BlackIoT: IoT Botnet of High Wattage Devices Can Disrupt the Power Grid

Saleh Soltan, Prateek Mittal, and H. Vincent Poor, *Princeton University*

https://www.usenix.org/conference/usenixsecurity18/presentation/soltan

This paper is included in the Proceedings of the
27th USENIX Security Symposium.

August 15–17, 2018 • Baltimore, MD, USA

ISBN 978-1-939133-04-5

BlackIoT: IoT Botnet of High Wattage Devices Can Disrupt the Power Grid

Saleh Soltan
Department of Electrical Engineering
Princeton University
ssoltan@princeton.edu

Prateek Mittal
Department of Electrical Engineering
Princeton University
pmittal@princeton.edu

H. Vincent Poor
Department of Electrical Engineering
Princeton University
poor@princeton.edu

Abstract

We demonstrate that an Internet of Things (IoT) botnet of high wattage devices–such as air conditioners and heaters–gives a unique ability to adversaries to launch large-scale coordinated attacks on the power grid. In particular, we reveal a new class of potential attacks on power grids called the Manipulation of demand via IoT (MadIoT) attacks that can leverage such a botnet in order to manipulate the power demand in the grid. We study five variations of the MadIoT attacks and evaluate their effectiveness via state-of-the-art simulators on real-world power grid models. These simulation results demonstrate that the MadIoT attacks can result in local power outages and in the worst cases, large-scale blackouts. Moreover, we show that these attacks can rather be used to increase the operating cost of the grid to benefit a few utilities in the electricity market. This work sheds light upon the interdependency between the vulnerability of the IoT and that of the other networks such as the power grid whose security requires attention from both the systems security and power engineering communities.

1 Introduction

A number of recent studies have revealed the vulnerabilities of the Internet of Things (IoT) to intruders [21, 49, 50]. These studies demonstrated that IoT devices from cameras to locks can be compromised either directly or through their designated mobile applications by an adversary [12, 28, 43]. However, most previous work has focused on the consequences of these vulnerabilities on personal privacy and security. It was not until recently and in the aftermath of the Distributed Denial of Service (DDoS) attack by the Mirai botnet, comprising six hundred thousand compromised devices targeting victim servers, that the collective effect of the IoT vulnerabilities was demonstrated [12]. In this paper, we reveal another substantial way that compromised IoT devices can be utilized by an adversary to disrupt one of the most essential modern infrastructure networks, the power grid.

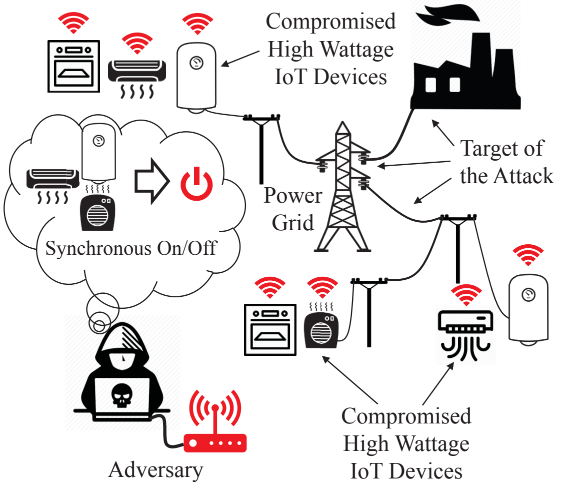

Figure 1: The MadIoT attack. An adversary can disrupt the power grid's normal operation by synchronously switching on/off compromised high wattage IoT devices.

Power grid security standards are all based on the assumption that the power demand can be predicted reliably on an hourly and daily basis [62]. Power grid operators typically assume that power consumers collectively behave similarly to how they did in the past and under similar conditions (e.g., time of the day, season, and weather). However, with the ubiquity of IoT devices and their poor security measures (as shown in [12]), we demonstrate that this is no longer a safe assumption.

There has been a recent trend in producing Wi-Fi enabled high wattage appliances such as air conditioners, water heaters, ovens, and space heaters that can now be controlled remotely and via the Internet [3] (for the power consumption of these devices see Table 1). Even older appliances can be remotely controlled by adding Wi-Fi enabled peripherals such as Tado° [8] and Aquanta [2]. A group of these devices can also be controlled remotely or automatically using smart thermostats or home assistants such as Amazon Echo [1] or Google Home [4]. Hence, once compromised, any of these devices can be used to control high wattage appliances remotely by an adversary to manipulate the power demand.

In this paper, we reveal a new class of potential attacks called the Manipulation of demand via IoT (MadIoT) attacks that *allow an adversary to disrupt the power grid's normal operation by manipulating the total power demand using compromised IoT devices* (see Fig. 1). These attacks, in the extreme case, can cause large scale blackouts. An important characteristic of MadIoT attacks is that unlike most of previous attacks on the power grid, they do not target the power grid's Supervisory Control And Data Acquisitions (SCADA) system but rather the loads that are much less protected as in load-altering attacks studied in [11, 41].

It is a common belief that manipulating the power demands can potentially damage the power grid. However, these speculations have mostly remained unexamined until our work. *We are among the first to reveal realistic mechanisms to cause abrupt distributed power demand changes using IoT devices–along with Dvorkin and Sang [24], and Dabrowski et al. [19]. Our key contribution is to rigorously study the effects of such attacks on the power grid from novel operational perspectives (for more details on the related work see Section 6).*

We study five variations of the MadIoT attacks and demonstrate their effectiveness on the operation of real-world power grid models via state-of-the-art simulators. These attacks can be categorized into three types:

(i) Attacks that result in frequency instability: An abrupt increase (similarly decrease) in the power demands–*potentially by synchronously switching on or off many high wattage IoT devices*–results in an imbalance between the supply and demand. This imbalance instantly results in a sudden drop in the system's frequency. If the imbalance is greater than the system's threshold, the frequency may reach a critical value that causes generators tripping and potentially a large-scale blackout. For example, using state-of-the-art simulators on the small-scale power grid model of the Western System Coordinating Council (WSCC), we show that a *30% increase in the demand results in tripping of all the generators. For such an attack, an adversary requires access to about 90 thousand air conditioners or 18 thousand electric water heaters within the targeted geographical area.* We also study the effect of such an attack during the system's restarting process after a blackout (a.k.a. the *black start*) and show that it can disrupt this process by causing frequency instability in the system.

(ii) Attacks that cause line failures and result in cascading failures: If the imbalance in the supply and demand after the attack is not significant, the frequency of the system is stabilized by the *primary controller* of the generators. Since the way power is transmitted in the power grid (a.k.a. the *power flows*) follows Kirchhoff's laws, the grid operator has almost no control over the power flows after the response of the primary controllers. Hence, even a small increase in the demands may result in line overloads and failures. These initial line failures may consequently result in further line failures or as it is called, a *cascading failure* [54]. For example, we show by simulations that *an increase of only* 1% *in the demand in the Polish grid during the Summer 2008 peak, results in a cascading failure with 263 line failures and outage in 86% of the loads. Such an attack by the adversary requires access to about 210 thousand air conditioners which is 1.5% of the total number of households in Poland [58].* During the Summer peak hours when most of the air conditioners are already on, decreasing their temperature set points [61] combined with the initiation of other high wattage appliances like water heaters, can result in the same total amount of increase in the demand.

Table 1: Home appliances' approximate electric power usage based on appliances manufactured by General Electric [3].

Appliance	**Power Usage (W)**
Air Conditioner	1,000
Space Heater	1,500
Air Purifier	200
Electric Water Heater	5,000
Electric Oven	4,000

We also show that an adversary can cause line failures by *redistributing the demand* via increasing the demand in some places (e.g., turning on appliances within a certain IP range) and decreasing the demand in others (e.g., turning off appliances within another IP range). These attacks, in particular, can cause failures in important high capacity *tie-lines* that connect two neighboring independent power systems–e.g., of neighboring countries.

(iii) Attacks that increase operating costs: When the demand goes above the day-ahead predicted value, conservatively assuming that there would be no frequency disturbances or line failures, the grid operator needs to purchase additional electric power from ancillary services (i.e., reserve generators). These reserve generators usually have higher prices than the generators committed as part of day ahead planning. Therefore, using the reserve generators can significantly increase the power generation cost for the grid operator but at the same time be profitable for the utility that operates the reserve generators. For example, we show by simulations that *a 5% increase in the power demand during peak hours by an adversary can result in a 20% increase in the power generation cost.* Hence, an adversary's attack may rather be for the benefit of a particular utility in the electricity market than for damaging the infrastructure.

The MadIoT attacks' sources are *hard to detect and disconnect* by the grid operator due to their distributed nature. These attacks can be *easily repeated* until being effective and are *black-box* since the attacker does not need to know the operational details of the power grid. These properties make countering the MadIoT attacks challenging. Nevertheless, we provide sketches of countermeasures against the MadIoT attacks from both the power grid and the IoT perspectives.

Overall, our work sheds light upon the interdependency between the vulnerability of the IoT and that of other networks such as the power grid whose security requires attention from both the systems security and the power engineering communities. We hope that our work serves to protect the grid against future threats from insecure IoT devices.

The rest of this paper is organized as follows. Section 2 provides a brief introduction to power systems. In Section 3, we introduce the MadIoT attack and its variations, and in Section 4, we demonstrate these attacks via simulations. In Section 5, we present countermeasure sketches against the MadIoT attacks. Section 6 presents a summary of the related work, and Section 7 discusses the limitations of our work. Finally Section 8 provides concluding remarks and recommendations. The central results of the paper are self-contained in the above sections. We refer the interested reader to the appendix for an overview of recent blackouts and their connection to MadIoT attacks, and additional experimental results.

2 Power Systems Background

In this section, we provide a brief introduction to power systems. For more details, refer to [26, 27, 31, 62].

2.1 Basics

Power systems consist of different components (see Fig. 2). The electric power is generated at power generators at different locations with different capacities and then transmitted via a high voltage *transmission network* to large industrial consumers or to the lower voltage *distribution network* of a town or a city. The power is then transmitted to commercial and residential consumers.

The main challenges in the operation and control of the power systems are in the transmission network. Moreover, since a distributed increase in power demand does not significantly affect the operation of the distribution network, we ignore the operational details of the distribution network and only consider it as an aggregated load within the transmission network. The term *power grid* mainly refers to the transmission network rather that the distribution network.

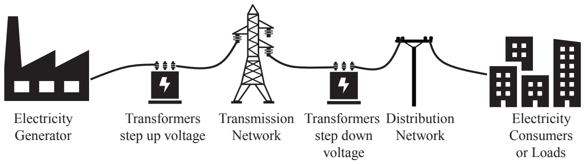
Electricity Generator
Transformers step up voltage
Transmission Network
Transformers step down voltage
Distribution Network
Electricity Consumers or Loads

Figure 2: Main components of a power system.

The transmission network can have a very complex topology. Each intersection point in the grid is called a *bus* which is a node in the equivalent graph.[1] Some of the buses may be connected to the distribution network of a city or a town and therefore represent the aggregated load within those places.

The instantaneous electric power generation and consumption are measured in watts (W) and are calculated based on electric voltages and currents. Almost all the power systems deploy Alternating Currents (AC) and voltages for transmitting electric power. This means that the electric current and voltage at each location and each point in time are equal to $I(t) = \sqrt{2}I_{\text{rms}}\cos(2\pi f t + \theta_I)$ and $V(t) = \sqrt{2}V_{\text{rms}}\cos(2\pi f t + \theta_V)$, in which f is the *nominal frequency* of the system, and $I_{\text{rms}}, V_{\text{rms}}$ and θ_I, θ_V are the root mean square (rms) values and the phase angles of the currents and voltages, respectively. In the U.S., Canada, Brazil, and Japan the power system frequency is $60Hz$ but almost everywhere else it is $50Hz$.

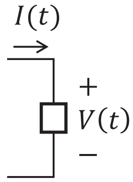

Figure 3

Given the voltages and the currents, the *active, reactive, and apparent power* amplitudes absorbed by a load can be computed as $P = V_{\text{rms}}I_{\text{rms}}\cos(\theta_V - \theta_I)$, $Q = V_{\text{rms}}I_{\text{rms}}\sin(\theta_V - \theta_I)$, and $S = V_{\text{rms}}I_{\text{rms}}$, respectively. $\cos(\theta_V - \theta_I)$ is called the *power factor* of a load.

2.2 Power Grid Operation and Control

Stable operation of the power grid relies on the persistent balance between the power supply and the demand. This is mainly due to the lack of practical large scale electrical power storage. In order to keep the balance between the power supply and the demand, power system operators use weather data as well as historical power consumption data to predict the power demand on a daily and hourly basis [27]. This allows the system operators to plan in advance and only deploy enough generators to meet the demand in the hours ahead without overloading any power lines. The grid operation should also comply with *the $N-1$ security standard*. The $N-1$ standard requires the grid to operate normally even after a failure in a *single* component of the grid (e.g., a generator, a line, or a transformer).

In power systems, the rotating speed of generators cor-

[1] The terms "bus" and "node" can be used interchangeably in this paper without loss of any critical information.

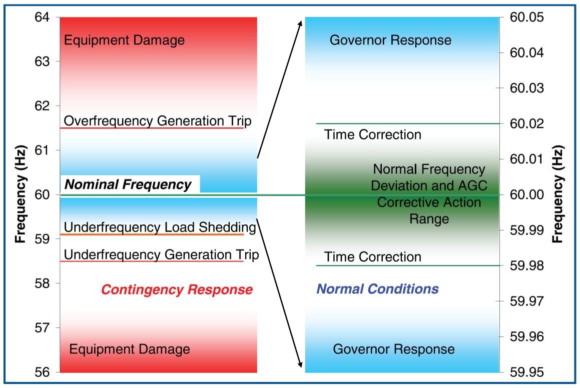

Figure 4: Normal and abnormal frequency ranges in North America. The figure is borrowed from [60].

respond to the frequency. When the demand gets greater than the supply, the rotating speeds of the turbine generators' rotors decelerate, and the kinetic energy of the rotors are released into the system in response to the extra demand. Correspondingly, this causes a drop in the system's frequency. This behavior of turbine generators corresponds to Newton's first law of motion and is calculated by the *inertia* of the generator. Similarly, the supply being greater than the demand results in acceleration of the generators' rotors and a rise in the system's frequency.

This decrease/increase in the frequency of the system cannot be tolerated for a long time since frequencies lower than the nominal value severely damage the generators. If the frequency goes above or below a threshold value, protection relays turn off or disconnect the generators completely (see Fig. 4 for normal and abnormal frequency ranges in North America). Hence, within seconds of the first signs of decrease in the frequency, the *primary controller* activates and increases the mechanical input which increases the speed of the generator's rotor and correspondingly the frequency of the system [26].

Despite stability of the system's frequency after the primary controller's response, it may not return to its nominal frequency (mainly due to the generators generating more than their nominal value). Hence, the *secondary controller* starts within minutes to restore the system's frequency. The secondary controller modifies the active power set points and deploys available extra generators and controllable demands to restore the nominal frequency and permanently stabilizes the system.

2.3 Power Flows

The equality of supply and demand is a necessary condition for the stable operation of the grid, but it is far from being sufficient. In order to deliver power from generators to loads, the electric power should be transmitted by the transmission lines. The power transmitted on each line in known as the *power flow* on that line.

Unlike routing in computer networks, power flows are almost entirely determined and governed by Kirchhoff's laws given the active and reactive power demand and supply values. Besides the constraints on the power flows enforced by Kirchhoff's laws, there are other limiting constraints that are dictated by the physical properties of the electrical equipment. In particular, each power line has a certain capacity of apparent power that it can carry safely.

Unlike water or gas pipelines, the capacity constraint on a power line is not automatically enforced by its physical properties. Once the power supply and demand values are set, the power flows on the lines are determined based on Kirchhoff's laws with no capacity constraints in the equations. Thus, an unpredicted supply and demand setting may result in electric power *overload* on some of the lines. Once a line is overloaded, it may be *tripped* by the protective relay, or it may break due to overheating–which should be avoided by the relay. Hence, the system operator needs to compute the power flows in advance–using the predicted demand values and optimal set of generators to supply the demand–to see if any of the lines will be overloaded. If so, the configuration of the generators should be changed to avoid lines overload and tripping.

