. In contrast, distribution systems are geographically smaller and their networks are mostly radial (i.e., non-redundant).

The **bulk** power system is designed and operated to satisfy the N-1 security criterion, which means that the system can lose any one of its N components (such as generators or transmission lines) and continue operating safely and serving the power supply to the customers in the large area. This operating criterion is mandatory and enforced by government entities, and therefore bulk power system operators have the incentives to make sure that their systems satisfy the N-1 criterion at any point in time, otherwise they get massive sanctions. In contrast, since distribution systems are usually non-redundant and serve customers in a regional area, they do not have to meet the same operating criterion.

The reason distribution systems do not have to meet the N-1 criterion is the scale of a system failure. A disruption in the **bulk** power grid will be the topic of national news headlines because it causes a blackout in a large part of the country

(sometimes even the whole country), while a disruption in the distribution system will usually only cause a localized outage (e.g., a neighborhood will be without electricity). Electric power in the distribution grid can also be more easily restored, while a system blackout of the bulk power system will require days of coordination in what is called **black start period.**

While distribution systems are not required to follow the N-1 criterion, there are separate criteria applied to them. For example, the hours of successful power supply to consumers as percentage of the total hours in a year is required to meet certain standard e.g. 99.999%. Other details of the distribution system will not be discussed as they go beyond the scope of this paper.

As we will show later in the paper, one of the protections embedded in the power system to prevent a bulk power outage is called **Under Frequency Load Shedding (UFLS)**, which is a mechanism where predetermined blocks of customers in the distribution system are automatically dropped from the system. This is a carefully selected procedure where electricity is not cut to safety-critical loads like Hospitals. We will show that some of the most severe MadIoT attacks will activate this protection and therefore can cause some controlled outages, but at the same time, these small outages are done in order to prevent that the bulk system goes into a cascading failure resulting in a system blackout.

2.2 Failure Analysis in the Bulk Power Grid

The power grid analysis tool we use in this paper was developed to address the limitations for modeling and analyzing cascading failures identified by the task force from the IEEE Power Engineering Society [14, 15]. As stated in these reports, most of the research in cascading failure analysis focuses on independent phenomenons, but these interactions are often ignored. In our recent work on cascading failures [33, 53–56] we have been developing a tool that captures the time interdependencies of all relevant protection equipment and stability studies in the power grid when multiple simultaneous (or quasi-simultaneous) contingencies occur. In this paper we adapt our tool to model MadIoT attacks. Before we discuss our approach in more detail, we now present related work in the analysis of failures in the power grid and discuss how our system compares to these approaches.

Cascading failure analysis has attracted a lot of attention from the research community [14, 44, 52]. There are two main approaches for studying cascading failures: stochastic models, and fine-grained simulations.

Stochastic models are used to evaluate the likelihood of a cascading event by giving us the probability of having incorrect settings for protection equipment in a given power system [26, 45]. To build these estimates, stochastic models perform a forensic analysis of previous cascading failures by looking at the properties of power systems just before they experienced a system blackout. Although these models provide a probabilistic insight of cascading events, they cannot be used to model the operation of a power system undergoing a cascade, which is particularly important when we want to understand how the system reacts to incidents in general (and cyber-attacks in particular). To understand the operation of the power system undergoing cascading failures we need to turn to detailed simulation models.

2.2.1 Power System Simulations

There are two main behaviors that we need to study when a system undergoes a failure:

1. **Transient Analysis** finds the behavior of the frequency in the power grid in the immediate aftermath of the incident. If the frequency deviates too far from 60Hz, some protection equipment will be activated. There are two options for transient analysis.

   (a) **No System Dynamics:** This is a very fast computational method where the behavior of all generators is simplified to only one generation machine. This allows us to evaluate how the frequency of the system behaves with big changes in electricity consumption. Several cascading studies use this method [35, 41]. This simplification cannot capture the frequency at every bus in the system (therefore it cannot model if a power system is partitioned into islands), nor model how each generator will react differently to cascading incidents (therefore it cannot model how the protection mechanism in each generator will activate).

   (b) **System Dynamics:** In this type of transient analysis we model all generators in the power system and all the frequencies in all the buses of the system. This is in line with one of the main objectives of a transient stability study—to determine whether the resulting angular separation between the machines in the system remains within certain bounds so that the system maintains synchronism [36]. Cascading analysis models with system dynamics are considered in [28, 34, 40].

2. **Steady-State Analysis** finds the voltages and currents of the system after all frequency equipment has tripped and can help us understand if the system ends up in a configuration where voltage protection or overcurrent protection equipment will activate. To compute these values, a power flow program uses Kirchhoff's physical laws to obtain the voltage magnitudes and phase angles at each bus of a power system. As a by-product of this calculation we can also compute real and reactive power flows in equipment such as transmission lines and transformers, as well as equipment losses [31]. There are two ways to perform steady state analysis:

(a) **DC Power Flow:** Direct Current (DC) Power Flow is a very fast way to compute voltages and phase angles. There are several cascading analysis studies that use DC Power Flow models [22, 27, 57]. DC power flow models however are approximations to AC models, and they do not show the variations on voltages that might trigger protection equipment, therefore DC methods are only valid when voltages are close to their nominal values, which rules out their use for modeling large-scale events such as MadIoT attacks.

(b) **AC Power Flow:** Alternating Current (AC) Power Flow is a more accurate (but computationally more expensive) way to analyze the steady state behavior of the power system. The only way to model voltage protection systems is with the use of AC power flow. Cascading analysis with AC power flow methods include [35, 41].

2.2.2 Power System Protections

In the previous subsection we have argued that the best practices for an accurate portrayal of power system behavior under *large-scale* events (i.e., events where voltages go beyond nominal values, and where individual generators might go beyond safety limits) is to use (1) System Dynamics for transient analysis, and (2) AC Power flow for steady-state analysis. In this section we describe how the results of our transient and steady-state simulations are used to evaluate how protection equipment in the power grid will react to changes in the operation of the system.

In particular, we model four protection mechanisms that are relevant for cascading analysis studies:

1. **Protection of Generators:** when the frequency of the system is too low or too high, the generator will be automatically disconnected from the power grid to prevent permanent damages to the generator.

2. **Under Frequency Load Shedding (UFLS):** if the frequency of the power grid is too low, controlled load shedding will be activated. As discussed before, this disconnection of portions of the distribution system is done in a controlled manner, while avoiding outages in safety-critical loads like hospitals. UFLS is activated in an effort to increase the frequency of the power grid, and prevent generators from being disconnected (as discussed in the point above).

3. **Overcurrent Protection:** if the current in a transmission line is too high, a protection relay will be triggered after time T. This activation time is based on an equation for current relays [10]. We will discuss in detail this equation when we describe our cascade outage analysis model.

4. **Over/Under Voltage Protection:** if the voltage of a bus is too low or too high, a voltage relay will be triggered after time T. This activation time depends on an equation modeling configuration thresholds and over/under voltage relay pick-up values [2].

2.2.3 Industry Practices

For day-to-day operations related to power grid failures, power operators focus on satisfying the N-1 criterion as this is the most important failure condition that is regulated and enforced by most electric regulatory agencies. Large-scale events such as a massive natural disaster, a terrorist attack, or a cyber-attack have not been a major priority for industry practices because the likelihood of these events is very small, and investment in preparing for these events has higher costs than responding to them when they happen [49].

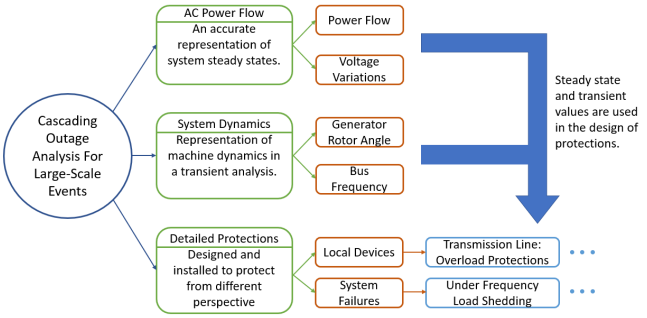

Figure 2: Analysis of Cascading Outages.

Therefore most of the industry efforts on cascading studies focus on smaller-scale events that initiate a cascade, and where the transient dynamics do not affect the cascade analysis too much. These efforts include the Transmission Reliability Evaluation of Large-Scale Systems (TRELSS) [32, 39] and the Oak Ridge-PSERC-Alaska (OPA) [18]. Similar problems have been studied by system operators like ERCOT [3]. Our tool on the other hand is designed for the study of the large disruptions in the operation of a power system like a deliberate cyber attack which can take hundreds of lines out in a short time, and therefore transient analysis has to be coupled with steady-state analysis.

The integration of (1) System dynamics, (2) AC power flow, and (3) the timing of protection equipment gives our tool a level of fidelity that goes beyond the current state-of-the-art practices [22, 27, 28, 34, 35, 35, 40, 41, 41, 57]. These three analysis techniques and their relationship are shown in Figure 2.

2.3 Contributions

Our contributions to the study of a MadIoT attacks compared to recent work [24, 47] include the following:

First, previous work considered transient and steady state simulation as separate use-cases (and in different inconsistent power systems), and as a result, the transient impacts on generators and system frequencies are not present in the power flow simulations. Therefore the predictions of cascading outages can differ between the two simulations. As we explain in Section 4.1, without the transient effect, the power flow solution will indicate a system blackout, while in reality Under Frequency Load Shedding will activate before generators start tripping and will prevent a system blackout.

Second, including the exact timing for the activation of a protection relay captures the realistic behavior of equipment in the power grid. Previous works on IoT attacks to the power grid [28, 34] do not represent the delay characteristic of protection equipment, but rather use models that appear to be based only on the immediate removal of an element after any amount of overload. Such a model violates NERC criteria for overload protection [1]. Our model is instead a discrete event simulator that does not assume that all relays will trip at the same time. In particular, we model equipment under stress, such as current overloads of 50-100% of the line rating. This model is based on the curves from manufacturers [2, 10] that relate the overload of the device to the time until it trips—e.g., if the overload of the line increases significantly, the trip time would be much shorter.

Third, we also perform the first large-scale transient analysis of MadIoT attacks on a real-world North American regional system with over 5,000 buses. This large-scale analysis shows that the most powerful MadIoT attacks can partition the bulk power system into three or more isolated islands. The power grid does not go into a system blackout, but each island will be more vulnerable to future contingencies. This is a new effect that has not been considered before.

Because by repeating the same attack conditions from previous work did not cause any blackout in our system, we introduce new variations of the MadIoT attacks where for example, the attacker systematically tries to create oscillations of demand in order to drive the system into a more vulnerable state before launching the second stage of the attack.

Finally, all our simulations are done in PowerWorld [4], which is an industry-standard transient and AC steady-state solver, as its basic building block, so the basic physics of the system are represented with industry-accepted fidelity.

These contributions are summarized in Table 1.

Table 1: Contributions

Contributions		**Our Work**	**Previous Work [24, 47]**
Simulations	Transient	PowerWorld	PowerWorld
	Steady-state	PowerWorld	Matlab
	Combined transient and steady-state analysis	Yes	No
Transient Analysis	**Under Frequency Protection**	Yes	No [47]
	Frequency in all buses	Yes	No [24]
Steady-state Analysis	**Power Flow**	AC	Not Specified
	Time for Over Current Protection	Yes	No
	Time for Voltage Protection	Yes	No
New MadIoT Attacks	**IoT Demand Increase and Decrease**	Yes	No
	IoT Repeat	Yes	No
Scale of Analysis	**Case used in Transient Simulation**	A North American regional system with over 5,000 buses	Up to WSCC 9-bus system
	Case used in Steady-state Simulation	A North American regional system with over 5,000 buses	Polish system with 3,120 Buses

3 Cascading Outage Analyzer

This section summarizes our Cascading Outage Analyzer (COA) tool. The COA model considers both steady-state and transient stability analysis in different time scales but coordinated so the transition of system stability from one steady-state operating point to another is present. The basic model checks for conditions that would trigger protective relays, and assesses the time when relays will be triggered.

The framework of the COA is described in Figure 3. The simulation has both transient and steady state parts. For each contingency, a transient simulation is run using the PowerWorld transient simulation tool. If the system reaches a stable state, then simulation results are sent to the steady state simulation as initial values, where an AC power flow is run. Based on the resulting line flows and voltage magnitudes, the timing for activating protection equipment is then computed.

If there are any new protection equipment activated from this steady state simulation, the new outage will be modeled and the next iteration of simulation will start using the PowerWorld transient simulation tool. This multi-time scale process continues until no outage occurs in both the transient and steady-state parts of the simulation, or until the transient simulation is unable to solve the problem, in which case an "algorithmic non-convergence" is declared to have occurred as a proxy to a system blackout.

We now describe how each of the four protection systems we consider are modeled.

3.1 Protection of Generators

If a mismatch between generation and load occurs, there will be a frequency deviation from the desired nominal value (if there is more load than generation, the frequency of the system will decrease, and if there is more generation than load, the frequency of the system will increase). A big frequency deviation may trigger generator under- and over frequency-protections.

Transient stability or rotor angle stability is the ability of

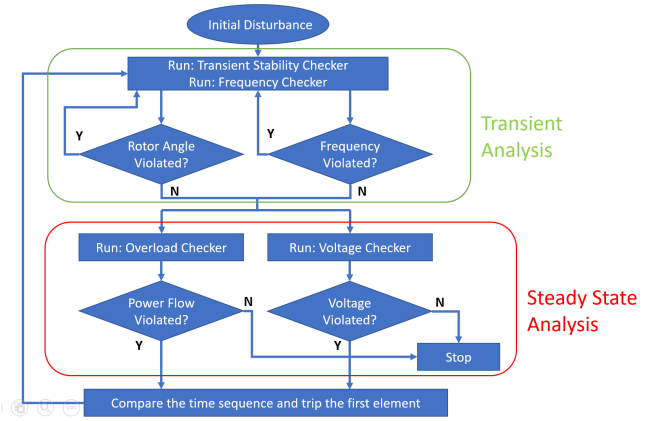

Figure 3: Overview of our Cascading Outage Analysis Tool.

the power system to remain in synchronism when subjected to large transient disturbances [37]. We choose to use time-domain simulation because the time-domain simulation takes into account the full system dynamic model and constantly checks that inter-machine rotor angle deviations lie within a specific range of values.

We use the PowerWorld transient stability solver to numerically calculate the system response after a fault. If the rotor angle deviation of a generator is bigger than a certain threshold, e.g., 100 degrees, the generator will be automatically tripped and removed from the power grid to prevent permanent damages. The disconnection of the generator won't be immediate after crossing a threshold, but it will be dependent on the amount of time that it remains in the unsafe region. We will discuss the exact configuration parameters for disconnecting a generator later in the paper.

3.2 Preventing the Tripping of Generators

When the system loses a generator or when there is a sudden increase in the load, the frequency of the power grid decreases rapidly. A countermeasure to prevent the activation of (more) under-frequency generator protections is a mechanism called Under Frequency Load Shedding (UFLS). The predominant system condition addressed by IEEE C37.117 involves the use of protective relays for under frequency shedding of connected load in the event of insufficient generation or transmission capacity within a power system. Therefore, we include UFLS along with over/under frequency generator tripping as frequency outage checkers in the COA model. Taking into consideration these protections embedded in power systems is one of the reasons we obtain different results when compared to previous work.

3.3 Overcurrent Protection

Disconnecting transmission lines because of a thermal limit violation is one of the most common events in cascading outages [53]. We trigger overcurrent protections based on the results from our steady-state results. The status and dispatch set points of units at the end of the PowerWorld transient analysis are used as starting points for the PowerWorld AC power flow simulator. An inverse-time overcurrent equation described in the Siemens SIPRO-TEC 5 Current Relay [10] is implemented in our model. The time when the over current relay trips the element is determined by equation (1),

$$T = \frac{0.14}{(\frac{I}{I_{th}})^{0.02} - 1} T_p[s], \tag{1}$$

where I_{th} is the current threshold value of the relay, and T_p is the setting value of the relay. Both values are set by the relay operator. I is the current on the monitored component such as a transmission line or a transformer. The value of T in (1) determines when the protection will be activated. It is important to understand that overloading the line past its nominal rating does not immediately result in a transmission outage. Simplified models that do not account for the detailed behavior of protection equipment are likely to consider that a line gets out of service when in reality it keeps operating (it just sags). This is another of the reasons we obtain different results from previous work.

3.4 Over/Under Voltage Protection

Another typical pattern associated with cascading outages is an under (or over) voltage problem. When the system is highly stressed, the voltage profiles of power systems may decline. Even if the AC power flow calculation converges, if a bus voltage stays below the lower limit in our simulations, a load-shedding protection mechanism will be triggered in order to return the bus voltages to their limits [53].

The bus voltages are required to be on a range for the safe operation of the connected generators. A generator may also be disconnected if the voltage of the connected bus goes out of limits for too long.

We implement in our simulator a standard inverse time characteristic equation described in ABB RXEDK 2H time over/under voltage relay [2] to find the timing for the activation of voltage protection equipment. The time duration until the under or over voltage relay trips is determined in equations (2) and (3),

$$T = \frac{k}{(\frac{U}{U_{th}}) - 1}[s], \tag{2}$$

$$T = \frac{k}{1 - (\frac{U}{U_{th}})}[s], \tag{3}$$

where k is the inverse time constant, U_{th} is the over/under voltage relay pick-up value, and U is the user defined relay operating value. The values of T in equations (2) and (3) determine when the protection will activate. As with the line overload model, over/under voltages do not immediately result in a bus outage.

4 Considerations for Modeling the Impact of IoT Attacks

This section will demonstrate the contribution of applying our cascading outage analyzer in the study of IoT demand attacks and in particular, this section will compare our results with previous work in order to show why we obtain different results. We will start our analysis with a relatively simple but standard Western System Coordinating Council (WSCC) model with 9 buses and 9 lines, as this is a model that has been used in previous work. We will also discuss in more detail some of our considerations for modeling the impact of IoT attacks. In the next section we will provide a detailed study on a model of a real-world North American system.

In this section we use the over/under frequency generation protection and Under-frequency load shedding parameters from Table 2 and Table 3. In the next section we will explain in more detail these parameters.