2.4 Voltage Stability

Besides power line thermal limits, the power flows on the lines are limited by their terminating buses' voltages. The voltages at the buses are controlled by maintaining the level of the reactive power (Q) supply. Voltage instability or as it is called *voltage collapse* occurs when the generated reactive power becomes inadequate. This is mainly due to changes in system configurations due to line failures, increase in active or reactive power demand, or loss of generators. Voltage collapse should be studied using V-Q (characterizing the relationship between the voltage at the terminating bus of a line to the reactive power flow) and P-V (characterizing the relationship between the voltage at the terminating bus of a line to the active power flow) analysis which is beyond the scope of this paper, but for more details see [62, Chapter 7].

Voltage collapse results in the infeasibility of the power flow equations. Hence, it can be detected when the power flow solver fails to find a solution to the power flow equation (usually after an initial change in the system). In such scenarios, the grid operator is forced to perform load shedding (i.e., outage in part of the grid) in order to recover the system from a voltage collapse and make the power flow equations feasible again. Hence, even failures in a few lines or an increase in the active/reactive power demands may result in large scale outages around the grid due to voltage collapse.

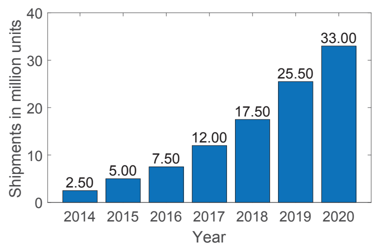

Figure 5: Estimated number of homes with smart thermostats in North America in millions. Data is obtained from Statista [56].

3 Attacking the Grid Using an IoT Botnet

In this section, we reveal attack mechanisms that can utilize an IoT botnet of high wattage devices to launch a large-scale coordinated attack on the power grid.

3.1 Threat Model

We assume that an adversary has already gained access to an IoT botnet of many high wattage smart appliances (listed in Table 1) within a city, a country, or a continent. Since most of the IoT devices are controlled using mobile phone applications, access to users' mobile phones or corresponding applications can also be used to control these devices [28]. This access can potentially allow the adversary to increase or decrease the demand in different locations remotely and synchronously. The adversary's power to manipulate the demand can also be translated into watts (W) using the numbers in Table 1 and based on the type and the number of devices to which it has access.

For example, if we consider only the houses with smart thermostats in 2018 as shown in Fig. 5 and assuming that each thermostat only controls two $1kW$ air conditioners, an attacker can *potentially* control $35GW$ of electric power[2]–even a fraction of which is a significant amount. Recall that in the case of the Mirai botnet, the attackers could get access to about 600 thousand devices within a few months [12].

The $35GW$ is computed by only considering the thermostats connected to a few air conditioners. By considering all the smart air conditioners as well as other high wattage appliances such as water heaters, this value would be much higher. Moreover, this amount will grow in the future as the trend shows in Fig. 5.

We call the attacks under this threat model the Manipulation of the demand via IoT (MadIoT) attacks. In the next subsection, we provide the details of various types of attacks that can be performed by an adversary.

3.2 MadIoT Attack Variations

MadIoT attacks can disrupt the normal operation of the power grid in many ways. Here, we present the most important and direct ways that such attacks can cause damage to the grid (summarized in Table 2):

1. Significant frequency drop/rise: As briefly described in Section 2, the normal operation of the power grid relies on the persistent balance between the supply and demand. Thus, an adversary's approach could be to disrupt this balance using an IoT botnet. An adversary can leverage an IoT botnet of high-wattage devices and synchronously switch on all the compromised devices. *If the resulting sudden increase in the demand is greater than a threshold, which depends on the inertia of the system, it can cause the system's frequency to drop significantly before the primary controllers can react. This consequently may result in the activation of the generators' protective relays and loss of generators, and finally a blackout.* Sudden decrease in the demand may also result in the same effect but this time by causing a sudden rise in the frequency.

An adversary can further increase its success by strategic selection of the timing of an attack using the online data available via the websites of Independent System Operators (ISOs)[3] (e.g., daily fuel mix and live updates of the demand values.) For example, we know that as the share of renewable resources in the power generation increases, the inertia of the system decreases. Therefore, an attack that is coordinated with the time that renewable penetration is highest, is more effective in causing large changes in the frequency. Similarly, *an attack during the peak hours can result in a slow yet persistent frequency drop in the system.* Such an attack may exhaust the controller reserves and force the system operator to perform load shedding. This *may result in power outages in several parts of the system if the situation is handled well by the operator, or in a large-scale blackout if it is mishandled and the system's frequency keeps dropping.* According to the European Network of Transmission System Operators for Electricity (ENTSOE) guidelines, if the frequency of the European grid goes below $47.5Hz$ or above $51.5Hz$, *a blackout can hardly be avoided [25].*

2. Disrupting a black start: Once there is a blackout, the grid operator needs to restart the system as soon as possible. This process is called a *black start*. Since the demand is unknown at the time of a black start, restarting the whole grid at the same time may result in frequency instability and system failure again. Hence, in a black start, the operator divides the system into smaller *islands* and tries to restart the grid in each island separately. The islands are then connected to increase the reliability of the system.

Since the grid is partitioned into smaller islands at

[2] For the sake of comparison, this amount is equal to 7% of the entire U.S. 2017 Winter peak demand (about $500GW$) [10].

[3] The system operators are given different names in different countries and continents, but here for the sake of simplicity, we refer to all of them as ISOs.

Table 2: MadIoT attack variations. The botnet size is in bots/MW which is the number of bots required to perform a successful variation of the MadIoT attack, if the total demand in the system is $1MW$. All the bots are assumed to be air conditioners.

#	Goal	Attack action	Initial impact	Botnet size	Simulation results
1	Grid frequency rise/drop	Synchronously switching on/off all the bots	Generation tripping	200–300	Figs. 8,7,9
2	Disrupting grid re-start	Synchronously switching on all the bots once the power restarts after a blackout	Generation tripping	100–200	Fig. 11
3	Line failures and cascades	Synchronously switching on or off the bots in different locations	Lines tripping	4–10	Figs. 12,13,15
4	Failure in tie-lines	Synchronously switching on (off) the bots in importing (exporting) end of a tie-line	Tie-lines tripping	10–15	Fig. 16
5	Increasing the operating cost	Slowly switching on the bots during power demand peak hours	Utilizing power generation reserve	30–50	Fig. 17

the time of a black start, the inertia of each part is low and therefore the system is very vulnerable to demand changes. Thus, an adversary can significantly hinder the black start process by suddenly increasing the demand using the IoT botnet once an island is up. This can cause a large frequency disturbance in each island and cause the grid to return to the blackout state.

3. Line failures and cascades: Recall from Section 2.3 that the power flows in power grids are determined by the Kirchhoff's laws. Therefore, most of the time, the grid operator does not have any control over the power flows from generators to loads. Once an adversary causes a sudden increase in the loads all around the grid, assuming that the frequency drop is not significant, the extra demand is satisfied by the primary controller. Since the power flows are not controlled by the grid operator at this stage, this may result in line overloads and consequent lines tripping.

After initial lines tripping or failures, the power flows carried by these lines are redistributed to other lines based on Kirchhoff's laws. Therefore, the initial line failures may subsequently result in further line failures or, as it is called, a *cascading failure* [54]. These failures may eventually result in the separation of the system into smaller unbalanced islands and a large-scale blackout.

Moreover, failure in a few lines accompanied by an increase in the power demand may result in a voltage collapse (recall from Section 2.4) which consequently would force the grid operator to perform load shedding. Hence, in some steps during the cascade, there are more outages due to load shedding.

An adversary may also start cascading line failures by redistributing the loads in the system by increasing the demand in a few locations and decreasing the demand in others in order to keep the total demand constant. This redistribution of the demand in the system may result in line failures without causing any frequency disturbances. The advantage of this attack is that it may have the same effect without attracting a lot of attention from the grid operator. It can be considered to be a *stealthier* version of the *demand increase only* attack.

4. Failures in the tie-lines: Tie-lines between the ISOs are among the most important lines within an interconnection. These tie-lines are usually used for carrying large amounts of power as part of an exchange program between two ISOs. Failure in one of these lines may result in a huge power deficit (usually more than $1GW$) in the receiving ISO and most likely a blackout due to the subsequent frequency disturbances or a large-scale outage due to load shedding by the grid operator.

Due to their importance, the tie-lines can be the target of an adversary. An adversary can observe the actual power flows on the tie-lines through ISOs' websites, and target the one that is carrying power flow near its capacity. In order to overload that line, all the adversary needs to do is to turn on the high wattage IoT devices in the area at the importing end of the line and turn off the ones at the exporting end (using the IP addresses of the devices).[4] This can overload the tie-line and cause it to trip by triggering its protective relay.

5. Increasing the operating cost: When the demand goes above the predicted value, the ISO needs to purchase additional electric power from ancillary services (i.e., reserve generators). These reserve generators usually have a higher price than the generators committed as part of the day ahead planning. Thus, using the reserve generators can significantly increase the power generation cost for the grid operator but at the same time be profitable for the utility that operates the reserve generator.

Hence, the goal of an adversary's attack may be to benefit a particular utility in the electricity market rather than to damage the infrastructure. The adversary can achieve this goal by slowly increasing the demand (e.g., switching on a few devices at a time) at a particular time of the day and in a certain location. Moreover, it may reach out

[4] A sudden increase in the demand, only at the importing end of the tie-line, may also result in its overload. This is due to the fact that once there is an imbalance between the supply and demand, all the generators within an interconnection (whether inside or outside of the particular ISO) respond to the imbalance which consequently results in an increase in the power flow on the tie-line.

to utilities to act in their favor in return for a payment.

Overall, the above attacks demonstrate that *an adversary as described in Section 3.1 has tremendous power to manipulate the operation of the grid in many ways which were not possible a few years ago in the absence of IoT devices.*

3.3 Properties and Defensive Challenges

The MadIoT attacks have unique properties that make them very effective and at the same time very hard to defend against. In this subsection, we briefly describe some of these properties.

First, the sources of the MadIoT attacks are *very hard to detect and disconnect* by the grid operator. The main reason is that the security breach is in the IoT devices, yet the attack is on the power grid. The grid operator cannot easily detect which houses are affected since it only sees the aggregation of the distributed changes in the demand around the grid. At the same time, the attack does not noticeably affect the performance of the IoT devices, especially if the smart thermostat is attacked. Moreover, the attack may not be noticeable by the households since the changes are temporary and can be considered as part of the automatic temperature control.

Second, the MadIoT attacks are *easy to repeat*. An adversary can easily repeat an attack at different times of the day and different days to find a time when the attack is the most effective. Moreover, this repeatability allows an adversary to cause a *persistent blackout* in the power grid by disrupting the black start process as described in the previous subsection.

Third, the MadIoT attacks are *black-box*. An adversary does not need to know the underlying topology or the detailed operational properties of the grid, albeit it can use the high-level information available on the ISOs' websites to improve the timing of its attack. It can also use the repeatability of these attacks and general properties of the power grids to achieve and perform a successful attack.

Finally, *power grids are not prepared to defend against the MadIoT attacks*, since abrupt changes in the demand are not part of the *contingency list* that grid operators are prepared for. As mentioned in Section 2, power grids are required to operate normally after a failure in a single component of the grid (the $N-1$ standard). Therefore, the daily operation of the grid is planned such that even a failure in the largest generator does not affect its normal operation.

The scenarios predicted by the $N-1$ standard, however, are quite different from the scenarios caused by the MadIoT attacks. Although an increase in the demand can be similar to losing a generator from the supply and demand balance perspective, these two phenomena result in completely different power flows in the grid. Hence, although losing a generator may not result in any issues as planned, increase in the demands by an adversary may result in many line overloads. Moreover, *the imbalance caused by an adversary may surpass the imbalance caused due to losing the largest generator*, and therefore results in unpredicted frequency disturbances. For example, the capacity of the largest operating generator in the system may be $1GW$ (usually a nuclear power plant) which can be surpassed by an attack comprising more than 100 thousand compromised water heaters.

Despite these difficulties, we provide sketches of countermeasures against the MadIoT attacks in Section 5.

3.4 Connection to Historical Blackouts

There have been several large-scale blackouts in the past two decades around the world. Although these events were not caused by any attacks, the chain of events that led to these blackouts could have been initiated by a MadIoT attack. For example, the initial reactive power deficit in Ohio in 2003 leading to the large-scale blackout in the U.S. and Canada [60], and the failures in the tie-lines connecting Italy to Switzerland in 2003 leading to the complete shutdown of the Italian grid [59], could have been caused by MadIoT attacks. Most of these events happened beacuse the systems' operators were *not prepared for the unexpected initial event*. Hence, the MadIoT attacks could result in similar unexpected failures. We reviewed a few of the recent blackouts in the power grids around the world and demonstrated how an adversary could have caused similar blackouts. The details of these events are relegated to Appendix A.

4 Experimental Demonstrations

In this section, we demonstrate the effectiveness of the MadIoT attacks on real-world power grid models via state-of-the-art simulators. Recall that the MadIoT attacks are black-box. Therefore, *the outcome of an attack highly depends on the operational properties of the targeted system at the time of the attack (e.g., generators' settings, amount of renewable resources, and power flows)*. We emphasize this in our simulations by changing the power grid models' parameters to reflect the daily changes in the operational properties of the system.

4.1 Simulations Setup

Our results are based on computer simulations. In particular, we use the MATPOWER [65] and the PowerWorld [7] simulators. MATPOWER is an open-source MATLAB library which is widely used for computing the power flows in power grids. PowerWorld, on the other hand, is an industrial-level software suite that is widely used by the industry for frequency stability analysis of power systems. We used the academic version of Power-

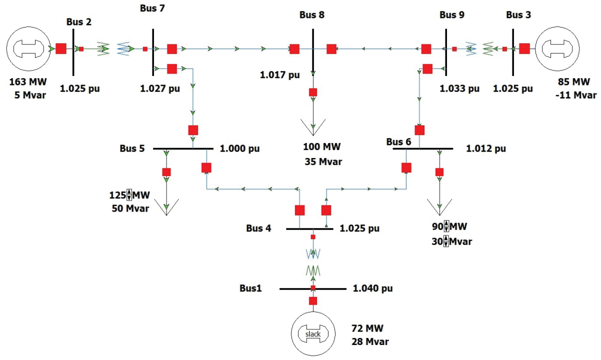

Figure 6: The WSCC 9-bus system. The generators at buses 2 and 3 are the buses with inertia, and the generator at bus 1 is a *slack bus* with no inertia. The slack bus is a bus in the system that can change its generation to make the power flow equations feasible. The load buses are buses 5, 6, and 8. We consider two operational settings of the WSCC system: (a) high inertia, in which both generators 2 and 3 have inertia constants (H) equal to $15s$, and (b) low inertia, in which generators 2 and 3 have inertia constants equal to $5s$ and $10s$, respectively [51, Chapter 3]. In all the simulations, the IEEE type-2 speed-governing model (IEEE-G2) is used for the generators [44].

World.

For frequency stability analysis in PowerWorld, to the best of our knowledge, there are no large-scale real-world power grids available for academic research. Hence, for evaluating the effects of the MadIoT attacks on the system's frequency, we use the WSCC 9-bus grid model that represents a simple approximation of the Western System Coordinating Council (WSCC)–with 9 buses, 9 lines, and $315MW$ of demand [35]. Despite its small size, due to the complexity of power systems transient analysis, it is widely used as a benchmark system [22, 48, 52].

For evaluating the effects of MadIoT attacks on the power flows, however, we use the Polish grid which is one of the largest and most detailed publicly available real-world power grids. To the best of our knowledge, there are no other real power grids at this scale and detail available for academic research.[5] We use the Polish grid data at its Summer 2004 peak–with 2736 buses, 3504 lines, and 18GW of demand–and at its Summer 2008 peak–with 3120 buses, 3693 lines, and 21GW of demand. Both are available through the MATPOWER library.

Since the total demand in the WSCC system is $315MW$, but the total demand in the Polish grid is about $20GW$, for comparison purposes, we focus on the percentage increase/decrease in the demand caused by an attack instead of the number of switching on/off bots. However, if we assume that all the bots are air conditioners, $1MW$ change in the demand corresponds simply to switching on/off 1,000 bots. Therefore, *we can define the* normalized botnet size *in bots/*MW *to be the number of bots required to perform a successful variation of the MadIoT attack, if the total demand in the system is* $1MW$. By this definition, it is easy to see that to increase the demand of any system by 1%, an adversary requires 10 bots/MW.