4.1 The Need for Combining Transient and Steady-State Simulations

Since the operation of a power system after a disturbance is a continuous process over a long time frame, a closed-loop structure of the cascading outage analyzer can better approximate the operations of the power system over various time scales after a disturbance. As previously discussed, the results and states of the system after the transient simulation are stored and set as the starting point of the steady-state simulations. The cascading outage generated from steady state simulations, if there is any, is then used as the initial condition in the transient simulation for the next loop.

Previous work considered transient and steady-state simulations as separate, and as a result, the transient impacts on generators and system frequencies are not present in the power flow simulations. Therefore the predictions of cascading outages can differ when compared to our work. Let us look at an example to see a possible inconsistency, while emphasizing the importance of a combined transient/steady-state simulation for the analysis of cascading outages caused by IoT demand attacks.

Figure 4 shows the WSCC 9-bus system considered by Soltan et al. [47]. Consider an IoT demand attack that increases all loads by 15% in the system. Now let us see what happens if a transmission line is removed if the power flow is over its rated capacity [20]. If the transient impacts of this

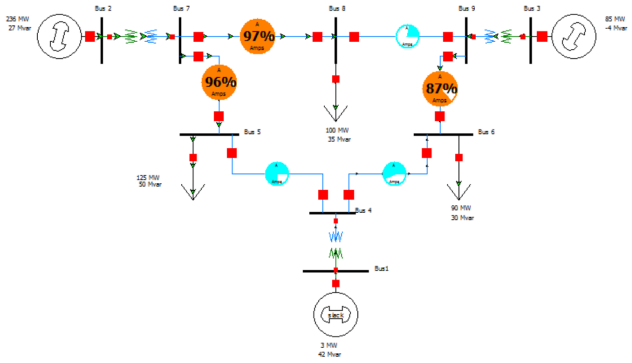

Figure 4: PowerWorld 9-bus system.

attack are not considered, the results from the steady-state power flow would indicate a line outage between bus 7 and bus 8, as highlighted with a red circle (shows the percentage of the rated capacity) in the top left corner in Figure 5.

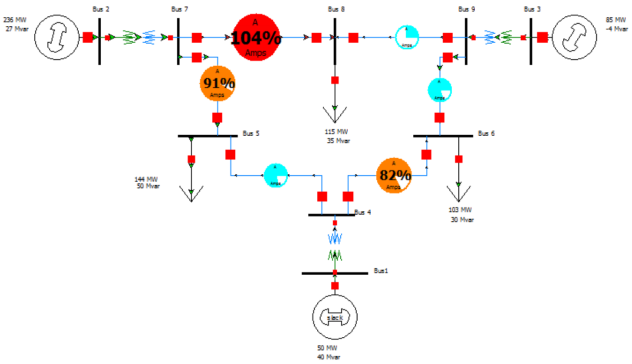

Figure 5: Power flow results of 15% of load increase.

However, because of the sudden load increase caused by the MadIoT attack, load and generation are not balanced and the frequency of the system will be affected. A frequency protection relay would disconnect a generator from the system if the frequency of the system stays lower or higher than the generator's threshold values for too long in order to prevent permanent damage to the generator. Figure 6 shows the frequency responses to the 15% load increase. We can see that the system frequency starts to decline after the attack starts (the attack starts after one second). The frequency relays then disconnect all the generators in the system two seconds after the frequency drops below the threshold of 58 Hz (table 2). Therefore, this results in a blackout in the transient simulation of the IoT demand attack. These transient stability results are different from the steady state stability study, which identifyied only one cascading line outage as discussed in the previous paragraph.

This is a motivating reason to include transient and steady state analysis together in a single simulation. Because transient and steady-state simulations are connected in a closed

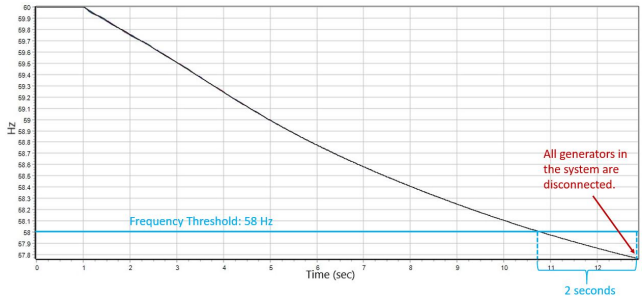

Figure 6: Frequency responses to the 15% of load increase in the transient simulation.

loop in our model, the transient solution at the end of the simulation time will be used as an initial condition for the steady-state power flow simulation. In this example, if the under frequency load shedding is not considered, which will be discussed in Section 4.2, the transient solution would include the fact that all three generators were disconnected from the system. Thus, the power flow solution would indicate a system blackout.

4.2 Under Frequency Load Shedding

Under Frequency Load Shedding (UFLS) is a countermeasure applied by bulk power system operators [5] to reduce the incidence of generator under-frequency tripping, which is a great danger to the reliable operation of the power systems. UFLS is a coordinated disconnection of small and non-critical (e.g., no Hospitals are ever disconnected) loads to prevent a large blackout.

To illustrate why it is important to consider UFLS in the simulation of IoT demand attacks, let us first take a second look at Figure 6. As observed, after the 15% load increase attack, the system frequency starts to decrease. Because there is no action that could relieve the imbalance between the increased load and unchanged generation, the system frequency declines fast until it drops below the thresholds of frequency protections at generators. Because the frequency stays below the thresholds for longer than the delay time set at the frequency protections, the generators are disconnected and there is a system blackout.

Now, let us compare the simulation results when we incorporate UFLS as defined by the parameters in Table 3. Figure 7 shows the frequency response to the 15% system demand increase attack on the WSCC 9-bus system. The system frequency declines after the IoT load increase attack starts at one second of the simulation time. The frequency of the system then reaches the first UFLS threshold at 59.3 Hz, and as a result, 5% of the system load is disconnected. However, this is not enough and the system frequency keeps declining until it reaches the second threshold: 58.9 Hz, and at that time a total of 15% of the system demand is disconnected and the frequency stops decreasing and starts to stabilize to its desired state. The system frequency reaches a new stable state and there are no generator disconnections from the system.

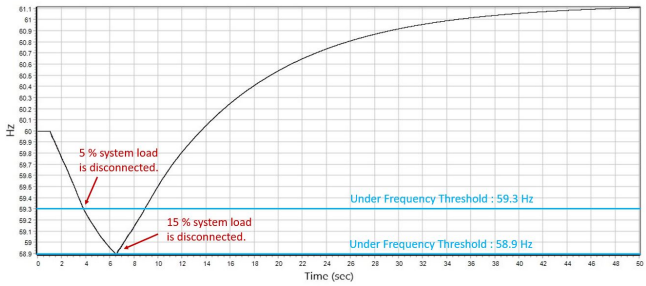

Figure 7: Frequency responses with Under Frequency Load Shedding to the 15% of load increase in the transient simulation.

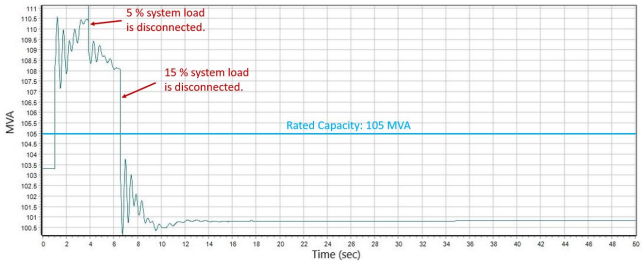

Figure 8: Power flow on the transmission line connected between bus 7 and bus 8 in the transient simulation

What is more, because of UFLS, the system load is reduced to a level where no transmission line is overloaded, and therefore there are no cascading outages. In Figure 8, we can see that the transmission line between bus 7 and bus 8 in Figure 4 is overloaded after the IoT demand increase attack begins at one second. However, the power flow on the line soon decreases following the load shedding event caused by UFLS and remains below its rated capacity at the end of the transient simulation. As discussed in Section 4.1, a power flow steady state simulation starts based on the solution of the transient simulation; the results of this new steady state stability analysis are shown in Figure 9. We can see that no line is overloaded and the combined transient and steady-state simulations end.

The example in this subsection shows that the simulation results will be significantly affected if UFLS protections are considered. In fact, by including UFLS, the closed-loop transient and steady state simulations used in this work generates a result suggesting that the system would shed some demand, but all the system transmission lines and generators will remain in operation. This result is different from the cascading line outage suggested by our steady-state simulation illustrated in Figure 5 and the complete system blackout suggested by previous work.

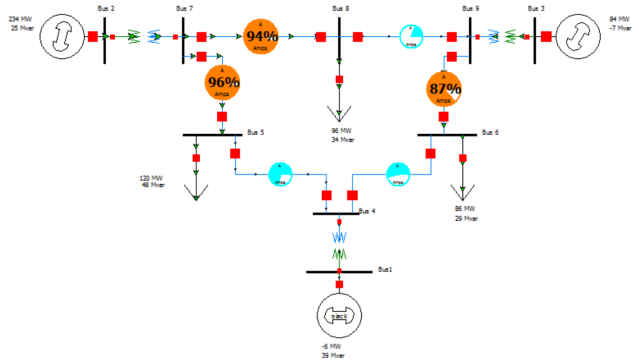

Figure 9: Power flow results after the transient simulation with UFLS.

4.3 Frequency Response Model

UFLS protections are indeed considered in some previous work [24]. However, the simplified frequency response model used by the authors is not a good fit to analyze IoT demand attacks. The system frequency responses used by Dabrowki et. al [24] model the power grid as a single large machine that represents an "aggregation" of all the synchronous generators in the system.

A synchronous machine is associated with a rotating magnetic field winding that induces alternating voltages in a armature windings of the stator. The frequency of the induced alternating voltages and of the resulting currents that flow in the stator windings when a load is connected depends on the speed of the rotor. The frequency of the stator electrical quantities is thus synchronized with the rotor mechanical speed, hence the designation "synchronous machine" [37]. When two or more synchronous machines are interconnected, the stator voltages and currents of all the machines must have the same frequency and the rotor mechanical speed of each is synchronized to this frequency. Therefore, the rotors of all interconnected synchronous machines must be in "synchronism" [37].

In contrast, the assumption of Dabrowski et. al [24] is that every generator in the system will respond to a disturbance exactly the same. In other words, the implicit assumption of this model is that all the generators in the system will always keep synchronism and respond identically. However, when the system is under a significant disturbance, generators will respond differently to the disturbance and the system will have the risk of losing synchronism in a short time after the disturbance. In some scenarios, the frequency protections will contribute to a lack of synchronism, and therefore, the frequencies at different buses will diverge from synchronism. All of this frequency diversity can not be reflected in the single machine mode [24]. A detailed discussion of why this phenomenon is important will be demonstrated in subsection 5.4, where we show how different parts of the grid start operating at different frequencies and therefore the system becomes a set of islands operating semi-independently.

4.4 Line Overloads

The line overload outage models also play an important role in understanding the impact of MadIoT attacks. Previous work [47] relied on the criteria described by Cetinay et al. [20], where a line will be removed from the system if the steady-state results indicate that the power flow on the line is greater than its rated capacity. When a transmission line is overloaded, the heat generated from the extra power flow on the line will sag the transmission line. Although it exposes the line to a possible outage from faults associated with ground element or vegetation, it does not necessarily cause any immediate real danger to the system. In fact, under an emergency, the system operator is allowed to use overloaded transmission lines for additional transmission capacity [6]. Therefore, instead of immediately removing the overloaded lines, we utilize a model that calculates the time of tripping given the overload level. The details are described in Section 3.3. The time inverse calculation in the outage protection mechanism will result in a quick tripping time for the lines that are heavily overloaded. In this way, we approximate the different actions taken at different levels of overload on transmission lines.

4.5 IoT Demand Attacks

In addition to fixed demand increase (or decrease) attacks, we also consider attacks that increase and then decrease the load. The intuition for this attack is that the first part of the attack will force automatic responses from the grid (such as UFLS) and therefore when the system starts operating with a reduced load, a reversal in the load (a big decrease) can drive the system to a potentially unstable state. After initial attack increasing the demand, the attackers will decrease the demand when they think the system frequency reverses due to UFLS and intend to overshoot the system frequency over the thresholds of generator frequency protections in the hopes of causing a generator disconnection.

This demand increase and decrease attack was studied by Dabrowski et al. [24]. However, our results will differ because of their simplification of the frequency model, as discussed in Section 4.3. In addition, if the attacker can cyclically increase and then decrease demand, it is reasonable to assume that the attacker is capable of repeating this attack. The simulation results and detailed discussions of the experiments are shown in Section 5.

5 Simulation Results in a Large Power System

The study case we use to analyze the impact of the IoT demand attacks is a large North American regional system with more than 5,000 buses, and as such it is the largest study

Table 2: Over/Under Frequency Generator Tripping. Source: Section 2.6.1 of [5].

Over Frequency Threshold	Time Delay	Under Frequency Threshold	Time Delay
60.6 Hz	9 min	59.4 Hz	9 min
61.6 Hz	30 sec	58.4 Hz	30 sec
61.8 Hz or above	0 sec	58.0 Hz	2 sec
		57.5 Hz	0 sec

done on the impact of IoT attacks on power systems. Unfortunately, because our close collaboration with the operator of this power systems we are required to maintain the confidentiality of this system and we are not allowed to share the name of the system or details of their network topology. Before we describe our simulation results we clarify our assumptions.

5.1 Assumptions

We state three main assumptions about an IoT demand attack:

1. IoT attackers have full and unlimited ability to control the compromised portion of loads;
2. The actions of attackers to increase or decrease the compromised loads are simultaneous;
3. The portion of the system demand compromised by the cyber attackers are evenly distributed at each demand connection point in the transmission system.

The third assumption is a speculation about the scalability of an IoT attack. For example, if the adversary is able to compromise one brand of air conditioner, they can systematically apply the attack to as many air conditioners as possible in the target system. Thus, if the total energy capacity of all such air conditioners is 10% of the system demand, this 10% of demand is likely to be spread to every demand connection point in the transmission system.

5.1.1 Parameters Used for Protection Equipment

There are two protections implemented in the transient simulation, namely Over/Under Frequency Generator Tripping (O/UFGT) and Under Frequency Load Shedding (UFLS). If the frequency at a bus deviates from a predefined threshold for more than a specific time period, the generator connected to that bus will be tripped, and a certain percentage of load connected to the bus will be shed. The details of O/UFGT and UFLS are shown in Table 2 and Table 3 specifically.

Since the current and voltage responses in the system are normally slower than frequency responses, the Time Inverse Overload, Time Inverse Under Voltage Load Shedding, and Time Inverse Over Voltage Generator Tripping are modeled in the steady state simulation. Each protection checker will calculate tripping times once the current flow on branches or the voltage at buses exceed the thresholds. The element (branch, generator, or load) with the shortest tripping time will be tripped as the initial conditions for the next iteration of transient simulation. The parameters of the steady state protection models described in equations (1-3) are listed in Table 4.

Table 3: Under Frequency Load Shedding. Source: Section 2.6.1 of [5].

Frequency Threshold	System Load Relief	Time Delay
59.3 Hz	5 %	0 sec
58.9 Hz	15 %	0 sec
58.3 Hz	25 %	0 sec

Table 4: Steady State Protections. Source: [53]

	Over Load	Over/Under Voltage	
		over	under
Threshold	$I_{th} = 2 \times line$ limit [amps]	$U_{th} = 1.3$ [pu]	$U_{th} = 0.8$ [pu]
Parameters	$T_p = 0.05$	$k = 0.5$	$k = 0.5$

5.2 Demand Increase Attacks

The most intuitive MadIoT attack against the power grid is a sudden increase of demand. This will attempt to overload the transmission lines and potentially cause cascading failures.

5.2.1 1% Demand Increase Attack

Previous work showed that a 1% increase attack against the Polish power grid in 2008 caused cascading failures. In their system, a 1% load increase corresponded to 210MW, requiring the adversary to compromise about 210,000 air conditioners. In our system, one percent of the load is equivalent to 822.7 MW, which would require the attacker to compromise approximately 822,000 air conditioners.

Figure 10 shows the bus frequency responses after 1% of load increase at second 1 and Figure 11 shows the power flow on branches as a percent of the branch rated capacity. We can observe that the bus frequencies shown in Figure 10, decline after the attack at second 1 except for very few buses that are connected to the region outside of the system with DC tie lines (the ones that remain at 60Hz on top of the diagram) and thereby remain less affected.

The rest of the frequencies decline from 60 Hz to 59.875 Hz in about 9 seconds and settle to a new stable state towards the end of the transient simulation. As indicated in table 2 and table 3, the system frequency doesn't violate any thresholds of

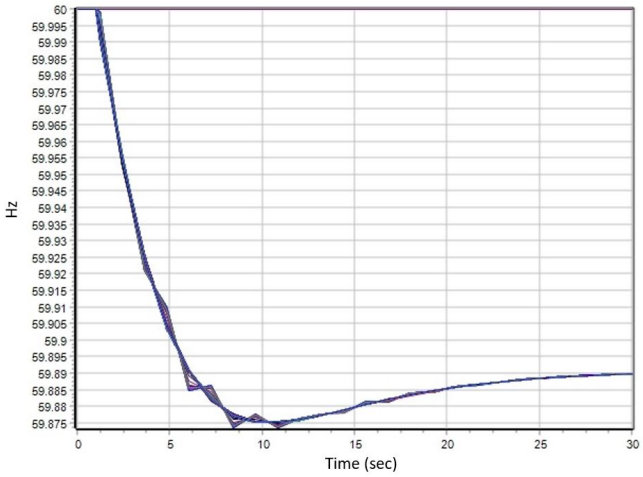

Figure 10: Frequency Response to 1% System Load Increase.

frequency protections on generators and loads. Notice that we focus our study in a short time window, since 30 seconds of transient simulation is enough to display the moving trends of the frequency in this case. In short, we can see the how the frequency is affected after the attack; however, as long as the bus frequency converges to a stable level, driving the frequency back to 60Hz can be accomplished either automatically or manually over a longer time scale.