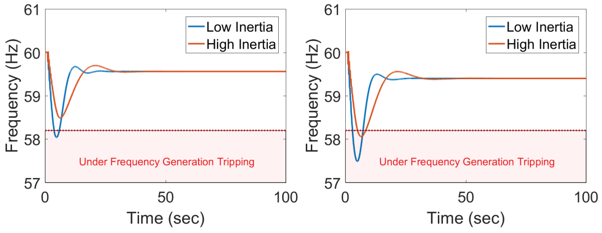

(a) (b)

Figure 7: Frequency disturbances due to unexpected demand *increases* in all the load buses in the WSCC system caused by an adversary, ignoring generators' frequency cut-off limit (shown by red dashed line). Increase by (a) $23MW$ and (b) $30MW$.

4.2 Frequency Disturbances

In this subsection, we evaluate the first two MadIoT attack variations described in Section 3.2. We consider two operational settings of the WSCC system: (a) high inertia and (b) low inertia (for details see Fig. 6).

4.2.1 200–300 Bots per MW Can Cause Sudden Generation Tripping

In order to show the frequency response of the system to sudden increases in the demand, we simulated the increase of (a) $23MW$ and (b) $30MW$ in all the loads for the high inertia and low inertia cases. These values can roughly be considered as 20% and 30% increases in the load buses, respectively. We similarly studied the frequency response of the system to sudden decreases of the demand. Figs. 7 and 8 present the results.

As mentioned in Section 2, the generators are protected from high and low frequency values by protective relays. These values depend on the type of a generator as well as the settings set by the grid operator. Here, we assume the safe frequency interval of $58.2Hz$ and $61.2Hz$ which is common in North America (see Fig. 4). Once a generator goes below or above these values, it gets disconnected from the grid by protective relays.

As can be seen in Figs. 7(b) and 8(b), sudden increase or decrease in the load buses by 30% or 20%, respectively, cause the system's frequency to go below or above the frequency cut-off limits. Hence, an adversary requires 200–300 bots/MW, or in this case 60–90 thousand bots, to perform these attacks.

As can be seen, however, the drop/rise in frequency is higher in the low inertia case (as predicted). Therefore, there are cases in which the frequency may go be-

[5] Topologies of other power grids may also be available through university libraries, but they are limited to the topology with no extra information on the operational details.

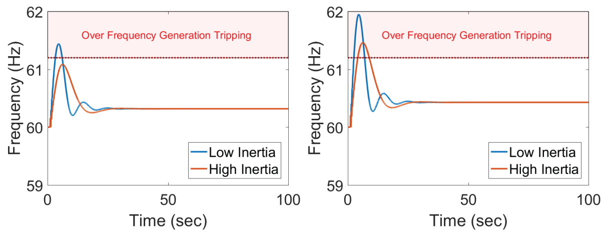

(a) (b)

Figure 8: Frequency disturbances due to unexpected demand *decreases* in all the load buses in the WSCC system by an adversary, ignoring generators' frequency cut-off limit (shown by red dashed line). Decrease by (a) $15MW$ and (b) $20MW$.

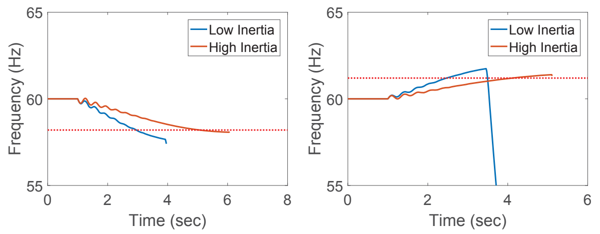

(a) (b)

Figure 9: Frequency disturbances due to unexpected demand changes in all the load buses in the WSCC system by an adversary, considering generators' frequency cut-off limits (shown by red dashed lines). (a) Demand increase of $30MW$ and (b) demand decrease of $20MW$.

low/above the critical frequency in the low inertia case but may remain in the safe interval in the high inertia case (see Figs. 7(a) and 8(a)). This suggests that *an attack that is not effective today, may be effective tomorrow* if the system's inertia is lower due to a higher rate of renewable generation.

In Figs. 7 and 8, the frequency cut-off limits of the generators are ignored. Hence, the generators are kept online even when the frequency goes beyond the safe operational limits. In reality, however, these generators are disconnected from the grid by the protective relays. Fig. 9 presents the frequency response of the system when the protective relays are enabled for the cases shown in Figs. 7(b) and 8(b). As can be seen, the grid completely shuts down and the simulations stop in less than 10 seconds due to disconnection of the generators.

Simulation results in this subsection demonstrate that the effectiveness of an attack in causing a critical frequency disturbance depends on the attack's scale as well as the system's total inertia at the time of the attack.

4.2.2 100–200 Bots per MW Can Disrupt the Grid Re-start

Once there is a blackout, the grid operator needs to restart the system as soon as possible (a.k.a. a black start). As mentioned in Section 3.2, due to frequency instability of the system at the black start, the restarting process is usually done by restarting the grid in parallel in disconnected islands and then reconnecting the islands.

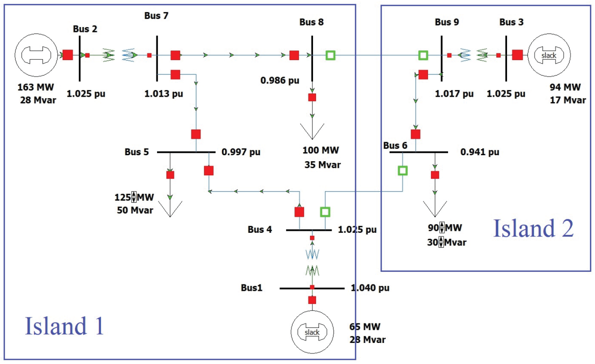

Figure 10: The WSCC 9-bus system during the black start.

Fig. 10 shows one way of partitioning the WSCC system into two islands. We assume that initially the grid operator could restart the two islands and stabilize the frequency at $60Hz$. Then, before the two islands are reconnected, an adversary increases the demand at all the load buses with the same amount (see Fig. 11).

The attack is performed at time 30 and the two islands are reconnected at time 50. As can be seen in Fig. 11(a), when there are no attacks, the two islands are reconnected with an initial small disturbance in the frequency and then the system reaches a stable state.

Fig. 11(b) shows the frequency of the system after $20MW$ increase in all the load buses at time 30. In this case, the frequency goes slightly below the minimum safe limit, but it is common in the black start process that the generators' lower (upper) frequency limits are set to lower (higher) levels than usual. Hence, the system may reach a stable state in this case as well.

As can be seen in Fig. 11(c), a $30MW$ increase in all the loads causes a large disturbance in the frequency, but as the two islands are reconnected the system's frequency is completely destabilized. These substantial deviations from safe frequency ranges can cause serious damage to the generators and are not permitted even in the black start process. Hence, in this case the system returns to the blackout stage. Even if the grid operator decides not to reconnect the two islands due to the frequency disturbances, Fig. 11(d) shows a significant drop in the second island's frequency that results in disconnection of the generators. Therefore, even if the big drop in frequency of island 1 ($1Hz$ below the safe limit) is acceptable during the black start, island 2 goes back to the blackout state.

For comparison purposes and to reflect on the role of the operational properties of the system on the outcome of an attack, we repeated the same set of simulations with different maximum power outputs for the generators' governors (see Fig. B.1 in the appendix). We observed that under the new settings, demand increases of

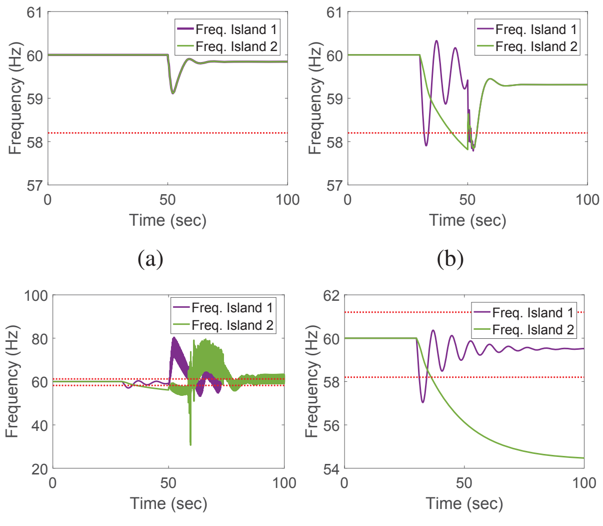

(a)

(b)

(c)

(d)

Figure 11: Frequency disturbances during the black start due to unexpected increases in all the load buses by an adversary, ignoring generators' frequency cut-off limits (shown by red dashed lines). (a) Normal black start in the absence of an adversary. (b) Demand increases of 20MW at the load buses before the reconnection of the two islands. (c) Demand increases of 30MW at the load buses before the reconnection of the two islands. (d) Demand increases of 30MW at the load buses without attempting to reconnect the two islands due to frequency instabilities.

up to 10MW results is a successful black start, unlike the previous case which could handle demand increases of 20MW at all the loads. Hence, an adversary requires at least 100–200 bots/MW, or in this case 30–60 thousand bots, to increase the demand at all the loads by 10–20% and disrupt the black start. Here again *we observe that the operational properties of the grid play an important role in the outcome of an attack.*

4.3 Line Failures and Cascades

In this subsection, we demonstrate the effectiveness of the third and the fourth variations of the MadIoT attacks described in Section 3.2. For simulating the cascading line failures, we use the MATLAB code developed by Cetinay et al. [18]. We had to slightly change the code to make it functional in the scenarios studied in this paper. To evaluate the severity of the cascade, we define *outage* as the percent of the demand affected by the power outage at the end of the cascade over the initial demand.

4.3.1 Only 10 Bots per MW Can Initiate a Cascading Failure Resulting in 86% Outage

As described in Section 3.2, once an adversary causes a sudden increase in the demand, if it does not result in a major frequency drop, the primary controllers at generators are automatically activated to compensate for the imbalance in the supply and demand. Despite balancing the supply and demand, since this balancing is unplanned, it may cause line overloads.

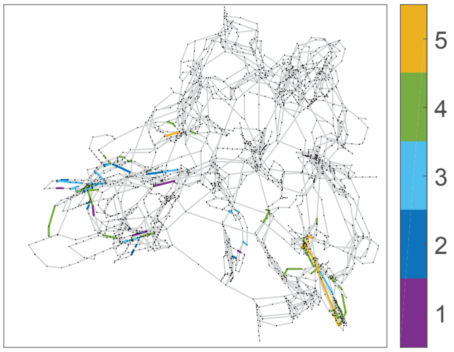

Figure 12: The cascading line failures initiated by a 1% increase in the demand in the Polish grid 2008 by an adversary (colors show the cascade step at which a line fails). It caused failures in 263 lines and 86% outage.

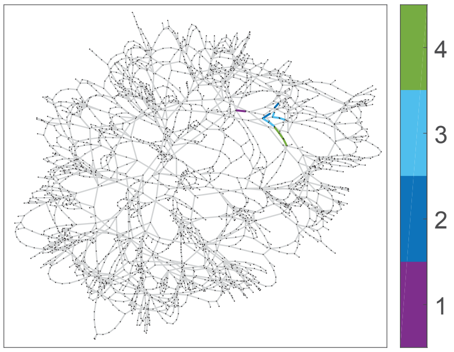

Figure 13: The cascading line failures initiated by a 10% increase in the demand in the Polish grid 2004 by an adversary (colors show the cascade step at which a line fails). It caused failures in 11 lines and 46% outage.

To demonstrate this, we assume that an adversary increases the demand at all the load buses by 1%. We also assume that all the generators contribute proportionally to their capacities to compensate for this sudden increase in the demand. This attack results in a single line failure in the Polish grid 2004 but no outages. However, as can be seen in Fig. 12, the same attack on the Polish grid 2008 results in the cascade of line failures that lasts for 5 rounds, causes 263 line failures, and 86% outage. The 1% increase in the total demand in the Polish grid 2008 is roughly equal to 210MW, requiring the adversary to access to 10 bots/MW which is about 210 thousand air conditioners in this case. This number is equal to 1.5% of the total number of households in Poland [58].

Since the Polish grid 2004 showed a good level of robustness against the 1% increase attack, we re-evaluated its robustness against a 10% increase in the demand. Fig. 13 shows the resulting line failures and the subsequent cascade caused by this attack. It can be seen that this attack causes much more damage with 11 line fail-

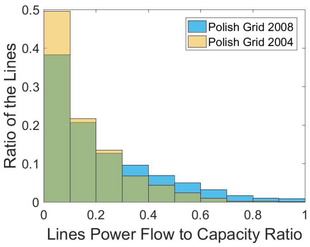

Figure 14: Histogram of the Polish grid lines' power flow to capacity ratio in Summer 2004 compared to Summer 2008.

ures and 46% outage. Despite the effectiveness of the second attack, the Polish grid 2004 shows greater level of robustness than the Polish grid 2008 even under a 10-time stronger attack. Although this may be due to many factors such as online generator locations and their values, topology of the grid, and even number of lines [54], one possible factor is *how initially saturated the power lines are*.

Fig. 14 presents the histogram of the Polish grid lines' power flow to capacity ratio in Summer 2004 compared to Summer 2008. There are about 10% more lines with flow to capacity ratio below 0.1 in the Polish grid 2004 compared to the Polish grid 2008. Consequently, there are more lines with power flow to capacity ratio greater than 0.3 in the Polish grid 2008 than in the Polish grid 2004 (to see the locations of the near saturated lines see Fig. B.2 in the appendix). This clearly demonstrates that a small increase in the demand is more likely to cause line overloads in the Polish grid 2008 than in the Polish grid 2004 (as observed in Figs.12 and 13).

Overall, as in the previous subsection, the results demonstrate that the effectiveness of an attack depends on the status of the grid at the time of the attack. However, *unlike the large botnet size (about 300 bots/MW) required to cause a blackout from frequency instability in the system, we observe here that even botnet size of 10 bots/MW can result in a significant blackout* depending on the grid's operational properties. Albeit the blackouts caused by frequency instabilities happen much faster (within seconds) than those caused by cascading line failures (within minutes or even hours).

4.3.2 Only 4 Bots per MW Can Initiate a Cascading Failure Resulting in 85% Outage by Redistributing the Demand

Another way of causing line failures and possibly cascading line failures in the grid is by redistributing the demand without increasing the total demand. As mentioned in Section 3.2, the advantage of this attack is that it may have a similar effect to the demand increase attack without attracting the grid operators' attention due to frequency disturbances.

Here, an adversary focuses only on the loads with demand greater than 20MW. This can be estimated by the adversary from the total number of IoT bots in a city or a town. The number of bots is correlated with the population of an area and therefore the total demand. Hence, an adversary detects these load buses and decreases or increases the demands by a random value such that the total demand increase and decrease sum up approximately to zero. We assume this can be done by randomly increasing or decreasing the demand by a Gaussian random variable with zero mean and selected standard deviation.

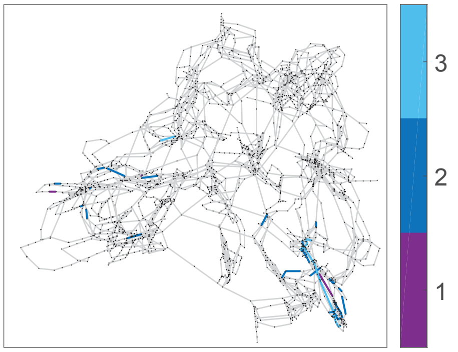

Figure 15: The cascading line failures initiated in the Polish grid 2008 by redistributing the demand by an adversary. Demand of the loads buses with demand greater than 20MW are changed by with a Gaussian distribution with zero mean and standard deviation 1MW (colors show the cascade step in which a line fails). It caused failures in 77 lines and 85% outage.

Again, the Polish grid 2004 showed a great level of robustness against these attacks. Even if an adversary decreases or increases the demand randomly by a Gaussian random variable with zero mean and standard deviation 10MW at loads with demand greater than 20MW, it only results in three line failures without any outages. However, the same attack with 10-time smaller changes, results in serious damage to the Polish grid 2008. As can be seen in Fig. 15, making only small changes with standard deviation of 1MW at load buses with demands greater 20MW results in cascading line failures with 77 line failures and outage of 85%. The total absolute value of the demand changes in this attack was about 80MW which means that *an adversary only requires 4 bots/MW, or in this case 80 thousand bots, to perform such an attack.*

Although these changes are made randomly, due to the stealthy nature of these attacks they can be repeated without attracting any attention until they are effective.