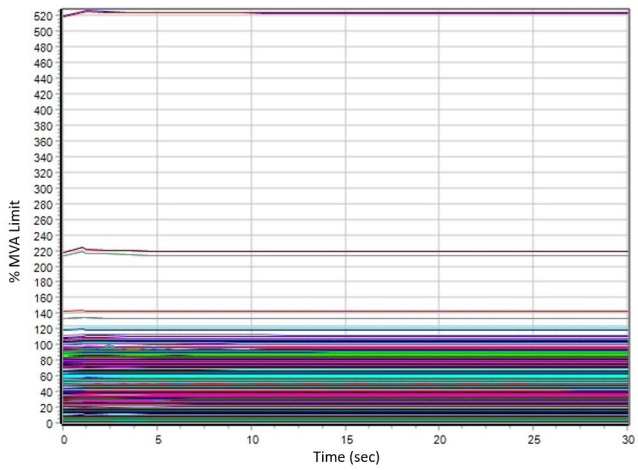

Figure 11: Branch Flow after 1% System Load Increase.

In Figure 11, we can see that the power flow of some branches slightly increases after the attack at second 1. However, no transmission line is overloaded resulting from the IoT 1% load increase attack. Note that some branches are initially overloaded before the simulation and remain unchanged during the simulation and the overload outage checker is not activated on those lines under the assumption that protection in the actual system would not have been activated under these conditions.

In summary, a 1% load increase attack does not affect our system and there is no need to activate any protection equipment as the transmission lines remain operating in their nominal values and the frequency of the system does not reach thresholds to activate any protection.

In contrast to our results, Soltan et al. [47] find that with a 1% increase in load there could be cascading outages in the summer peak of the Polish grid. We are surprised that a sudden 1% increase in load can lead to cascades in a power system. The reason for our surprise is the N-1 security criterion.

The N-1 criterion requires that electricity systems be operated to be able to withstand sudden step changes in the supply-demand balance due to outages of generation. The NERC disturbance control performance standard [8] requires any system to be able to withstand "the most severe single contingency" which may include certain common-model double outages. For ERCOT, for example, (the Power Grid of Texas) this amounts to always having 2700 MW or more of reserves to cope with a simultaneous outage of nuclear units having total production of around 2700 MW. To put that in perspective, peak load in ERCOT is around 70GW, and 1% of 70GW is 700MW, which is much smaller than the 2700MW of reserves carried in ERCOT to satisfy the N-1 criterion.

While an increase by 700MW in load due to an IoT attack (and the reaction by generation reserves) would result in somewhat different changes in transmission flows compared to the effect of a 700MW decrease in generation (and the reaction by generation reserves), we believe that it is unlikely that an increase in load of 1% would result in any unacceptably adverse conditions on the transmission system. This is because load is geographically distributed around the system, so that it is unlikely for there to be a more than a 1% increase in most transmission flows, and it is unlikely that the system is operating such that a 1% increase in current would immediately trigger the overload protection.

In the Eastern and Western Interconnections of North America, the total load is much larger (several hundred GW) but even 1% of this would only amount to slightly more than the double outage of a nuclear unit (plus it would require millions of compromised IoT devices). To summarize, the results of the Polish power grid reported by Soltan et al. [47] suggest that the system being modeled is not N-1 secure.

5.2.2 10% Demand Increase Attack

Ten percent of system load in our case study is equivalent to 8,227.3 MW, which would be equivalent to an adversary controlling over eight million air conditioners. Figure 12 shows the bus frequency responses after a 10% load increase attack at 1s and Figure 13 shows the power flow on branches as a percent of the branch rated capacity.

To better understand the variations of power flow depicted in Figure 13, let's first take a look at Figure 12. From Figure 12, we can observe that the bus frequencies plummet after the

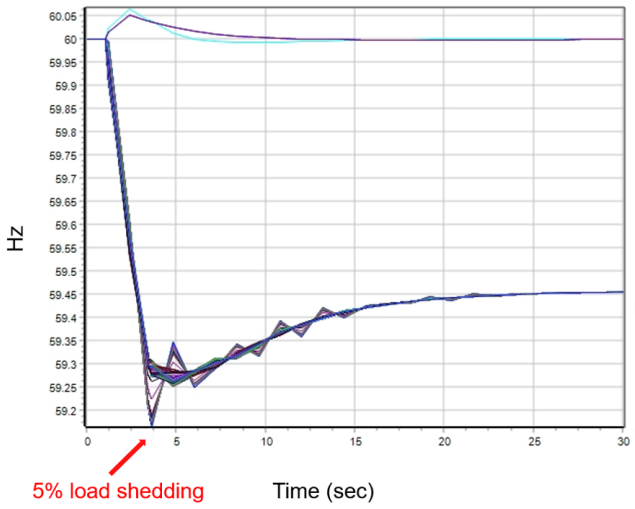

Figure 12: Frequency Response to 10% System Load Increase.

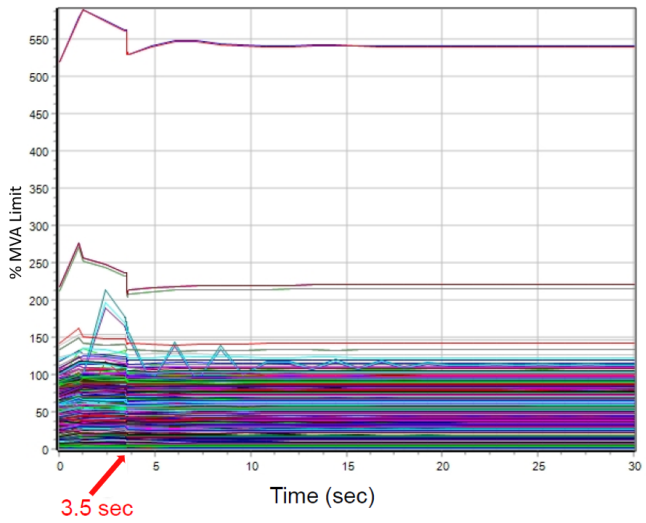

Figure 13: Branch Flow after a 10% System Load Increase.

attack begins (1s). The only lines that are not affected are the few buses that connected our power grid to another region outside the system with DC tie lines (the frequencies at the top of the figure).

In contrast to the previous 1% attack, with a 10% demand increase the power system needs to activate protection algorithms; in particular, 5% UFLS is activated at 3.5 seconds by shedding 5% of the system load. Again, as long as the bus frequency converges to a stable level, the differences between the converged value and its initial value of 60 Hz can be fixed either automatically or manually over a longer time scale.

Although the under frequency shedding has no deliberate time delay as indicated in Table 3, a 0.02 second of relay operation time is included in the simulation. Therefore, the load shedding occurs 0.02 seconds after the time frequency falls below the first UFLS threshold of 59.3 Hz.

In Figure 13, we can see that the power flows of some branches increase after the attack starts (1 sec.). However, the power flows of those branches drop to or gradually decrease to roughly their initial values after the under frequency load shedding protection is activated at 3.5 seconds. Therefore, at the end of the simulation there is no additional transmission line overloaded. Note that some branches are initially overloaded before the simulation and remain unchanged during the simulation and so, as in the previous example, the protection mechanisms for these transmission lines are not activated.

Even with the assumption of millions of compromised IoT devices to affect 10% of our load, our results show that the power grid protections to prevent generators from disconnecting from the system are effective in mitigating any further problem. The amount of UFLS is intended to reflect ERCOT standards. The Eastern and Western Interconnections may have overall lower levels of UFLS than ERCOT; however, they have much larger levels of inertia than ERCOT.

5.3 Increase and Decrease Attack

One of the characteristics of IoT attacks, is that they are highly distributed they are hard to detect. Once the load is compromised, the compromised devices are unlikely to be removed from the grid (or the Internet) in a short time after they launch the first attack. Therefore attackers can launch a sequence of attacks, the first as an attempt to drive the system to a vulnerable state, and the second to exploit that vulnerability.

In the last attack we saw how under frequency load shedding successfully prevented a cascading failure of transmission lines from a single 10% load increase attack. However, a sophisticated attacker can identify when the system frequency starts rebounding after the initial drop, and can attempt to make this trend continue by immediately decreasing electricity consumption. This can cause a frequency overshoot that may trigger the action of over-frequency protection relays on the generators and disconnect them from the power grid; creating another cycle of frequency decrease along with new load shedding etc.

A straightforward approach in this experiment is to increase the load at the first attack and decrease the same amount of load at the second attack. However, we investigate a potentially worse scenario where in the second attack, we decrease by twice the amount of the load increase in the first attack (minus the percentage of the load that the attacker loses control of after the under frequency load shedding implementation).

The result in Figure 14 shows that the frequency does overshoot after the loads decrease at second 20, however the system frequency tends to stabilize at 61.7 Hz, which happens about 10 seconds later. From Table 2, we can observe that

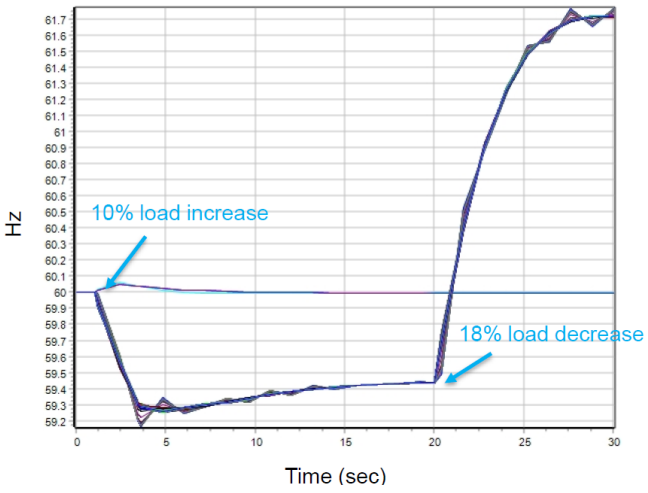

Figure 14: Frequency Response to a Cycle of Load Increase and Decrease

61.7 Hz will not cause an immediate generation trip by the frequency protections at the generators.

As mentioned in Section 5.2, a 10% system load compromised by the adversary is already a pessimistic assumption. We take this even further to 20% of the system load in this simulation to see if the IoT attack can cause a cascading result. However, we still do not observe an immediate generation trip after this demand increase and decrease attack in a system that is intended to reflect ERCOT standards for UFLS.