4.3.3 Only 15 Bots per MW Can Fail a Tie-line by Increasing (Decreasing) the Demand of the Importing (Exporting) ISOs

In order to demonstrate an attack on the tie-lines as described in Section 3.2, since we do not have access to the European grid or the U.S. Eastern Interconnection, we modified the Polish grid 2008 in a principled manner to

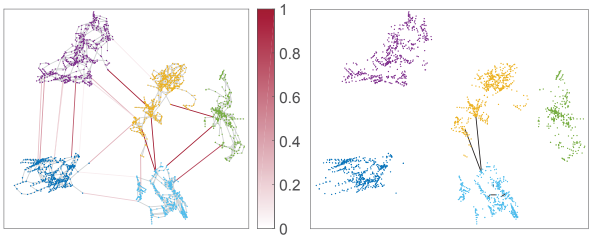

(a) (b)

Figure 16: Tie-line vulnerabilities in the partitioned Polish grid 2008. (a) The ratios of tie-lines' power flows to their nominal capacity. (b) Failures in the tie-lines between the yellow area and the light blue area caused by decreasing the demand by 1.5% in the former and increasing the demand by 1.5% in the latter by an adversary. Failed lines are shown in black.

represent a few neighboring ISOs in Europe connected by a few tie-lines.

First, we used a spectral clustering method to partition the Polish grid into 5 areas with a few connecting tie-lines. This is done using MATLAB's Community Detection Toolbox [34, 36]. Since the Polish grid does not inherently have 5 areas, however, the number of tie-lines between areas is slightly more than those of the European grid or Eastern Interconnection. Therefore, we removed one fifth of the tie-lines. In order to make the power flows feasible then, we reduced the total supply and demand by 60% and increased the capacity on the lines that were overloaded.

Fig. 16(a) shows the modified grid along with the ratios of tie-lines' power flows to their nominal capacities. As can be seen, similarly to the real grid operation, some of these tie-lines are carrying power flows near their capacities. These lines–*which can be detected through some of the ISOs' websites [5]*–are the most vulnerable to this variation of the MadIoT attacks.

For example, as can be seen in Fig. 16(a), the two lines that are connecting the yellow area to the light blue area are carrying power flows near their capacities. Therefore, increasing the demand in the light blue area and decreasing the demand in the yellow area (corresponding to the direction of the power flow on the lines) can potentially result in those lines tripping. It can be seen in Fig. 16(b) that a 1.5% decrease in the demand of the yellow area and a 1.5% increase of the demand in the light blue area by an adversary results in the failure of the two tie-lines (additional attacks on the other tie-lines are demonstrated in Figs. B.3(a) and B.3(b) in the appendix). Hence, an adversary can cause a failure in a tie-line by only a botnet of size 15 bots/MW, or in this case 60 thousand bots (30 thousand bots at each end of the tie-line).

Since the tie-lines usually carry substantial amounts of power, failure in these lines can result in cascade of line failures in other lines and eventually in disconnection of an ISO from the interconnection. Such a disconnection may result in a huge imbalance in the supply and demand values and in uncontrollable frequency drop leading to an inevitable blackout.

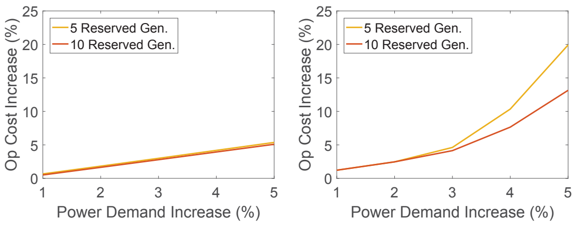

(a) (b)

Figure 17: Increase in the operating cost of the Polish grid 2004 by an adversary. The initial demand is 10% higher than the original demand during the Summer 2004 morning peak. (a) If the operating costs of the reserve generators are linear functions $c_1(x) = 100x$, and (b) if the operating costs of the reserve generators are quadratic functions $c_2(x) = 5x^2 + 100x$.

Attacks on the tie-lines are an effective approach when an adversary has a limited number of bots. By disconnecting an ISO from its neighboring ISOs, an adversary can cause a huge demand deficit in the targeted ISO and possibly a large-scale blackout.

4.4 Increasing the Operating Cost

In this final subsection, we evaluate the last variation of the MadIoT attacks described in Section 3.2. In this variation of the attacks, an adversary increases the demand not to necessarily cause a blackout, but rather to significantly increase the operating cost of the grid in favor of a utility in the electricity market.

4.4.1 50 Bots per MW Can Increase the Operating Cost by 20%

For these simulations, we use the Polish grid in Summer 2004. However, in order to mainly focus on the cost related issues, we increase the line capacities to make sure that the attack causes no line overloads. To simulate the system in its peak demand state, we increase the initial demand by 10% to make the demand before the attack close to the online generators' generation capacity.

We assume that the sudden increase in the demand caused by the attack can temporarily be handled by the primary controller and no large frequency drops as in Section 4.2 happen in any of the scenarios here. Therefore, our focus is on the cost of the required reserve generators for providing the additional power and returning the system's frequency back to 60Hz (or 50Hz).

We consider two cases, one with 5 reserve generators, and the other one with 10. We also consider two possible cost functions for the reserve generators: $c_1(x) = 100x$ and $c_2(x) = 5x^2 + 100x$, in which x is in MW and the $c_i(x)$s are in $\$/hr$. The linear and quadratic cost functions are the most common functions for approximating the generation costs [62, Chapter 3]. The $c_1(x)$ is selected similarly to cost function of the high-cost online generators in the grid before the attack and the $c_2(x)$ is selected to capture the start-up cost of the reserve generators as well as their higher cost compared to the online generators.

Fig. 17 shows the increase in the total cost given the two cost functions. As can be seen, in the worst-case scenario, a 5% increase in the demand–which requires 50 bots/MW, or in this case 1 million bots–can result in about a 20% increase in the operating cost of the grid (see the yellow line in Fig. 17(b)). This is four times higher than the best-case scenario (see the orange line in Fig. 17(a)) which is similar to the normal increase in the operating cost when no reserve generators are needed.

We observe that the effectiveness of the attack in increasing the cost depends on the total number of reserve generators as well as their generation cost functions.

5 Countermeasure Sketches

Although we are not aware of any rigorous countermeasures against the MadIoT attacks, in this section, we briefly provide a set of suggestions both in the power grid operation side and in the IoT design side to reduce the effectiveness of these attacks.

5.1 Power Grid Side

One of the most important properties of the MadIoT attacks, as mentioned in Section 3.3, is that grid operators, in general, are not prepared for these types of attacks. Hence, these types of attacks are not part of the contingency list of the power grid operators. Our first suggestion is for the grid operators to consider the MadIoT attacks in their contingency list and prepare for them. Such preparations can be directly incorporated into their already existing day-ahead planning tools to ensure that their systems have for example enough *inertia* (or *spinning reserve*) and the power lines have enough extra capacity to minimize the effects of a potential attacks. Although this might initially increase the grid operating cost, by developing more efficient planning tools and applying recent advances in designing *virtual inertia* for power systems [32], these costs can be reduced in the future. Thus, our suggestion for system operators is to push for more research in that direction in order to make their systems more robust to potential MadIoT attacks.

To minimize costs, the grid operators should also *have an accurate estimate of the total number of high wattage IoT devices in their system and accordingly the scale of a potential attack*, without being overprotective.

Since this is a new type of attack, enabled by the ubiquity of IoT devices, our last suggestion for the systems operators is *to revisit their online data and to find secure ways to release their data without revealing any critical information* that can be used by an adversary to improve the effectiveness of an attack.

5.2 IoT Side

The security challenges facing IoT devices are much more difficult to deal with. There are many ways an adversary can access a smart appliance. An adversary can directly get access to the device, or get access to the mobile phone, tablet, or a thermostat that controls that device, or with the ubiquity of digital home assistant devices such as Amazon Alexa or Google Home, an adversary can control smart appliances by getting access to these devices. Any of these devices can be a breaching point for an adversary. Hence, *coherent security measures are needed to protect almost all the devices within a home network against an adversary.*

Thus, in the IoT side, more research is required to study the vulnerability of IoT devices and networks, and to protect them against cyber attacks.

6 Related Work

The security and vulnerability of the IoT against cyber attacks has been widely studied [21,42,45,50,53,57,63]. In a recent study of the DDoS attack by the Mirai botnet [12], Antonakakis et al. showed that due to poor security measures in the IoT devices, such as easy to guess default passwords, an attacker could get access to about 600 thousand devices from cameras to DVRs and routers in a very short period. Similar studies had previously shown that Honeywell home controllers (including thermostats) could easily be compromised due to a pair of bugs in their authentication system [6]. It was also shown by Hernandez et al. that the lack of proper hardware protections in Nest thermostats allows attackers to install malicious software on these devices [33]. The vulnerability of Arduino Yun microcontrollers–used in some IoT devices–to cyber attacks was also revealed by Pastrana et al. [47].

In an interesting recent work [64], Zhang et al. demonstrated that home assistant devices can be controlled by an adversary using inaudible voice commands. In another recent work [49], Ronen et al. demonstrated that the smart lights within a city can potentially be compromised by creating a worm that can affect all the lamps using Zigbee. The security of mobile applications that control IoT devices has also been studied [28, 43]. In a comprehensive work [28], Fernandes et al. studied security of all Samsung-owned SmartThings apps and demonstrated that due to the security flaws in these applications, they could perform attacks like disabling vacation mode of a smart home. Naveed et al. also demonstrated that malicious apps on Android devices can freely *mis-bond* with any external IoT devices and control them [43].

Power systems' vulnerability to failures and attacks has been widely studied in the past few years [14, 17, 18, 23, 54]. In a recent work [29], Garcia et al. introduced Harvey, malware that affects power grid control systems and can execute malicious commands. Theoretical methods for detecting cyber attacks on power grids and recovering information after such attacks have also been developed [15, 20, 37, 39, 40, 55]. However, most of the previous work has focused on the attacks that directly target the power grid's physical infrastructure or its control system.

The interdependency between failures in power grids and communication networks, and their propagation has also been recently studied [16, 38, 46], but these works focused on attacks and failures that target both the power grid's and the communication network's physical infrastructure at the same time.

Load altering attacks on smart meters and large cloud servers has been first introduced by Mohsenian et al. [41]. Their work was mostly focused on the cost of protecting the grid against such attacks at loads. In contrast, we have analyzed the consequence of such attacks and introduced practical ways that they can be performed. Amini et al. [11] have also recently studied the effects of load altering attacks on the dynamics of the system and ways to use the system's frequency as feed-back to improve an attack. In two very recent papers, Dvorkin and Sang [24], and Dabrowski et al. [19] independently revealed the possibility of exploiting compromised IoT devices to disrupt normal operation of the power grid. Dvorkin and Sang [24] modeled their attack as an optimization problem for the attacker–with complete knowledge of the grid–to cause circuit breakers to trip in the distribution network. In contrast, we have focused on black-box attacks on transmission networks. Dabrowski et al. [19] studied the effect of demand increases caused by remotely activating CPUs, GPUs, hard disks, screen brightness, and printers on the frequency of the European power grid. *To the best of our knowledge, however, the work presented in this paper provides the most coherent and complete study on the effects of potential attacks on the power grid using high wattage IoT devices.*

There is another line of research that focuses on privacy of the customers in the presence of smart power meters which is beyond the scope of our paper [30].

7 Limitations and Future Work

In this work, we have analyzed the potential consequences of the MadIoT attacks on the operation of the power grid. However, our study has some limitations, and by addressing them one can provide a clearer picture of the threats facing the grid now and in the future. First, as mentioned in Section 4, we have only used publicly available data sets that may not exactly reflect the characteristics of all existing power grids. Therefore, the number of bots listed in Table 2 may not be enough to cause significant damage to all power grids. More detailed analysis of MadIoT attacks should be performed by system operators with access to the details of their systems.

Second, in our studies, we have not fully considered the existing control mechanisms for minimizing the subsequent effects of an initial failure (e.g., preventive load-shedding mechanisms). Hence, our cascading failures analysis may only reflect the worst case scenario.

Third, some of these high wattage IoT devices like air conditioners, have very large capacitors. Hence, it takes these devices 10 to 15 seconds to reach their maximum capacities. Therefore, it might be challenging to cause an abrupt increase in the demand and subsequently sudden drop in the frequency using these devices. Nevertheless, other smart devices like water heaters that are *resistive* loads can still be used for such purposes. Moreover, other varieties of the MadIoT attacks that do not require *synchronicity* on the scale of seconds (e.g., line failures) can still be performed using air conditioners.

Finally, unlike DDoS attacks, for the MadIoT attacks, the IoT bots should all be geographically located within boundaries of a power system. Hence, although the numbers of bots in Table 2 are achievable considering recent botnet sizes (e.g., the Mirai botnet), it might be much more challenging to reach these numbers within a targeted geographical location.

8 Conclusions

We have studied the collective effects of vulnerable high wattage IoT devices and have shown that once compromised, an adversary can utilize these devices to perform attacks on the power grid. We have revealed a new class of attacks on the power grid using an IoT botnet called Manipulation of demand via IoT (MadIoT) attacks. We have demonstrated via state-of-the-art simulators that these attacks can result in local outages as well as large-scale blackouts in the power grid depending on the scale of the attack as well as the operational properties of the grid. Moreover, we have shown that the MadIoT attacks can also be used to increase the operating cost of the grid to benefit a few utilities in the electricity market.

We hope that our work raises awareness of the significance of these attacks to grid operators, smart appliance manufacturers, and systems security experts in order to make the power grid (and other interdependent networks) more secure against cyber attacks. This is especially critical in the near future when more *smart* appliances with the ability to connect to the Internet are going to be manufactured. In particular, our work leads to following recommendations for the research community:

Power systems' operation: Power systems' operators should rigorously analyze the effects of potential MadIoT attacks on their systems and develop preventive methods to protect their systems. Initiating a data sharing platform between academia and industry may expedite these developments in the future.

IoT security: As shown by both presented MadIoT attacks and the Mirai botnet, insecure IoT devices can have devastating consequences that go far beyond individual security/privacy losses. This necessitates a rigorous pursuit of the security of IoT devices, including regulatory frameworks.

Interdependency: Our work demonstrates that interdependency between infrastructure networks may lead to hidden vulnerabilities. System designers and security analysts should explicitly study threats introduced by interdependent infrastructure networks such as water, gas, transportation, communication, power grid, and several other networks.

Acknowledgments

This work was supported in part by the U.S. National Science Foundation under Grants DMS-1736417, ECCS-1549881, and CNS-1553437, and Office of Naval Research YIP Award. We also thank anonymous reviewers for their helpful comments.

References

[1] Amazon Echo. https://www.amazon.com/all-new-amazon-echo-speaker-with-wifi-alexa-dark-charcoal/dp/B06XCM9LJ4. Accessed: Jan. 2018.

[2] Aquanta: Heat water when you need it, save money when you don't. https://aquanta.io/. Accessed: Jan. 2018.

[3] GE Wi-Fi connect appliances. http://www.geappliances.com/ge/connected-appliances/. Accessed: Jan. 2018.

[4] Google Home. https://store.google.com/product/google_home. Accessed: Jan. 2018.

[5] New York Independent System Operator (NYISO). http://www.nyiso.com/public/index.jsp. Accessed: Jan. 2018.

[6] Pair of bugs open Honeywell home controllers up to easy hacks. https://threatpost.com/pair-of-bugs-open-honeywell-home-controllers-up-to-easy-hacks/113965/. Accessed: Jan. 2018.

[7] PowerWorld Simulator. https://www.powerworld.com/. Accessed: Jan. 2018.

[8] Tado intelligent AC control. https://www.tado.com/us/. Accessed: Jan. 2018.

[9] The Federal Energy Regulatory Comission (FERC) and the North American Electric Reliability Corporation (NERC). Arizona-Southern California Outages on September 8, 2011. http://www.ferc.gov/legal/staff-reports/04-27-2012-ferc-nerc-report.pdf. Accessed: Jan. 2018.

[10] U.S. Energy Information Administration (EIA). https://www.eia.gov/. Accessed: Jan. 2018.

[11] Amini, S., Pasqualetti, F., and Mohsenian-Rad, H. Dynamic load altering attacks against power system stability: Attack models and protection schemes. *IEEE Trans. Smart Grid 9*, 4 (2018), 2862–2872.

[12] Antonakakis, M., April, T., Bailey, M., Bernhard, M., Bursztein, E., Cochran, J., Durumeric, Z., Halderman, J. A., Invernizzi, L., Kallitsis, M., et al. Understanding the Mirai botnet. In *Proc. USENIX Security Sympsion'17* (Aug. 2017).