5.3.1 Under Frequency Load Shedding in a Repeated IoT Attack

We have explored the results of an attack "cycle" of load increase and decrease. The adversary could continue repeating this attack cycle of increasing and reducing the compromised load as long as their capabilities are not disabled by the load shedding mechanism.

The under frequency load shedding would disconnect some amount of demand each time when the IoT attack causes the frequency drop below any thresholds. Once the load is disconnected by Under-Frequency Load Shedding (UFLS) systems, the restoration of shed load is coordinated between the Independent System Operator (ISO), Transmission Service Providers (TSPs) and Distribution Service Providers (DSPs) [5]. It is fair to assume that such restoration, which requires coordination between different entities may take a relatively long time to complete. Therefore, a potential negative effect of such repeated attacks is that they can deplete the under frequency load shedding resources before they are restored, which might eventually lead to having no more UFLS protections against the attacks and will eventually cause a generator to trip.

The result in Figure 12 shows that although the system frequency needs additional measures to be brought back to its initial frequency of 60 Hz, the frequency decline caused by 10% of system load increase can be stopped by only 5% of system load shedding. In Table 3 we can see that in ERCOT, 25% of the system load is contracted as UFLS. Under this condition, the adversary needs to apply the attack at least five times to deplete the UFLS resources. What is more, additional under-frequency relays may be installed on transmission facilities with the approval of the ISO provided the relays are set at 58.0 Hz or below in the real system [5]. That means, in reality, the adversary may need to apply the attack even more times to deplete the UFLS and cause a possible system failure.

Therefore, it may take many cycles of IoT demand increase and decrease attacks to deplete the UFLS resources, and these cycles will not only deplete the resources from a defensive stand point, but also the resources available to the adversary as each activation of UFLS will remove loads controlled by the attacker. Therefore, the efficiency, even the feasibility of the approach of using up the UFLS by such repeated IoT demand attack remains unclear.

5.4 Bifurcations, and Generator Tripping

5.4.1 30% Load Increase Attack

In Section 5.3, we briefly discussed the potential threats of generator disconnections caused by over frequency protections. In this section, we extend this discussion to IoT attacks that specifically target disturbing frequency and causing generator disconnections by frequency protection. In order to observe the response of frequency protection at generators, we study the impact of a MadIoT attack consisting of a load increase or decrease by 30%.

In previous work [47], this 30% load change attack was able to disconnect all generators of the (simplified) North American Western Interconnection, causing a complete system blackout. In our system, a 30% load increase attack would require the attacker to compromise about 24 million air conditioners.

Figure 15 shows the frequency response of our system to a MadIoT attack that increases the system load by 30%. First, we can observe that due to the sudden load increase, the bus frequencies decline dramatically and some of them drop quickly below the first UFLS threshold of 59.3 Hz. At this point 5% of the system load is disconnected by UFLS.

We notice that the frequency in some buses decline at a slower rate than others and they do not reach any UFLS thresholds. For convenience, we name this set of buses Group 1. The buses with DC tie lines are again less affected, and we call this set of buses Group 2. The group of buses whose frequencies decline faster and drop below UFLS thresholds are named Group 3. The group names are indicated in Figure 15.

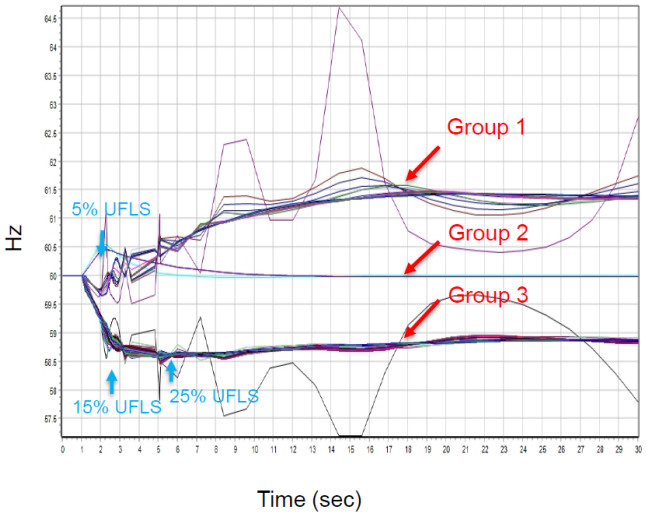

Figure 15: Frequency Response to 30% Load Increase.

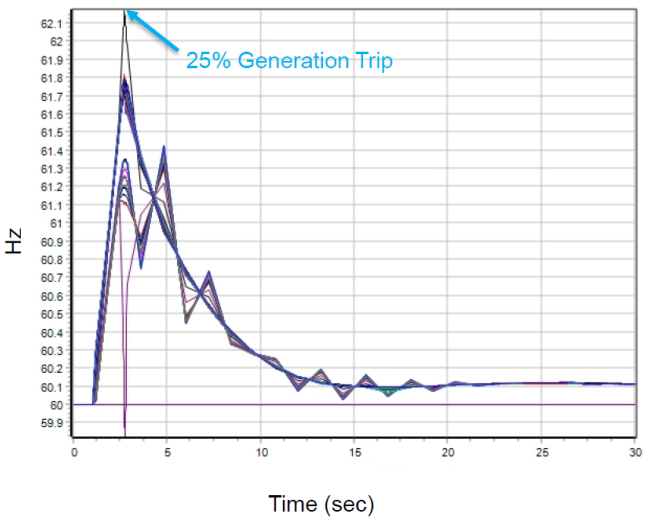

Figure 16: Frequency Response to 30% Load Decrease.

Notice that, even within a group, the frequency responses are not exactly the same. Because of the first UFLS action, the frequency deviation between buses increases. After the 5% load shed, the frequency of Group 1 starts to increase–potentially this group has more generators in their region—while the frequency of Group 3 keeps declining–indicating that this region of the grid has insufficient generation of electrical power.

Shortly afterwards, the frequency of Group 3 declines to the point where the second and third UFLS thresholds, 58.9Hz and 58.3Hz, need to be activated (at around 2.6 seconds and 5.6 seconds respectively). An additional 10% of system load is disconnected in each occasion. The frequency deviation between Group 1 and Group 3 gets larger after the two UFLSs. What is more, the frequency deviation between buses in a group, especially in Group 1, increases after the actions of UFLS.

After the three activations of UFLSs for group 3, which disconnect a total 25% of system load, the frequency decline at Group 3 is stopped. Because there is no additional load shedding, the frequency at Group 1 stops increasing as well. Thus, although the bus frequencies have not converged at the end of the simulation, they stop diverging and there is no need to activate frequency protections to disconnect generators.

5.4.2 30% Load Decrease Attack

We now study what happens if instead of increasing the load by 30%, we decrease the load by 30%. In this case we expect the frequencies in all buses to increase dramatically; furthermore, because UFLS can only be activated when the frequency is decreasing, then we know that there are no immediate protections to prevent a generator from disconnecting from the grid because of its over-frequency protections.

Figure 16 shows the system frequency response to a MadIoT attack that decreases the system load by 30%. We can see that the bus frequencies increase after the attack and a few of them go above the threshold of immediate over frequency protections at generators, which is 61.8 Hz within 5 seconds. The over frequency protections then disconnect generators, resulting in a 25% reduction of system generation. After the tripping of generators, the bus frequencies reduce and converge to a value close to 60 Hz and no more protection actions or failures are observed.

Because we model the time in which each generator is disconnected, we can see that not all of them are disconnected synchronously, as suggested in prior work, but at different times, depending on their configuration settings. When some generators are disconnected, then the frequency drops and is stabilized by the remaining generators.

Compared to the system frequency response to an IoT attack that increases the system load, we find that the bus frequencies react differently to the IoT attack that decreases the system load. In Figure 16, although the frequencies of some buses increase faster than those in some other buses, the frequencies gradually converge after 25% of the system generation is tripped. One of the conclusions we can draw from this comparison is that a quick protection reaction in big scales like the generation tripping in Figure 16 performs better than the gradual protection actions like the load shedding in Figure 15 in terms of the system frequency restoration.

We also find that the tripped generations in this simulation consist of a significant amount of wind generation. The benefit of disconnecting the wind generation or any generation that doesn't provide inertia in this condition is that the system loses less inertia after the over frequency protection action. Therefore, the system doesn't become weaker in terms of

maintaining frequency stability. This phenomenon suggests that generation that doesn't provide inertia could be included in the over frequency protection to protect the system against any following attacks targeted at disturbing the system frequency after an IoT attack.