[13] Australian Energy Market Operator (AEMO). Black system South Australia 28 september 2016. https://www.aemo.com.au/-/media/Files/Electricity/NEM/Market_Notices_and_Events/Power_System_Incident_Reports/2017/Integrated-Final-Report-SA-Black-System-28-September-2016.pdf. Accessed: Jan. 2018.

[14] Bienstock, D. *Electrical Transmission System Cascades and Vulnerability: An Operations Research Viewpoint*. SIAM, 2016.

[15] Bienstock, D., and Escobar, M. Computing undetectable attacks on power grids. *ACM PER 45*, 2 (2017), 115–118.

[16] Buldyrev, S., Parshani, R., Paul, G., Stanley, H., and Havlin, S. Catastrophic cascade of failures in interdependent networks. *Nature 464*, 7291 (2010), 1025–1028.

[17] Carreras, B., Lynch, V., Dobson, I., and Newman, D. Critical points and transitions in an electric power transmission model for cascading failure blackouts. *Chaos 12*, 4 (2002), 985–994.

[18] Cetinay, H., Soltan, S., Kuipers, F. A., Zussman, G., and Van Mieghem, P. Analyzing cascading failures in power grids under the AC and DC power flow models. In *Proc. IFIP Performance'17* (Nov. 2017).

[19] Dabrowski, A., Ullrich, J., and Weippl, E. R. Grid shock: Coordinated load-changing attacks on power grids: The non-smart power grid is vulnerable to cyber attacks as well. In *Proc. ACM ACSAC'17* (Dec. 2017).

[20] Dán, G., and Sandberg, H. Stealth attacks and protection schemes for state estimators in power systems. In *Proc. IEEE SmartGridComm'10* (2010).

[21] Denning, T., Kohno, T., and Levy, H. M. Computer security and the modern home. *Commun. ACM 56*, 1 (2013), 94–103.

[22] Dobakhshari, A. S., and Ranjbar, A. M. A novel method for fault location of transmission lines by wide-area voltage measurements considering measurement errors. *IEEE Trans. Smart Grid 6*, 2 (2015), 874–884.

[23] Dobson, I. Cascading network failure in power grid blackouts. *Encyclopedia of Systems and Control* (2015), 105–108.

[24] Dvorkin, Y., and Garg, S. IoT-enabled distributed cyber-attacks on transmission and distribution grids. In *Proc. NAPS'17* (Sept 2017).

[25] European Network of Transmission System Operators for Electricity (ENTSOE). Frequency stability evaluation criteria for the synchronous zone of continental Europe. https://www.entsoe.eu/Documents/SOC%20documents/RGCE_SPD_frequency_stability_criteria_v10.pdf. Accessed: Jan. 2018.

[26] European Network of Transmission System Operators for Electricity (ENTSOE). Continental Europe operation handbook, 2004. https://www.entsoe.eu/publications/system-operations-reports/operation-handbook/Pages/default.aspx. Accessed: Jan. 2018.

[27] Federal Energy Regulatory Commission and others. *Energy Primer, a Handbook of Energy Market Basics*. 2012.

[28] Fernandes, E., Jung, J., and Prakash, A. Security analysis of emerging smart home applications. In *Proc. IEEE S&P'16* (2016), pp. 636–654.

[29] Garcia, L., Brasser, F., Cintuglu, M. H., Sadeghi, A.-R., Mohammed, O., and Zonouz, S. A. Hey, my malware knows physics! attacking PLCs with physical model aware rootkit. In *Proc. NDSS'17* (2017).

[30] GIACONI, G., GÜNDÜZ, D., AND POOR, H. V. Privacy-aware smart metering: Progress and challenges. *IEEE Signal Process. Mag. (to appear)* (2018).

[31] GLOVER, J. D., SARMA, M. S., AND OVERBYE, T. *Power System Analysis & Design, SI Version*. Cengage Learning, 2012.

[32] GROSS, D., BOLOGNANI, S., POOLLA, B. K., AND DÖRFLER, F. Increasing the resilience of low-inertia power systems by virtual inertia and damping. In *Proc. IEEE IREP'17* (2017).

[33] HERNANDEZ, G., ARIAS, O., BUENTELLO, D., AND JIN, Y. Smart nest thermostat: A smart spy in your home. *Black Hat USA* (2014).

[34] HESPANHA, J. P. An efficient Matlab algorithm for graph partitioning. *Technical Report* (2004). https://www.ece.ucsb.edu/~hespanha/published/tr-ell-gp.pdf. Accessed: Jan. 2018.

[35] ILLINOIS CENTER FOR A SMARTER ELECTRIC GRID (ICSEG). Power test cases. http://icseg.iti.illinois.edu/power-cases/. Accessed: Jan. 2018.

[36] KEHAGIAS, A. Community detection toolbox. https://www.mathworks.com/matlabcentral/fileexchange/45867-community-detection-toolbox. Accessed: Jan. 2018.

[37] KIM, J., TONG, L., AND THOMAS, R. J. Subspace methods for data attack on state estimation: A data driven approach. *IEEE Trans. Signal Process. 63*, 5 (2015), 1102–1114.

[38] KORKALI, M., VENEMAN, J. G., TIVNAN, B. F., BAGROW, J. P., AND HINES, P. D. Reducing cascading failure risk by increasing infrastructure network interdependence. *Sci. Rep. 7* (2017).

[39] LI, S., YILMAZ, Y., AND WANG, X. Quickest detection of false data injection attack in wide-area smart grids. *IEEE Trans. Smart Grid 6*, 6 (2015), 2725–2735.

[40] LIU, Y., NING, P., AND REITER, M. K. False data injection attacks against state estimation in electric power grids. *ACM Trans. Inf. Syst. Secur. 14*, 1 (2011), 13.

[41] MOHSENIAN-RAD, A.-H., AND LEON-GARCIA, A. Distributed internet-based load altering attacks against smart power grids. *IEEE Trans. Smart Grid 2*, 4 (2011), 667–674.

[42] NAEINI, P. E., BHAGAVATULA, S., HABIB, H., DEGELING, M., BAUER, L., CRANOR, L., AND SADEH, N. Privacy expectations and preferences in an IoT world. In *Proc. SOUPS'17* (2017).

[43] NAVEED, M., ZHOU, X.-Y., DEMETRIOU, S., WANG, X., AND GUNTER, C. A. Inside job: Understanding and mitigating the threat of external device mis-binding on android. In *Proc. NDSS'14* (2014).

[44] NEPLAN-POWER SYSTEMS ANALYSIS. Turbine-governor models. http://www.neplan.ch/wp-content/uploads/2015/08/Nep_TURBINES_GOV.pdf. Accessed: Jan. 2018.

[45] NIA, A. M., AND JHA, N. K. A comprehensive study of security of internet-of-things. *IEEE Trans. Emerg. Topics Comput. 5*, 4 (2017), 586–602.

[46] PARANDEHGHEIBI, M., AND MODIANO, E. Robustness of interdependent networks: The case of communication networks and the power grid. In *Proc. IEEE GLOBECOM'13* (2013).

[47] PASTRANA, S., RODRIGUEZ-CANSECO, J., AND CALLEJA, A. ArduWorm: A functional malware targeting Arduino devices. *COSEC Computer Security Lab* (2016).

[48] RAMIREZ, L., AND DOBSON, I. Monitoring voltage collapse margin with synchrophasors across transmission corridors with multiple lines and multiple contingencies. In *Proc. IEEE PES-GM'15* (2015).

[49] RONEN, E., SHAMIR, A., WEINGARTEN, A.-O., AND O'FLYNN, C. IoT goes nuclear: Creating a ZigBee chain reaction. In *Proc. IEEE S&P'17* (2017).

[50] SACHIDANANDA, V., TOH, J., SIBONI, S., SHABTAI, A., AND ELOVICI, Y. Poster: Towards exposing internet of things: A roadmap. In *Proc. ACM CCS'16* (2016).

[51] SAUER, P., AND PAI, M. *Power System Dynamics and Stability*. Prentice Hall, 1998.

[52] SHARMA, A., SRIVASTAVA, S., AND CHAKRABARTI, S. Testing and validation of power system dynamic state estimators using real time digital simulator (RTDS). *IEEE Trans. Power Syst. 31*, 3 (2016), 2338–2347.

[53] SIMPSON, A. K., ROESNER, F., AND KOHNO, T. Securing vulnerable home IoT devices with an in-hub security manager. In *Proc. IEEE PerCom'17* (2017).

[54] SOLTAN, S., MAZAURIC, D., AND ZUSSMAN, G. Analysis of failures in power grids. *IEEE Trans. Control Netw. Syst. 4*, 3 (2017), 288–300.

[55] SOLTAN, S., YANNAKAKIS, M., AND ZUSSMAN, G. Joint cyber and physical attacks on power grids: Graph theoretical approaches for information recovery. In *Proc. ACM SIGMETRICS'15* (June 2015).

[56] STATISTA. Number of homes with smart thermostats in North America from 2014 to 2020 (in millions). https://www.statista.com/statistics/625868/homes-with-smart-thermostats-in-north-america/. Accessed: Jan. 2018.

[57] SURBATOVICH, M., ALJURAIDAN, J., BAUER, L., DAS, A., AND JIA, L. Some recipes can do more than spoil your appetite: Analyzing the security and privacy risks of ifttt recipes. In *Proc. WWW'17* (2017).

[58] THE UNITED NATIONS. Demographic yearbook, 2017. https://unstats.un.org/unsd/demographic-social/products/dyb/dybcensusdata.cshtml. Accessed: Jan. 2018.

[59] UNION FOR THE COORDINATION OF THE TRANSMISSION OF ELECTRICITY (UCTE). Final report of the investigation committee on the 28 September 2003 blackout in Italy. http://www.rae.gr/old/cases/C13/italy/UCTE_rept.pdf. Accessed: Jan. 2018.

[60] U.S.-CANADA POWER SYSTEM OUTAGE TASK FORCE. Report on the August 14, 2003 blackout in the United States and Canada: Causes and recommendations. https://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/BlackoutFinal-Web.pdf. Accessed: Jan. 2018.

[61] WANG, N., ZHANG, J., AND XIA, X. Energy consumption of air conditioners at different temperature set points. *Energy and Buildings 65* (2013), 412–418.

[62] WOOD, A. J., AND WOLLENBERG, B. F. *Power Generation, Operation, and Control*. John Wiley & Sons, 2012.

[63] YU, T., SEKAR, V., SESHAN, S., AGARWAL, Y., AND XU, C. Handling a trillion (unfixable) flaws on a billion devices: Rethinking network security for the internet-of-things. In *Proc. ACM HotNets'15* (2015).

[64] ZHANG, G., YAN, C., JI, X., ZHANG, T., ZHANG, T., AND XU, W. DolphinAttack: Inaudible voice commands. In *Proc. ACM CCS'17* (2017).

[65] ZIMMERMAN, R. D., MURILLO-SÁNCHEZ, C. E., AND THOMAS, R. J. MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Trans. Power Syst. 26*, 1 (2011), 12–19.

Appendix

A Historical Blackouts Details

In this appendix, we briefly review a few of the recent blackouts in the power grids around the world to further demonstrate the potential effectiveness of the MadIoT attacks.

A.1 The 2003 Blackout in the U.S. and Canada

The August 14, 2003, blackout in the U.S. and Canada is one of the largest blackouts in history. It affected an area with an estimated 50 million people and $61,800MW$ of power in the states of Ohio, Michigan, Pennsylvania, New York, Vermont, Massachusetts, Connecticut, New Jersey and the Canadian province of Ontario. According to the aftermath report [60], the failure started with a generator failure in Ohio due to an underpredicted reactive load to serve high air conditioning demand. After the initial failure, the Ohio grid operators were forced to import power which caused more line failures due to overloads and lines touching nearby trees. Within hours, the line failures cascaded and caused failure in major tie-lines between ISOs. This resulted in disconnection of the Eastern interconnection into East and West parts which caused further frequency and voltage instabilities and a large-scale blackout. The details of the events leading to the blackout can be found in [60].

How an adversary could have initiated a similar scenario? In a relatively hot summer day (but not the hottest day), an adversary could have initiated the same event by overloading the Ohio system by increasing the reactive power demand by remotely starting several air conditioners. This could cause an unexpected shortage in reactive power generation and possibly the same generator failure and consequent voltage collapse events.

A.2 The 2003 Blackout in Italy

The September 28, 2003, blackout was the most serious blackout in Italy and caused an outage almost everywhere in Italy. At around 3pm in the afternoon, Italy was importing $3,610MW$ and $2,212MW$ of power from Switzerland and France, about $600MW$ and $400MW$ above their scheduled exchange agreements, respectively. At this time, one of the tie-lines between Switzerland and Italy tripped due to an overload and touching a tree. This resulted in an overload in another tie-line between the two countries and tripping of the second line. After, the second line failure, further lines between Italy and France, Austria, and Slovenia tripped due to overloads and caused the Italian grid to be disconnected from the continental European grid. This resulted in a huge imbalance between supply and demand within Italy and a frequency drop that could not be recovered despite further aggressive load shedding. The details of the events leading to this blackout can be found in [59].

How an adversary could have initiated a similar scenario? An adversary could actively monitor the power flow on the tie-lines through European grids' websites and overload the tie-lines by increasing power demand in Italy and possibly decreasing power demand in Switzerland or France. This could have resulted in the failure of the same tie-lines and subsequent failures.

A.3 The 2011 Blackout in Arizona-Southern California

The September 8, 2011, Arizona-Southern California affected approximately 2.7 million people. It started with a single high voltage line failure due to a fault which redistributed power towards the San Diego area on *a hot day during hours of peak demand*. Within minutes this redistribution of power resulted in more line and transformer failures (which are modeled as line failures in simulations in the previous section) and eventually separation of the San Diego area from rest of the Western Interconnection. This separation resulted in a huge imbalance between the supply and demand in the San Diego area and a frequency drop which caused generation tripping and a blackout. The details of the events can be found in [9].

How an adversary could have initiated a similar scenario? An adversary could have caused the same initial line failure (which was operating within 78% of its capacity) by increasing the demand in the San Diego area and possibly reducing the demand in Arizona.

A.4 The 2016 Blackout in South Australia

The September 28, 2016, blackout in South Australia affected approximately 1 million customers. Extreme weather conditions on September 28 caused failure in three transmission lines. Following these failures, there was a $456MW$ reduction in wind generation in the South Australia grid which resulted in an increase in imported power and further tripping of the tie-lines. As a result, the South Australia grid was separated from rest of the Australian grid. This resulted in $900MW$ imbalance is supply and demand, and a sudden drop in the frequency which caused a blackout in the system. The details of these events can be found in [13].

What is special about this blackout is that a big portion of the electric power in South Australia in generated by wind turbines and solar panels (about 75%) which have very low inertia. This is the main reason for the very quick drop in the frequency after the separation of the South Australian grid from the rest of the interconnection, without the grid operator having a chance to respond to the imbalance by load shedding. This event, in particular, shows that in places or times that renewable resources have a higher share of the power generation, the grid is much more vulnerable to the MadIoT attacks that cause sudden increases in the demand.

How an adversary could have initiated a similar scenario? Due to the low inertia of the South Australian grid, the sudden increase in the demand by an adversary in the area should be compensated by the tie-lines. This, depending on the amount of the increase, can potentially result in the overload of the tie-lines and their failure. Once they fail and the system is islanded, it may collapse because of the supply and demand imbalance and a quick frequency drop.

B Extra Simulations and Details

In this appendix, we present supplemental simulation results.

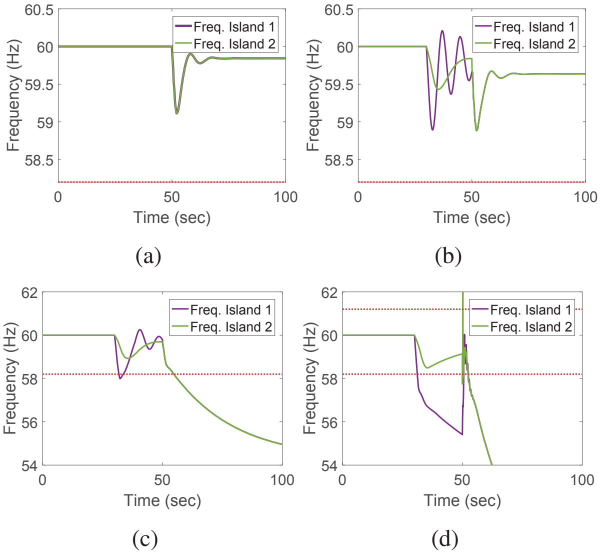

(a)

(b)

(c)

(d)

Figure B.1: Frequency disturbances during the black start due to unexpected increases in all the load buses by an adversary (as described in Section 4.2.2), ignoring generators' frequency cutoff limits (shown by red dashed lines). The maximum power outputs for the generators' governors are different in this figure from that of the generators in Fig. 11. (a) Normal black start operation in the absence of an adversary. (b) Demand increases of 10MW at the load buses before the reconnection of the two islands. (c) Demand increases of 20MW at the load buses before the reconnection of the two islands. (d) Demand increases of 30MW at the load buses before the reconnection of the two islands.