6 Limitations

Our results also assume that all grid operators satisfy the N-1 security criterion. This is the general practice and should be expected as operators can get massive fines if they are found to be in violation of this criterion. Having said that, some blackouts have occurred because operators believe they are satisfying the N-1 criterion but a misconfigured protection device that should have been activated during an event was not activated, and this created an unanticipated N-2 event that initiated a cascading failure. As discussed in our summary of related work on cascading analysis, stochastic models can complement our approach by establishing the risk or likelihood that one of our protection devices does not work as expected and causes a series of cascading events.

We believe the type of protections considered in this study is the subset of the protections in power systems that would contribute to a cascading outage the most after a disturbance in the system. However, future work can be done to explore the impacts from other protections that are commonly equipped in the power systems e.g. differential and distance protections on buses [31]. In addition, in this study, we considered only an IoT demand attack that is evenly distributed across all the load points in the system. However, in future work, we will consider how feasible it is to compromise a large-scale set of high wattage IoT devices in a specific geographical area such that that target only a part of the system.

7 Related Work

The importance of stronger cyber security requirements in SCADA systems is highlighted by recent experiences in Ukraine. On December 23rd 2015, a third party illegally accessed the computer and Supervisory Control and Data Acquisition (SCADA) systems of three regional electricity distribution companies in Ukraine. Investigations revealed that a malware named BlackEnergy had infected the SCADA systems after successful spear phishing attacks. Seven 110 kV and twenty-three 35 kV substations were disconnected for three hours resulting in several outages that caused approximately 225,000 customers to lose power across various areas [19]. The following year, on December 17th 2016, a second power outage occurred in Ukraine and deprived part of its capital, Kiev, of power for over an hour. An assessment was made that a more advanced form of malware called "Industroyer", was used in the second cyber attack against the power grid in Ukraine [23].

While both security researchers and industry practitioners have worked on the security of the power grid for a decade, their focus has been on understanding and preventing attacks to devices in the bulk of the power grid [7, 17, 38, 48, 51], i.e., the components controlling the operation of the electrical transmission system in large geographical areas and the Supervisory Control and Data Acquisitions (SCADA) systems.

While in the U.S. the bulk power system is regulated to maintain a minimal set of cybersecurity standards [7], there is a growing push to start improving the security of systems in the distribution network. On October 19th 2017, the Federal Energy Regulatory Commission (FERC) proposed new mandatory cybersecurity controls to address the risk posed by, for example, smaller grid control centers that are typically less critical than major control centers, but which are nonetheless vulnerable to attacks [9].

Load-altering attacks have been previously studied in demand-response systems [12, 16, 21, 30, 42, 50]. Demand-response programs provide a new mechanism for controlling the demand of electricity to improve power grid stability and energy efficiency. In their basic form, demand-response programs provide incentives (e.g., via dynamic pricing) for consumers to reduce electricity consumption during peak hours. Currently, these programs are mostly used by large commercial consumers and government agencies managing large campuses and buildings, and their operation is based on informal incentive signals via as phone calls by the utility or by the demand-response provider (e.g., a company such as Enel X) asking the consumer to lower their energy consumption during the peak times. As these programs become more widespread (targetting residential consumers) and automated (giving utilities or demand-response companies the ability to directly control the load of their customers remotely) the attack surface for load altering attacks will increase.

8 Conclusions

This paper presents a study of the impacts of IoT demand attacks on power systems using the cascading outage analysis in a North American Regional Interconnection System.

From the simulation results, we show that, 1% of load increase attack does not interrupt any generator, load, or transmission line in the system. We also find that, thanks to under frequency load shedding protections, a 10% of sudden IoT load increase does not cause a cascading failure on the transmission lines.

A "frequency swing attack" is defined as a cycle of load increase and decrease attacks with the aim to push the frequency outside the safety limits of the generators. However, the frequency swing attack doesn't show an ability to cause an immediate disconnection of generators. We also discussed a possible repeated frequency swing attack and the potential impact of depleting the UFLS resources. Our analysis shows that the effectiveness of such attack would be impacted by

any additional frequency protection measures in the system, and by the diminishing resources that the adversary would have to continue the attacks.

We also considered high-impact attacks with control of 30% of the system load. The simulation results show that under a sudden IoT attack increasing 30% of the system demand, load shedding by UFLS would split the frequencies of the buses into islands of different operating regions of the grid. In contrast, a 30% decrease of the load would cause the frequency of the system to increase above the thresholds for over-frequency protections, and will result in the disconnection of some (but not all) generators. Our results show that the actions of UFLS and over frequency protection are sufficient to prevent an immediate system failure over a short time after the attack. Additional actions may be needed over a longer time scale to restore the stable operation of the system, but the main point is that a system blackout will likely not occur in this situation. In addition, we discover that including generations that are not providing inertia in the over-frequency protections would benefit the system in case of following IoT attacks targeted at disturbing the system frequency.

Our results show a different perspective on the risks of IoT attacks to the power grid and will hopefully serve as a starting point for new discussions to assess this threat. We show that while immediate cascading failures or a total system blackout will be very hard to achieve, the power system will still suffer negative consequences. First, UFLS will disconnect various consumers from the power grid. This is done to prevent further damage to the grid, but several consumers will be affected. Second, our attacks show that with millions of high-energy IoT devices, the attacker can potentially cause a bifurcation of the frequency in the power grid, forcing the grid to operated as separate islands and driving it to a more vulnerable state.

Acknowledgments

This work is supported by NSF CRISP awards CMMI-1541159 and CMMI-1925524, by a grant from the University of Texas National Security Network, and by a Defense Threat Reduction Agency award HDTRA1-14-1-0021.

References

[1] NERC Standard PRC-023-4. https://www.nerc.com/pa/Stand/Reliability%20Standards/PRC-023-4.pdf, 2015.

[2] Time over/under voltage relay and protection assemblies model rxedk 2h and raedk user manual. www.abb.com/product/us/9AAC30405217.aspx, ABB Inc, 2004.

[3] 2018 regional transmission plan scope and process. http://www.ercot.com/content/wcm/key_documents_lists/108892/2018_RTP_Scope_and_Process_draft_clean.pdf, Accessed, 2018.

[4] Powerworld simulator 20. https://www.powerworld.com/, Accessed, 2019.

[5] ERCOT nodal operating guides section 2 system operations and control requirements. www.ercot.com/content/wcm/current_guides/53525/02_030116.doc, ERCOT, 2016.

[6] ERCOT nodal operating guides section 4 emergency operations. http://www.ercot.com/content/wcm/libraries/147359/February_1__2018_Nodal_Operating_Guide.pdf, ERCOT, 2018.

[7] Cyber risk preparedness assessment table-top exercise 2012 report. May, 2013.

[8] NERC Standard BAL-002-1a. https://www.nerc.com/files/bal-002-1a.pdf, NERC, 2012.

[9] FERC sets rules to protect grid from malware spread through laptops. Washington Examiner, October,2017.

[10] Siprotec 5 distance protection and line differential protection and overcurrent protection for 3-pole tripping 7sa84, 7sd84, 7sa86, 7sd86, 7sl86, 7sj86 technical data. www.energy.siemens.com, Siemens AG, 2012.

[11] Omar Alrawi, Chaz Lever, Manos Antonakakis, and Fabian Monrose. Sok: Security evaluation of home-based iot deployments. In *SoK: Security Evaluation of Home-Based IoT Deployments*, page 0. IEEE.

[12] Sajjad Amini, Fabio Pasqualetti, and Hamed Mohsenian-Rad. Dynamic load altering attacks against power system stability: Attack models and protection schemes. *IEEE Transactions on Smart Grid*, 9(4):2862–2872, 2018.

[13] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Luca Invernizzi, Michalis Kallitsis, et al. Understanding the mirai botnet. In *USENIX Security Symposium*, pages 1092–1110, 2017.

[14] Ross Baldick, Badrul Chowdhury, Ian Dobson, Zhaoyang Dong, Bei Gou, David Hawkins, Henry Huang, Manho Joung, Daniel Kirschen, Fangxing Li, et al. Initial review of methods for cascading failure analysis in electric power transmission systems ieee pes cams task force on understanding, prediction, mitigation and restoration of cascading failures. In *2008 IEEE Power and Energy Society General Meeting-Conversion and Delivery of Electrical Energy in the 21st Century*, pages 1–8. IEEE, 2008.

[15] Ross Baldick, Badrul Chowdhury, Ian Dobson, Zhaoyang Dong, Bei Gou, David Hawkins, Zhenyu Huang, Manho Joung, Janghoon Kim, Daniel Kirschen, et al. Vulnerability assessment for cascading failures in electric power systems. In *2009 IEEE/PES Power Systems Conference and Exposition*, pages 1–9. IEEE, 2009.

[16] Carlos Barreto, Alvaro A Cárdenas, Nicanor Quijano, and Eduardo Mojica-Nava. Cps: Market analysis of attacks against demand response in the smart grid. In *Proceedings of the 30th Annual Computer Security Applications Conference*, pages 136–145. ACM, 2014.

[17] Carlos Barreto, Jairo Giraldo, Alvaro A Cardenas, Eduardo Mojica-Nava, and Nicanor Quijano. Control systems for the power grid and their resiliency to attacks. *IEEE Security & Privacy*, 12(6):15–23, 2014.

[18] Benjamin A Carreras, Vickie E Lynch, Ian Dobson, and David E Newman. Critical points and transitions in an electric power transmission model for cascading failure blackouts. *Chaos: An interdisciplinary journal of nonlinear science*, 12(4):985–994, 2002.

[19] Defense Use Case. Analysis of the cyber attack on the ukrainian power grid. *Electricity Information Sharing and Analysis Center (E-ISAC)*, 2016.

[20] Hale Cetinay, Saleh Soltan, Fernando A Kuipers, Gil Zussman, and Piet Van Mieghem. Analyzing cascading failures in power grids under the ac and dc power flow models. *SIGMETRICS Performance Evaluation Review*, 45(3):198–203, 2017.

[21] Bo Chen, Nishant Pattanaik, Ana Goulart, Karen L Butler-Purry, and Deepa Kundur. Implementing attacks for modbus/tcp protocol in a real-time cyber physical system test bed. In *Communications Quality and Reliability (CQR), 2015 IEEE International Workshop Technical Committee on*, pages 1–6. IEEE, 2015.

[22] Jie Chen, James S Thorp, and Ian Dobson. Cascading dynamics and mitigation assessment in power system disturbances via a hidden failure model. *International Journal of Electrical Power & Energy Systems*, 27(4):318–326, 2005.

[23] Anton Cherepanov. Win32/industroyer, a new threat for industrial control systems. *White paper, ESET (June 2017)*, 2017.

[24] Adrian Dabrowski, Johanna Ullrich, and Edgar R Weippl. Grid shock: Coordinated load-changing attacks on power grids: The non-smart power grid is vulnerable to cyber attacks as well. In *Proceedings of the 33rd Annual Computer Security Applications Conference*, pages 303–314. ACM, 2017.

[25] Tamara Denning, Tadayoshi Kohno, and Henry M Levy. Computer security and the modern home. *Communications of the ACM*, 56(1):94–103, 2013.

[26] I. Dobson. Estimating the extent of cascading transmission line outages using standard utility data and a branching process. In *2011 IEEE Power and Energy Society General Meeting*, pages 1–3, July 2011.

[27] Margaret J Eppstein and Paul DH Hines. A "random chemistry" algorithm for identifying collections of multiple contingencies that initiate cascading failure. *IEEE Transactions on Power Systems*, 27(3):1698–1705, 2012.

[28] D. Fabozzi and T. Van Cutsem. Simplified time-domain simulation of detailed long-term dynamic models. In *2009 IEEE Power Energy Society General Meeting*, pages 1–8, July 2009.

[29] Earlence Fernandes, Jaeyeon Jung, and Atul Prakash. Security analysis of emerging smart home applications. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 636–654. IEEE, 2016.

[30] Jairo Giraldo, Alvaro Cárdenas, and Nicanor Quijano. Integrity attacks on real-time pricing in smart grids: impact and countermeasures. *IEEE Transactions on Smart Grid*, 8(5):2249–2257, 2016.

[31] J Duncan Glover, Mulukutla S Sarma, and Thomas Overbye. *Power System Analysis & Design, SI Version*. Cengage Learning, 2012.

[32] R. C. Hardiman, M. Kumbale, and Y. V. Makarov. Multiscenario cascading failure analysis using trelss. In *CIGRE/IEEE PES International Symposium Quality and Security of Electric Power Delivery Systems, 2003. CIGRE/PES 2003.*, pages 176–180, Oct 2003.

[33] Bing Huang, Mohammad Majidi, and Ross Baldick. Case study of power system cyber attack using cascading outage analysis model. *IEEE PES GM, Portland OR*, 2018.

[34] S. K. Khaitan, Chuan Fu, and J. McCalley. Fast parallelized algorithms for on-line extended-term dynamic cascading analysis. In *2009 IEEE/PES Power Systems Conference and Exposition*, pages 1–7, March 2009.

[35] Daniel S Kirschen, Dilan Jayaweera, Dusko P Nedic, and Ron N Allan. A probabilistic indicator of system stress. *IEEE Transactions on Power Systems*, 19(3):1650–1657, 2004.

[36] Prabha Kundur, Neal J Balu, and Mark G Lauby. *Power system stability and control*, volume 7. McGraw-hill New York, 1994.

[37] Prabha Kundur, John Paserba, Venkat Ajjarapu, Göran Andersson, Anjan Bose, Claudio Canizares, Nikos Hatziargyriou, David Hill, Alex Stankovic, Carson Taylor, et al. Definition and classification of power system stability. *IEEE transactions on Power Systems*, 19(2):1387–1401, 2004.

[38] Yao Liu, Peng Ning, and Michael K. Reiter. False data injection attacks against state estimation in electric power grids. In *Proceedings of the 16th ACM Conference on Computer and Communications Security*, CCS '09, pages 21–32, New York, NY, USA, 2009. ACM.

[39] F. Xia M. Kumbale, T. Rusodimos and R. adapa. Trelss: A computer program for transmission reliability evaluation of large-scale systems. *User's Referecne Manual*, 2, 1997.

[40] Hong Tao Ma and Badrul H Chowdhury. Dynamic simulations of cascading failures. In *2006 38th North American Power Symposium*, pages 619–623. IEEE, 2006.

[41] Shengwei Mei, Yixin Ni, Gang Wang, and Shengyu Wu. A study of self-organized criticality of power system under cascading failures based on ac-opf with voltage stability margin. *IEEE Transactions on Power Systems*, 23(4):1719–1726, 2008.

[42] Amir-Hamed Mohsenian-Rad and Alberto Leon-Garcia. Distributed internet-based load altering attacks against smart power grids. *IEEE Transactions on Smart Grid*, 2(4):667–674, 2011.

[43] Muhammad Naveed, Xiao-yong Zhou, Soteris Demetriou, XiaoFeng Wang, and Carl A Gunter. Inside job: Understanding and mitigating the threat of external device mis-binding on android. In *NDSS*, 2014.

[44] Milorad Papic, Keith Bell, Yousu Chen, Ian Dobson, Louis Fonte, Enamul Haq, Paul Hines, Daniel Kirschen, Xiaochuan Luo, Stephen S Miller, et al. Survey of tools for risk assessment of cascading outages. In *2011 IEEE Power and Energy Society General Meeting*, pages 1–9. IEEE, 2011.

[45] M. Rahnamay-Naeini, Z. Wang, N. Ghani, A. Mammoli, and M. M. Hayat. Stochastic analysis of cascading-failure dynamics in power grids. *IEEE Transactions on Power Systems*, 29(4):1767–1779, July 2014.

[46] Eyal Ronen, Adi Shamir, Achi-Or Weingarten, and Colin O'Flynn. Iot goes nuclear: Creating a zigbee chain reaction. In *Security and Privacy (SP), 2017 IEEE Symposium on*, pages 195–212. IEEE, 2017.

[47] Saleh Soltan, Prateek Mittal, and H Vincent Poor. Blackiot: Iot botnet of high wattage devices can disrupt the power grid. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pages 15–32, 2018.

[48] Siddharth Sridhar, Adam Hahn, and Manimaran Govindarasu. Cyber–physical system security for the electric power grid. *Proceedings of the IEEE*, 100(1):210–224, 2011.

[49] Nassim Nicholas Taleb. *The black swan: The impact of the highly improbable*, volume 2. Random house, 2007.

[50] Rui Tan, Varun Badrinath Krishna, David KY Yau, and Zbigniew Kalbarczyk. Impact of integrity attacks on real-time pricing in smart grids. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 439–450. ACM, 2013.

[51] Chee-Wooi Ten, Chen-Ching Liu, and Govindarasu Manimaran. Vulnerability assessment of cybersecurity for scada systems. *IEEE Transactions on Power Systems*, 23(4):1836–1846, 2008.

[52] Marianna Vaiman, Keith Bell, Yousu Chen, Badrul Chowdhury, Ian Dobson, Paul Hines, Milorad Papic, Stephen Miller, and Pei Zhang. Risk assessment of cascading outages: Methodologies and challenges. *IEEE Transactions on Power Systems*, 27(2):631, 2012.

[53] Yezhou Wang and Ross Baldick. Cascading outage analysis using sequential outage checkers. *Modeling, Simulation, And Optimization for the 21st Century Electric Power Grid*, 2013.

[54] Yezhou Wang and Ross Baldick. Case study of an improved cascading outage analysis model using outage checkers. In *Power and Energy Society General Meeting (PES), 2013 IEEE*, pages 1–5. IEEE, 2013.

[55] Yezhou Wang and Ross Baldick. Interdiction analysis of electric grids combining cascading outage and medium-term impacts. *IEEE Transactions on Power Systems*, 29(5):2160–2168, 2014.

[56] Yezhou Wang, Chen Chen, Jianhui Wang, and Ross Baldick. Research on resilience of power systems under natural disasters—a review. *IEEE Transactions on Power Systems*, 31(2):1604–1613, 2015.

[57] Jun Yan, Yufei Tang, Haibo He, and Yan Sun. Cascading failure analysis with dc power flow model and transient stability analysis. *IEEE Transactions on Power Systems*, 30(1):285–297, 2015.