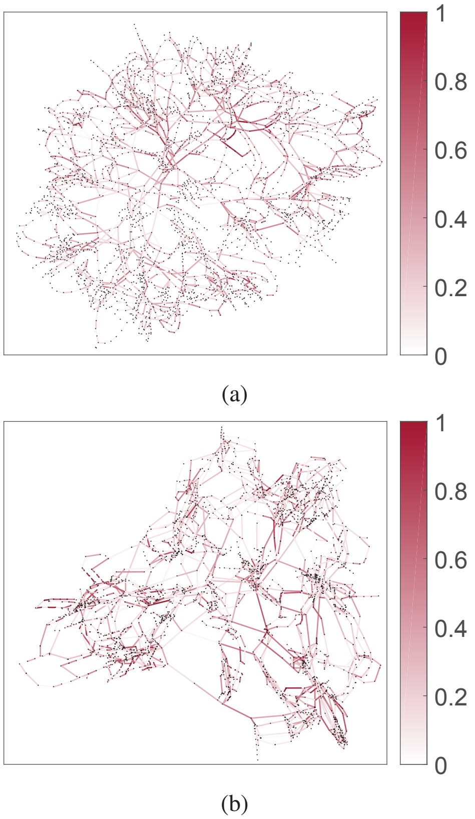

(a)

(b)

Figure B.2: Polish grid lines' power flow to capacity ratio in (a) Summer 2004 and (b) Summer 2008.

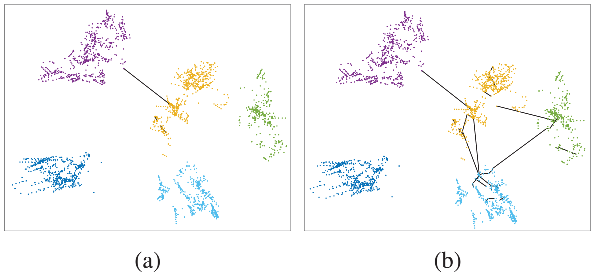

(a)

(b)

Figure B.3: Tie-line vulnerabilities in the partitioned Polish grid 2008. (a) Failures in the tie-lines between the yellow area and the purple area caused by decreasing the demand by 1% in the former and increasing the demand by 1% in the latter. All the failed lines are shown in black. (b) Failures in several tie-lines caused by decreasing the demand by 1% in the yellow area and increasing the demand by 0.3% in the purple, dark blue, and light blue areas. All the failed lines are shown in black.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TNSE.2019.2922131, IEEE Transactions on Network Science and Engineering

Protecting the Grid against MAD Attacks

Saleh Soltan, *Member, IEEE*, Prateek Mittal, *Senior Member, IEEE*, H. Vincent Poor, *Fellow, IEEE*

Abstract—Power grids have just recently been shown to be vulnerable to MAnipulation of Demand (MAD) attacks using high-wattage IoT devices. In this paper, we introduce two forms of defenses against line failures caused by these attacks: (1) we develop two algorithms named SAFE and IMMUNE for finding efficient operating points for generators during the normal operation of the grid such that no lines are overloaded instantly after any potential MAD attacks, and (2) assuming lines can temporarily tolerate overloads, we develop efficient methods to verify in advance if such overloads can quickly be cleared by changing the operating points of the generators after any attacks. We then define the novel notion of αD-robustness for a grid indicating that line overloads can either be prevented or cleared after any attacks based on the two forms of introduced defenses if an adversary can increase/decrease the demands by at most α fraction. We demonstrate that practical upper and lower bounds on the maximum α for which a grid is αD-robust can be found efficiently in polynomial time. Finally, we evaluate the performance of the developed algorithms and methods on realistic power grid test cases.

Index Terms—Power grid, IoT, cyber attacks, demand manipulation, control

1 Introduction

Power grids, as one of the most essential infrastructure networks, have been repeatedly shown in the past couple of years to be vulnerable to cyber attacks. The most infamous example of these attacks was on Ukrainian grid that affected about 225,000 people in December 2015 [1]. However, smaller scale attacks on regional power grids have been shown in a recent report to be more common and pervasive [2]. As indicated in the report, *"Hackers are developing a penchant for attacks on energy infrastructure because of the impact the sector has on people's lives"* [2].

Because of this ever-growing threat, there has been a significant effort by researchers in recent years to protect the grid against cyber attacks. These efforts have been mainly focused on potential attacks that directly affect different components of power grids' Supervisory Control And Data Acquisition (SCADA) systems. Many system operators prefer to completely disconnect their SCADA systems from the Internet in the hope that their systems remain unreachable to hackers.

Despite these efforts, the *power demand* side of the grid operation which is not controlled by SCADA has been neglected to be directly susceptible to attacks in security assessments due to their predictable nature. However, as we [3] and Dabrowski et al. [4] have recently revealed, the universality and growth in the number of high-wattage Internet of Things (IoT) devices, such as air conditioners and water heaters, have provided a unique way for adversaries to *disrupt the normal operation of power grid, without any access to the SCADA system [5], [6].* In particular, an adversary with access to sufficiently many of such high-wattage devices (i.e., a *botnet*), can abruptly increase or decrease the total demand in the system by synchronously turning these devices on or off, respectively. We call these attacks MAnipulation of Demand (MAD) attacks (see Fig. 1).

An abrupt increase/decrease in the total demand results in abrupt drop/rise in the system's frequency. If this drop/rise is significant, generators will be automatically disconnected from the grid and a large scale blackout occurs within seconds [3], [4]. If the drop/rise in the frequency is not significant, the extra demand/generation can automatically be compensated by generators' primary controllers, and the frequency of the system will be stabilized. As a result of this automatic change in the generation–and demand by the adversary–the power flows in the transmission network change based on power flow equations. Since the power flows are not controlled by the grid operator at this stage, this change in the power flows may result in line overloads and consequent line-trippings. These initial line failures can initiate a cascading line failure and result in a large scale blackout in the grid [3]. For example, it has been demonstrated that only a 1% increase in the demands at certain scenarios may initiate a cascading failure leading to 86% power outage in the system.

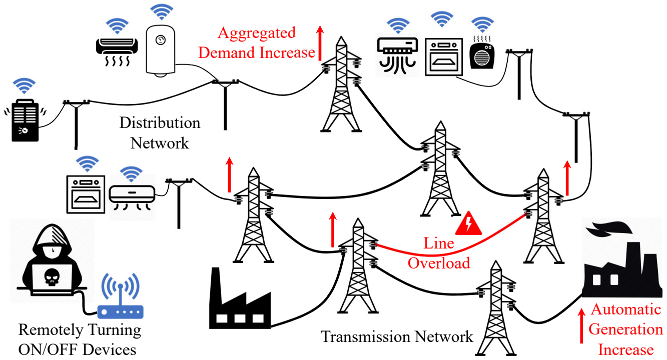

Fig. 1: The MAD attack. An adversary with access to an IoT botnet of high-wattage devices can remotely and synchronously switch on/off these devices in order to change power flows on the lines in power transmission network and cause line overloads and failures.

The grid operator can protect the grid against initial drop/rise in the system's frequency caused by a MAD attack by ensuring that the system has enough *inertia* (mostly through rotating generators) and there is enough available *spinning reserve* (i.e., generators have enough extra gener-

Authors are with the Department of Electrical Engineering at Princeton University, Princeton, NJ. Emails: {`ssoltan, pmittal, poor`*}*`@princeton.edu`*.*

ation capacity) [3]. However, protecting the grid against possible line overloads and failures after a MAD attack, which is the main focus of this paper, is more analytically and computationally challenging. Such defenses require the grid operator to analyze all possible MAD attacks and their consequences on the power flows and select operating points for the generators (i.e., their power generation output) to satisfy the power demands such that no lines are overloaded after any MAD attacks.

We first focus on finding operating points (namely *robust operating points*) with the minimum cost for the generators such that no lines are overloaded after the automatic primary response of the generators to any MAD attacks. Since changes in power flows after a MAD attack directly depend on generators' operating points, finding the optimal operating points for the generators requires solving a nonconvex and nonlinear optimization problem which is hard in general. Despite this hardness, we develop two algorithms named Securing Additional margin For generators in Economic dispatch (SAFE) Algorithm and Iteratively MiniMize and boUNd Economic dispatch (IMMUNE) Algorithm for finding suboptimal yet robust operating points for the generators efficiently. The SAFE Algorithm provides robust operating points for the generators by solving a single Linear Program (LP). The IMMUNE Algorithm, on the other hand, requires a few iterations until it converges, but it provides robust operating points with lower costs than the ones obtained by the SAFE Algorithm.

In situations that the operating cost of the grid in a robust state is costly (or no robust operating points exists due to lack of enough resources), the grid operator may decide to allow temporary line overloads–by increasing thresholds on circuit breakers–in the case of a MAD attack, and clear the overloads during the *secondary control*. During the secondary control, which comes right after the automatic primary control, the grid operator can directly change generators' operating points in order to bring back the system's frequency to its nominal value and clear any line overloads. To make sure that line overloads can be cleared during the secondary control, the grid operator needs to verify in advance whether for any potential MAD attack, there exist operating points for the generators satisfying demands such that no lines are overloaded (namely, the grid is *secondary controllable*). However, due to the extent of the attack space, checking all possible attack scenarios is numerically impossible. Hence, we develop several predetermined control policies that can be used to verify the secondary controllability of the grid in most scenarios with no false positives.

We then evaluate the *robustness* of grids against MAD attacks with different magnitudes. The magnitude of an attack can be determined by the fraction of demand (denoted by α) that the adversary can increase or decrease at each location. We call a grid αD-robust if either line overloads can be prevented (i.e., robust operating points exists for generators) or they can be cleared during the secondary control (i.e., grid is secondary controllable) after any MAD attacks by *an adversary that can change the demands by at most α fraction*. In general, finding the maximum α such that a given grid is αD-robust, is hard. However, by focusing on grid secondary controllability and the developed predetermined control policies, we provide efficient methods to compute practical upper and lower bounds for the maximum α in polynomial time.

Finally, we numerically evaluate the performance of the developed algorithms and controllers. For example, in New England 39-bus system, we show that the SAFE and IMMUNE Algorithms find operating points for the generators with at most 6 and 2 percent increase in the total operating cost such that the grid is robust against MAD attacks of magnitude $\alpha = 0.08$. We also evaluate the performance of the developed methods for approximating the maximum α that grid is αD-robust and show that for example in New England 39-bus system, the provided lower and upper bounds are tight and are equal to the maximum $\alpha^{\max} = 0.0962$.

To the best of knowledge, our work is the first to study the effects of potential MAD attacks on the *power flows* in the grid and provide efficient preventive algorithms to avoid line failures after the primary control response, and also efficient methods to verify if the line overloads can be cleared during the secondary control. These algorithms and methods can be adopted by grid operators to protect their systems against MAD attacks now and in the near future.

The rest of this paper is organized as follows: Section 2 provides related work and Section 3 presents a brief introduction to the power system's operation and control. In Section 4, we introduce the MAD attacks and provide their basic properties. In Section 5, we present the SAFE and IMMUNE algorithms and in Section 6, we provide efficient methods for verifying secondary controllability of a grid. Section 7 provides methods to evaluate the robustness of grids against MAD attacks and Section 8 presents numerical results. Finally, Section 9 provides concluding remarks and future directions. To improve the readability of the paper, some of the proofs are moved to Section 10.

2 Related Work

Power systems' vulnerability to failures and attacks has been widely studied in the past few years [7], [8], [9], [10], [11], [12]. In a recent work [13], Garcia et al. introduced Harvey malware that affects power grid control systems and can execute malicious commands. Theoretical methods for detecting cyber attacks on power grids and recovering information after such attacks have also been developed [14], [15], [16], [17], [18], [19], [20], [21], [22]. Another related type of cyber attacks called *load redistribution attacks* has been studied by Yuan et al. [23]. However, these type of attacks *change only the measurements* at the loads in order to force the grid operator into problematic corrective actions rather than actually changing the loads as have been studied in our work. Overall, most of the previous work on protecting the grid against attacks have focused on attacks that directly target the power grid's physical infrastructure or its control system.

The possibility of load altering attacks on smart meters and large cloud servers has been first introduced by Mohsenian et al. [24]. Their work was mostly focused on minimizing the total cost of protecting the loads (which is not always possible, especially for distributed IoT devices) against such attacks. Amini et al. [25] have also recently studied the effects of load altering attacks on the system's dynamics

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TNSE.2019.2922131, IEEE Transactions on Network Science and Engineering

and ways to use the system's frequency as a feedback to improve an attack. However, until very recently, practical ways to perform such attacks in a large-scale and their consequences on power flows were not fully studied [3]. Hence, little attention has been given to protecting the grid against line failures caused by these type of attacks.

In three very recent papers, Dvorkin and Sang [26], Dabrowski et al. [4], and our work [3] revealed the possibility of exploiting compromised IoT devices to manipulate the demands and to disrupt the normal operation of the power grid. Dvorkin and Sang [26] modeled their attack as an optimization problem for the adversary—with complete knowledge of the grid—to cause circuit breakers to trip in the distribution network. Dabrowski et al. [4] studied the effect of demand increases caused by remote activation of CPUs, GPUs, hard disks, screen brightness, and printers on the frequency of the European power grid. In [3], we analyzed the effects of sudden increase and decrease in the demand via an IoT botnet of high-wattage devices from various operational perspectives and demonstrated that besides frequency instability, such attacks can also result in widespread *cascading line failures* in the transmission network leading to large-scale blackouts. Nevertheless, practical preventive defenses against possible *line failures* caused by these attacks have not been developed yet.

Finally, while there have been extensive efforts in recent years to develop efficient algorithms for solving the Optimal Power Flow (OPF) problem [27], [28], [29] and its different variations including *Security Constrained* OPF (SC-OPF) [30] (which considers grid robustness against possible line outages) and *Chance Constrained* OPF (CC-OPF) [31] (which considers uncertainty in the output of the renewable resources), since these works do not consider grid robustness against *adversarial changes in the demands*, our work is different from previously studied variations of the OPF problem. Moreover, the second part of this work deals with secondary controllability of the grid after an attack which is a totally different problem from the OPF and its variations.

3 Model and Definitions

In this section, we provide a brief introduction to power systems' operation and control. Our focus is on the power transmission network.

Throughout this paper, we use bold uppercase characters to denote matrices (e.g., $\mathbf{A}$), italic uppercase characters to denote sets (e.g., V), and italic lowercase characters and overline arrow to denote column vectors (e.g., $\vec{\theta}$). For a matrix $\mathbf{Q}$, $\mathbf{Q}_i$ denotes its i^{th} row, q_{ij} denotes its $(i,j)^{\text{th}}$ entry, and $\mathbf{Q}^T$ denotes its transpose. For a column vector $\vec{y}$, $\vec{y}^T$ denotes its transpose, and $\|\vec{y}\|_1 := \sum_{i=1}^n |y_i|$ is its l_1-norm. For a variable x, $\text{sgn}(x)$ denotes its sign, and $\overline{x}$ and $\underline{x}$ denote its upper and lower limits, respectively. For a vector $\vec{y}$, for simplicity of notation, we drop the vector sign $\vec{}$ in denoting vectors of upper and lower limits on the entries of $\vec{y}$ as $\overline{y}$ and $\underline{y}$, respectively. Finally, $\vec{e}_1, \ldots, \vec{e}_n$ denote the fundamental basis of $\mathbb{R}^n$ and $\vec{1} = \sum_{i=1}^n \vec{e}_i$ denotes the all ones vector.

3.1 Power Flows

Power flows are governed by a set of differential equations. In the steady-state, using *phasors*, these differential equations can be reduced to a set of algebraic equations on complex numbers known as *Alternating Current (AC)* power flow model. Due to the nonlinearity of AC power flow equations and the computational complexity of solving these equations, in practice and in day-ahead power grid contingency analysis and planning, the linearized version of these equations known as *Direct Current (DC)* power flow model is widely being used [27]. Hence, in this work, we also use the DC power flow model for our analysis. This allows us to focus on the complexities of MAD attacks instead of nonlinearity of AC power flows. Nevertheless, the *main ideas* of the algorithms developed in this work can be extended to the AC power flow model as well (e.g., by combining them with the recently introduced convex relaxation methods for solving the AC Optimal Power Flow (ACOPF) problem [28]), albeit not effortlessly.

We represent the power grid by a connected directed graph $G = (V, E)$ where $V = \{1, 2, \ldots, n\}$ and $E = \{e_1, \ldots, e_m\}$ are the set of nodes and edges corresponding to the *buses* and *transmission lines*, respectively (the definition implies $|V| = n$ and $|E| = m$). Each edge e is a set of two nodes $e = (i, j)$. (Direction of the edges are arbitrary.) $\vec{p_d} \geq 0$ and $\vec{p_g} \geq 0$ denote the vector of power demand and supply values, respectively. Accordingly, $\vec{p} = \vec{p_g} - \vec{p_d}$ denotes the vector of total supply and demand values. Since the sum of supply should be equal to the sum of demand,

$$\vec{1}^T \vec{p} = 0, \tag{1}$$

in which $\vec{1}$ is an all ones vector. In the DC model, lines are also assumed to be *purely reactive*, implying that each edge $e = (i, j) \in E$ is characterized by its *reactance* $x_e = x_{ij} > 0$.

Given the power supply/demand vector $\vec{p} \in \mathbb{R}^{n \times 1}$ and the reactance values, the vector of power flows on the lines $\vec{f} \in \mathbb{R}^{m \times 1}$ can be computed by solving the following linear equations:

$$\mathbf{A}\vec{\theta} = \vec{p}, \tag{2}$$

$$\mathbf{YD}^T \vec{\theta} = \vec{f}, \tag{3}$$

where $\vec{\theta} \in \mathbb{R}^{n \times 1}$ is the vector of voltage phase angles at nodes, $\mathbf{D} \in \{-1, 0, 1\}^{n \times m}$ is the *incidence matrix* of G defined as,

$$d_{ik} = \begin{cases} 0 & \text{if } e_k \text{ is not incident to node } i, \\ 1 & \text{if } e_k \text{ is coming out of node } i, \\ -1 & \text{if } e_k \text{ is going into node } i, \end{cases}$$

$\mathbf{Y} := \text{diag}([1/x_{e_1}, 1/x_{e_2}, \ldots, 1/x_{e_m}])$ is a diagonal matrix with diagonal entries equal to the inverse of the reactance values, and $\mathbf{A} = \mathbf{DYD}^T$ is the *admittance matrix* of G.[1]

Since $\mathbf{A}$ is not a full-rank matrix, we follow [8] and use the *pseudo-inverse* of $\mathbf{A}$, denoted by $\mathbf{A}^+$ to solve (2) as $\vec{\theta} = \mathbf{A}^+\vec{p}$. Once $\vec{\theta}$ is computed, $\vec{f}$ can be computed from (3) as $\vec{f} = \mathbf{YD}^T\mathbf{A}^+\vec{p}$. For the convenience of notation, we define $\mathbf{B} := \mathbf{YD}^T\mathbf{A}^+$. Hence, $\vec{f} = \mathbf{B}\vec{p}$.

3.2 Power Grid Operation

Stable operation of the power grid relies on the persistent balance between the power supply and demand. In order

[1] The admittance matrix $\mathbf{A}$ is also known as the *weighted Laplacian matrix* of the graph [32] in graph theory.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TNSE.2019.2922131, IEEE Transactions on Network Science and Engineering

to keep the balance between the power supply and the demand, power system operators use weather data as well as historical power consumption data to predict the power demand on a daily and hourly basis [33]. This allows the system operators to plan in advance and only deploy enough generators to meet the demand in the hours ahead without overloading any power lines. This planning ahead consists of two parts: *unit commitment* and *economic dispatch*.

In unit commitment which is mainly performed daily, the grid operator selects a set of generators to *commit* their availability during the day-ahead operation of the grid. But the actual operating points of the generators (i.e., generation outputs) are determined by the operator during the day and in the process known as *economic dispatch*. The main goal of the operator during economic dispatch is to ensure reliable operation of the grid with minimum power generation cost. When feasibility of the power flows is also considered during economic dispatch, the process is also known as *Optimal Power Flow (OPF)* problem. Since in practice feasibility of power flows is always being considered, these two terms can be used interchangeably most of the times.

In this work, we mainly focus on ensuring the robustness of the grid during the economic dispatch. Extending our methods to the unit commitment process is beyond the scope of this paper and is part of the future work. Hence, here we assume that the set of available generators are given. The main challenge is to obtain a favorable operating point for these generators.

3.2.1 Optimal Power Flow

In the OPF problem, given the vector of predicted demand values $\vec{p_d}$, the grid operator needs to find the operating point vector $\vec{p_g}$ for the generators such that supply matches the demand (i.e., $\vec{1}^T(\vec{p_g} - \vec{p_d}) = 0$), the operating and physical constraints are satisfied, and the operating cost of the generators are minimized.

In particular, each line f_{ij} has a thermal power flow limit $\overline{f_{ij}}$ limiting the amount of power that a line can *safely* carry. If the power flow on a line goes above this limit (i.e., *overloads*), in most of the cases, it will be tripped by a circuit breaker in order to keep the line from breaking due to overheating. Hence, during the normal operation of the grid

$$|f_{ij}| \leq \overline{f_{ij}}, \quad \forall (i,j) \in E. \tag{4}$$

The amount of power that each generator p_{gi} is generating is also limited by a maximum ($\overline{p_{gi}}$) and a minimum ($\underline{p_{gi}}$) value. If there are no generators at node i, then $\overline{p_{gi}} = \underline{p_{gi}} = 0$. Hence,

$$\underline{p_g} \leq \vec{p_g} \leq \overline{p_g}. \tag{5}$$

The generation cost at each generator is a given by a cost function $c_i(x)$ in $\$/hr$. Given these cost functions, the OPF problem can be formulated as follows:

$$\begin{aligned} \min_{\vec{\theta},\vec{f},\vec{p_g}} \quad & \sum_{l=1}^{n} c_l(p_{gl}), \\ \text{s.t.} \quad & (1),(2),(3),(4),(5), \\ & \vec{p} = \vec{p_g} - \vec{p_d}. \end{aligned} \tag{6}$$

Several methods for finding an optimal solution to (6) depending on the cost functions exist in the literature [27]. Here, we assume that the cost functions are convex and therefore the OPF problem can be solved optimally in polynomial time. *Our main focus in Section 5 is on how to add additional constraints to the OPF problem to ensure grid robustness against MAD attacks without making the problem nonconvex.*

3.3 Frequency control

In power systems, the rotating speed of generators corresponds to the frequency. When demand becomes greater than supply, the rotating speeds of turbine generators' rotors decelerate, and the kinetic energy of the rotors is released into the system in response to the extra demand. Correspondingly, this causes a drop in the system's frequency. This behavior of turbine generators corresponds to Newton's first law of motion and is calculated by the *inertia* of the generators. Similarly, the supply being greater than the demand results in acceleration of the generators' rotors and a rise in the system's frequency.

This decrease/increase in the frequency of the system cannot be tolerated for a long time since frequencies lower than their nominal value severely damage the generators. If the frequency goes above or below a threshold value, protection relays turn off or disconnect the generators completely. Hence, in case of a demand increase, within seconds of the first signs of a decrease in the frequency, the *primary controllers* at generators activate and increase the mechanical input to the generators which increase the speed of the generator's rotor and correspondingly the generator's output and frequency of the system [34]. The rate of decrease/increase in the frequency of the system, before activation of the primary controllers, directly depends on the total *inertia* of the system. Systems with a higher number of rotating generators have higher inertia and therefore are more robust against sudden demand changes or generation losses.

The rate of increase in the output generation of generator i during the primary control is determined by its *governor droop characteristic* denoted by R_i [35, Chapter 9]. In particular, after a change in the total demand by $S_{\Delta p_d}$, the primary controller of each generator i increases its output with rate $1/R_i$ until the total generation is equal to the demand again. In particular, if none of the generators reach their generation limit, each generator i will increase its generation by $1/R_i \times S_{\Delta p_d}/(\sum_{l=1}^{n} 1/R_l)$. The amount of power that generators can provide during the primary control is called the *spinning reserve* of the generators.

Despite the stability of the system's frequency after the primary controllers' response, it may not return to its nominal value (since generators generating more than their generating set points). Hence, the *secondary controller* starts within minutes to restore the system's frequency. The secondary controller modifies the power set points and deploys available extra generators and controllable demands to restore the nominal frequency and permanently stabilizes the system.[2] Fig. 2 presents an example of the way frequency of the system changes after a sudden increase in the demand (or loss of generation) at time 0.

[2] Part of these controls can be done during the *tertiary control*. However, for simplicity and without loss of generality we refer to them as the secondary control.

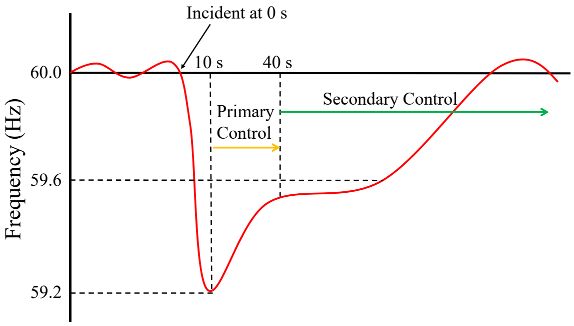

Fig. 2: A sample frequency response of the power grid to a sudden increase in the demand (or loss of generation).

4 MAD Attacks

In this work, we follow the threat model that we have initially introduced in [3]. In particular, we assume that an adversary has already gained access to an IoT botnet of many high-wattage smart appliances within a city, a country, or a continent. Such access can potentially allow the adversary to increase or decrease the demand at different locations *remotely and synchronously* at a certain time. We call the attacks under this threat model the MAnipulation of the Demand (MAD) attacks.

Since the focus of this work is to develop defenses against MAD attacks rather than dealing with complexities of performing such an attack (as extensively studied in [3]), we abstract the threat model by the adversary's power to manipulate the demands at each node. In particular, we assume the demand changes at node l by an adversary are bounded by $-\overline{\Delta p_{dl}} \leq \Delta p_{dl} \leq \overline{\Delta p_{dl}}$. Notice that from defensive point of view, there are no differences between an adversary with the total knowledge of the system (a.k.a *white-box* attacks) and an adversary with no knowledge of the system (a.k.a *black-box* attacks), since the operator needs to make sure that the grid is robust against *any possible attacks*.

The initial effect of a MAD attack, as described in Section 3.3 is on the frequency of the system. However, the system operator can make the system robust against frequency disturbances caused by MAD attacks by ensuring that enough generators with inertia and spinning reserve are committed to operate during the unit commitment process [3]. The minimum required inertia and spinning reserve should be computed based on the potential attack size and the properties of the grid. Devices that provide virtual inertia such as batteries, super-capacitors, and flywheels can also be integrated into the system to increase the total inertia [36].

Hence, the main challenge in protecting the grid against the initial effects of MAD attacks is at the hardware level. However, the effects of MAD attacks are not limited to frequency disturbances. Recall from Section 3.1 that the power flows in power grids are determined uniquely given supply and demand values. Therefore, most of the time, the grid operator does not have any control over the power flows from generators to loads. Once an adversary causes a sudden increase in the loads all around the grid, assuming that the frequency drop is not significant, the extra demand is satisfied automatically by generators through their primary controllers as described in Section 3.3. Since the power flows are not controlled by the grid operator at this stage, this change in supply and demand may result in line overloads and consequent line-trippings [3].

If the primary controllers' response results in line overloads, assuming that these overloads can barely be tolerated for a short period of time, these line overloads can be cleared during the secondary control. However, the system operator needs to ensure in advance that possible line overloads can indeed be cleared during the secondary control after any MAD attacks.

In this work, we focus on the effects of MAD attacks on the power flow changes on the lines which are more challenging from the system planning perspective. Our objectives are: (i) to develop algorithms for finding efficient operating points for the generators during the economic dispatch such that no lines are overloaded after the primary control response to any potential MAD attacks, and (ii) to design methods to efficiently examine if line overloads after the primary control–if any–can be cleared during the secondary control.

Notice that we assume the system have enough inertia and reaches a steady-state after the primary controllers' response to a MAD attack (as in Fig. 2). Moreover, since power lines can normally withstand sudden but momentary power surges, in analyzing power flows after a contingency, the transient power flows are usually neglected [27]. Therefore, it is reasonable to use the steady-state power flow equations as described in Section 3.1 for our analysis.

5 Power Flows: Primary Control

In this section, we provide two algorithms for finding operating points for the generators during the economic dispatch process such that no lines are overloaded after the automatic response of the primary controllers to any MAD attacks. We call such operating points, *robust operating points*.

5.1 Power Flow Changes

In this subsection, we present a couple of examples in order to demonstrate the complexity of power flow analysis after the *primary controller's response to a MAD attack*.

First, as can be seen in Fig. 3 the relationship between the power flow changes on the lines and the demand changes is not intuitive. For example, flow on line $(2,3)$ is maximized when only the demand at node 3 increases (Fig. 3(c)), whereas when demands at both nodes 1 and 3 increase, flow on line $(2,3)$ increases less (Fig. 3(d)).

Another important factor affecting the amount of power flow changes on the lines is the amount of spinning reserve at each generator. For example, as can be seen in Fig. 4, an increase in the demand at node 1 by 3 units may result in power flow *decrease* on line $(2,3)$ if all the generators have enough spinning reserves (Fig. 4(a)). The same scenario, however, results in power flow *increase* on line $(2,3)$, if only generators 2 and 4 have spinning reserves (Fig. 4(b)).

Fig. 5 presents the relationship between power flow changes on lines $(2,3)$ and $(5,3)$ versus power demand increase at node 1 during two different spinning reserve availability scenarios in the grid shown in Fig. 3(a). As can be seen in Fig. 5(a), if all generators have enough spinning

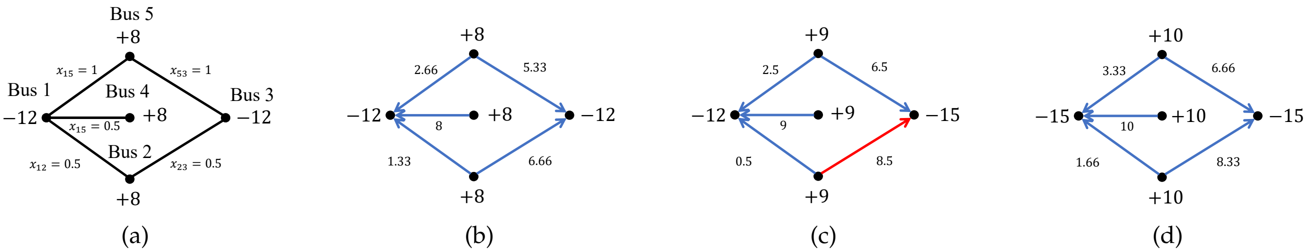

Fig. 3: An example demonstrating that increasing all demands may not necessarily result in the maximum flow on the lines. (a-b) Initial setting and power flows, (c) power flows if demand at bus 3 increases, and (d) power flows if demand at both buses 1 and 3 increases. All generators have the same droop characteristic and they all have enough spinning reserve.

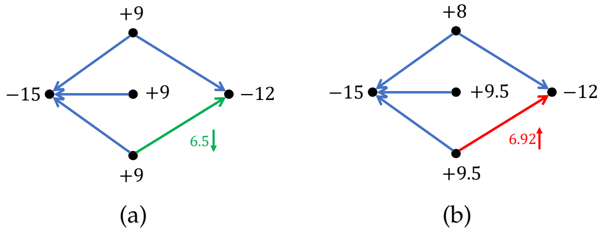

Fig. 4: Dependency of power flow changes on the location of the spinning reserves. (a) If all generators have spinning reserves, demand increase at bus 1 results in power flow decrease on line $(2,3)$. (b) If only generators 2 and 4 have spinning reserves then demand increase at bus 1 results power flow increase on line $(2,3)$.

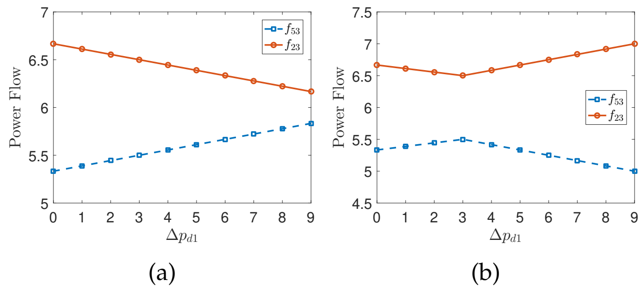

Fig. 5: Power flows on lines $(5,3)$ and $(2,3)$ in the grid shown in Fig. 3(a) as demand at bus 1 increases. (a) If all the generators have enough spinning reserve, and (b) if generator 5 has only 1 unit of spinning reserve.

reserve the power flows change monotonically with the demand change. However, as can be seen in Fig. 5(b), limited spinning reserve at generator 5 results in a nonlinear relationship between the power flows and the demand change.

Following the examples provided in this subsection, it is clear that power flow changes on the lines after a MAD attack highly depend on the initial operating point of the grid and is a nonlinear problem in most cases. Despite the difficulties, however, in the next two subsections, we provide efficient algorithms for finding efficient and robust operating points for the generators.

5.2 SAFE Algorithm

In order to avoid line overloads after the primary control response to a potential MAD attack, the grid operator needs to compute the maximum possible power flow changes on the lines following an attack (based on $\overline{\Delta p_{dl}}$ values) and enforce the power flows on the lines in OPF to be below their capacity minus the maximum possible changes. As shown in the previous subsection, however, the maximum power flow changes on the lines depend on the operating point of the generators and their spinning reserve. Therefore, one cannot compute the maximum power flow changes on the lines independent of the operating points to be used in the OPF problem.

One way to circumvent this problem, is to enforce all the generators to have enough spinning reserves to keep the relationship between the power flow changes and demand changes linear (as in Fig. 5(a)), and use this linear relationship to compute the maximum power flow changes on the lines based on the operating point of the generators. These values can then be added to the OPF problem without making the problem nonlinear and nonconvex. Recall that since here we use DC power flows with convex cost functions, the OPF problem is convex. Hence, when we mention the nonconvexity of the problem, it is due to additional constraints on the power flows.

For each load i, define $\vec{v}_i = [v_{i1}, v_{i2}, \ldots, v_{in}]^T$ to denote the primary controllers' response to a unit demand increase at load i. If all generators have enough spinning reserve, each generator j will increase its generation by $v_{ij} := (1/R_j)/(\sum_{l=1}^n 1/R_l)$ to compensate for a unit demand increase at node i (as described in Section 3.3). Hence, by defining $\vec{w}_i := \vec{v}_i - \vec{e}_i$ (recall from Section 3 that $\vec{e}_i$ is the ith fundamental basis of $\mathbb{R}^n$) one can compute the change in the flow of line $e = (i,j)$ solely in terms of changes in the demands (Δp_{di}s):

$$\Delta f_{ij} = 1/x_{ij}(\mathbf{A}_i^+ - \mathbf{A}_j^+)\sum_{l=1}^n \Delta p_{dl}\vec{w}_l. \qquad (7)$$

Recall that $-\overline{\Delta p_{dl}} \leq \Delta p_{dl} \leq \overline{\Delta p_{dl}}$ based on the grid operator's estimation of the adversary's power. Hence, the maximum flow change on line (i,j) can be computed using (7) as:

$$\Delta f_{ij}^{\max} = 1/x_{ij}\sum_{l=1}^n \overline{\Delta p_{dl}}|(\mathbf{A}_i^+ - \mathbf{A}_j^+)\vec{w}_l|, \qquad (8)$$

since for each l, Δp_{dl} can be selected by the adversary to be equal to $-\overline{\Delta p_{dl}}$, if $(\mathbf{A}_i^+ - \mathbf{A}_j^+)^T\vec{w}_l < 0$, and equal to $\overline{\Delta p_{dl}}$, if $(\mathbf{A}_i^+ - \mathbf{A}_j^+)\vec{w}_l \geq 0$. Now, to ensure that no lines are overloaded after a MAD attack, all the system operator needs to do is to replace the capacity of each line (i, j) in the OPF problem by $\overline{f_{ij}} - \Delta f_{ij}^{\max}$. The only other constraint that needs to be added to the OPF problem is to make sure that each generator i with $0 < 1/R_i$ has enough spinning reserve to increase its generation according to its governor droop. For this, define $\overline{S_{\Delta p_d}} := \sum_{l=1}^{n} \overline{\Delta p_{dl}}$. Hence, each generator's operating point should be within the following limits:

$$\forall 1 \leq i \leq n:$$
$$\underline{p_{gi}} + \frac{1/R_i}{\sum_{l=1}^{n} 1/R_l} \overline{S_{\Delta p_d}} \leq p_{gi} \leq \overline{p_{gi}} - \frac{1/R_i}{\sum_{l=1}^{n} 1/R_l} \overline{S_{\Delta p_d}}. \quad (9)$$

Therefore, the robust OPF problem can be written as follows:

$$\begin{aligned} \min_{\vec{\theta}, \vec{f}, \vec{p_g}} \quad & \sum_{l=1}^{n} c_l(p_{gl}), & (10) \\ \text{s.t.} \quad & (1), (2), (3), (8), (9), & \\ & |f_{ij}| \leq \overline{f_{ij}} - \Delta f_{ij}^{\max}, \quad \forall (i,j) \in E & \\ & \vec{p} = \vec{p_g} - \vec{p_d}. & \end{aligned}$$

We call the algorithm for finding a robust operating point for generators by limiting their operating points—to be able to analytically compute $\Delta f_{ij}^{\max}$s—and solving (10), the Securing Additional margin For generators in Economic dispatch (SAFE) Algorithm. Since this algorithm limits the operating points of the generators by adding conditions (9) to the OPF problem, it is obvious that it may not obtain the *minimum cost* robust operating points for the generators. In the next subsection, we provide an algorithm, albeit computationally more expensive, for finding robust operating points for the generators without limiting their operating points—as in (9).

5.3 IMMUNE Algorithm

In (7), we assumed that none of the generators reach their maximum/minimum capacity as they increase/decrease their generation according to their droop characteristics. However, by allowing some generators to reach their maximum/minimum capacity, one may find robust operating points for the generators with a lower cost.

In this subsection, for brevity and to avoid repetition, we assume that the total demand change $S_{\Delta p_d} := \sum_{i=1}^{n} \Delta p_{di}$ can only be positive. Hence, we focus mainly on the generators' maximum capacity. However, the same set of equations can similarly be derived for the case $S_{\Delta p_d} < 0$ which should also be considered separately in computing the maximum power flow changes on the lines. In particular, whenever there is a minimization/maximization problem with $S_{\Delta p_d} \geq 0$ constraint, one should also solve a similar optimization problem with $S_{\Delta p_d} < 0$ and take the minimum/maximum of the optimal value of the two optimization problems. In Section 8, we consider both cases for numerical evaluations.

Once a generator reaches its maximum capacity, it cannot increase its generation anymore, and therefore other generators should generate more to compensate for the extra demand. The following lemma provides the amount each generator generates based on its spinning reserve and governor droop characteristic to compensate for the extra demand after a MAD attack.

Lemma 1. *Suppose generators are ordered such that if $i < j$, $R_i(\overline{p_{gi}} - p_{gi}) \leq R_j(\overline{p_{gj}} - p_{gj})$. Define $t_i := R_i(\overline{p_{gi}} - p_{gi})$ and $S_i := \sum_{l=1}^{i} t_l/R_l + \sum_{l=i+1}^{n} t_i/R_l$. If $S_i < S_{\Delta p_d} \leq S_{i+1}$, to compensate for the extra demand, generators 1 to i reach their maximum capacity and each generator $j > i$ generates $\frac{1/R_j}{\sum_{l=i+1}^{n} 1/R_l} \left(S_{\Delta p_d} - \sum_{l=1}^{i} (\overline{p_{gl}} - p_{gl})\right)$.*

In general, as demonstrated in Figs. 4 and 5, due to power generation limits, power flow on a line may not change monotonically as demand changes in a specific node–as in (7). Hence, the maximum change in the power flows cannot be found in a closed form as in (8). However, one may be able to find an upper bound on the maximum power flow change on a line.

Upper bounds on the maximum power flow changes after a MAD attack can be computed by assuming the worst case initial operating points and also assuming that generators can be arbitrarily assigned to provide extra required generation. In particular, an upper bound $\widehat{\Delta f_{ij}}$ for the power flow changes on line (i, j) can be computed by finding the worst initial operating points for the generators $\vec{p_g}$ and the worst possible way to increase the power generations $\Delta \vec{p_g}$ (in oppose to the automatic primary controller's response) in response to the worst possible way to increase the demands by an adversary $\Delta \vec{p_d}$ as follows:

$$\begin{aligned} \widehat{\Delta f_{ij}} := \max_{\vec{p_g}, \Delta\vec{p}_d, \Delta\vec{p}_g} \quad & \left| 1/x_{ij}(\mathbf{A}_i^+ - \mathbf{A}_j^+)(\Delta\vec{p}_g - \Delta\vec{p}_d) \right| & (11) \\ \text{s.t.} \quad & \vec{1}^T(\vec{p_g} - \vec{p_d}) = 0, & \\ & \vec{1}^T(\Delta\vec{p}_g - \Delta\vec{p}_d) = 0, & \\ & -\overline{\Delta p_{dl}} \leq \Delta p_{dl} \leq \overline{\Delta p_{dl}}, \quad 1 \leq l \leq n & \\ & \underline{p_g} \leq \vec{p_g} \leq \overline{p_g}, & \\ & 0 \leq \Delta p_{gl} \leq \overline{p_{gl}} - p_{gl}, \quad 1 \leq l \leq n, & \\ & S_{\Delta p_d} \geq 0. & \end{aligned}$$

Optimization (11) is a Linear Program (LP) that can be solved efficiently for each line (i, j). Using these upper bounds, we can limit the power flows on the lines in the OPF problem (6) as $|f_{ij}| \leq \overline{f_{ij}} - \widehat{\Delta f_{ij}}$ to leave enough margin for the lines in case of a MAD attack. Hence, the solution to the following modified OPF problem provides robust operating points for the generators:

$$\begin{aligned} \min_{\vec{\theta}, \vec{f}, \vec{p_g}} \quad & \sum_{l=1}^{n} c_l(p_{gl}), & (12) \\ \text{s.t.} \quad & (1), (2), (3), (5), & \\ & |f_{ij}| \leq \overline{f_{ij}} - \widehat{\Delta f_{ij}}, \quad \forall (i,j) \in E & \\ & \vec{p} = \vec{p_g} - \vec{p_d}. & \end{aligned}$$

Enforcing the power flows on all the lines, such as (i, j), to be less than $\overline{f_{ij}} - \widehat{\Delta f_{ij}}$ as in (12) ensures that none of the lines will be overloaded after a potential MAD attack. However, the solution to (12) may not provide the optimal robust operating points for the generators since $\widehat{\Delta f_{ij}}$s only provide an upper bound on the maximum power flow changes on the lines. To achieve more efficient robust operating points, we introduce an iterative algorithm that solves the OPF problem and updates the lines' required safety margins to ensure that none of the lines get overloaded after a MAD attack. We will then use the upper bounds $\widehat{\Delta f_{ij}}$s to prove that the algorithm will converge to a local optimal solution.

First, given the operating points $p_{g1}, \dots, p_{gn}$ to the OPF problem, the maximum power flow change on line (i, j) (denoted by $\Delta f_{ij}^{\max}$) after an attack can be computed based on the power flow solution $\vec{f} = \mathbf{Y}\mathbf{D}^T\mathbf{A}^{+}\vec{p}$ by solving the following optimization problem:

$$\begin{aligned}\Delta f_{ij}^{\max} = \max_{\Delta \vec{p}_d} \quad & \operatorname{sgn}(f_{ij})\Big(1/x_{ij}\sum_{l=1}^{n} -\Delta p_{dl}(a_{il}^{+} - a_{jl}^{+}) \\ & + 1/x_{ij}\sum_{l=1}^{n} f_l(S_{\Delta p_d})(a_{il}^{+} - a_{jl}^{+})\Big) \\ \text{s.t.} \quad & -\overline{\Delta p_{dl}} \leq \Delta p_{dl} \leq \overline{\Delta p_{dl}}, \quad 1 \leq l \leq n \\ & S_{\Delta p_d} \geq 0.\end{aligned} \qquad (13)$$

in which $f_l(\cdot)$s denote piecewise linear functions that determine the extra output of the generators based on the total demand change $S_{\Delta p_d}$. Since we assumed that $p_{g1}, \dots, p_{gn}$ are given, functions $f_l(\cdot)$ can be uniquely determined using Lemma 1. $\operatorname{sgn}(f_{ij})$ in the objective of (13) is to ensure that the maximum changes are in the direction of *increase* in the power flow on line (i, j). Hence, for all lines $\Delta f_{ij}^{\max} \geq 0$.[3]

Lemma 2. *Optimization (13) can be solved in polynomial time for each $(i, j) \in E$.*

Proof: Without loss of generality, assume that generators are ordered such that $t_1 \leq t_2 \leq \cdots \leq t_n$ as defined in Lemma 1. It is easy to see that by using Lemma 1 and defining $S_0 := 0$, one can solve (13) in different linear regions of $f_l(\cdot)$s by considering additional conditions for $S_{\Delta p_d}$ (for $0 \leq z < n$):

$$S_z \leq S_{\Delta p_d} < S_{z+1}. \qquad (14)$$

Under condition (14), $f_l(\cdot)$s can be determined as follows:

$$f_l(S_{\Delta p_d}) = \begin{cases} \overline{p_l} - p_l & l \leq z, \\ \frac{1/R_l\left(S_{\Delta p_d} - \sum_{w=1}^{z}(\overline{p_w} - p_w)\right)}{\sum_{w=z+1}^{n} 1/R_w} & l > z. \end{cases} \qquad (15)$$

Hence, all the $f_l(\cdot)$ are either constant or linear functions in (13) and therefore (13) can be solved efficiently using LP. Hence, by solving (13) at most n times (once for every condition (14) for different z) $\Delta f_{ij}^{\max}$ can be found in polynomial time. □

After computing $\Delta f_{ij}^{\max}$ values, one can use them to verify if any of the lines will be overloaded after an attack (e.g., by checking if $\overline{f_{ij}} < |f_{ij}| + \Delta f_{ij}^{\max}$). If yes, then update the required margins for the lines that may get overloaded in the OPF problem to ensure that those lines will not be overloaded. The OPF problem can then be solved again with new power flow margins for the lines and the process continues until no additional updates for the line margins are required at the obtained operating point (which means that the obtained operating point is robust). We call this algorithm Iteratively MiniMize and boUNd Economic dispatch (IMMUNE) Algorithm (summarized in Algorithm 1).

Algorithm 1: Iteratively MiniMize and boUNd Economic dispatch (IMMUNE)

```
Input: $G$
1:  flag = 1
2:  Define $c_{ij} := \overline{f_{ij}}$ for all $(i,j) \in E$
3:  while flag do
4:      Solve the OPF problem (6) such that
        $\forall (i,j) \in E : |f_{ij}| \leq c_{ij}$
5:      if OPF is not feasible then
6:          return none
7:      Compute $\Delta f_{ij}^{\max}$ by solving (13) for all $(i,j) \in E$
8:      flag = 0
9:      for $(i,j) \in E$ do
10:         if $\overline{f_{ij}} < |f_{ij}| + \Delta f_{ij}^{\max}$ then
11:             $c_{ij} = \overline{f_{ij}} - \Delta f_{ij}^{\max}$
12:             flag = 1
13: return $p_{g1}, p_{g2}, \dots, p_{gn}$
```

Lemma 3. *If (12) is feasible, then the IMMUNE Algorithm converges to a local optimum solution.*

Lemma 3 provides a sufficient condition such that the IMMUNE Algorithm converges to a local optimum. However, even if (12) is not feasible, the system operator can still run the IMMUNE Algorithm to obtain a local optimum solution if the OPF problem remains feasible at each iteration of the algorithm.

We can also provide an upper bound on the number of iterations that IMMUNE algorithm requires to converge. For this reason, the algorithm needs to change discrete changes to the capacities at each iteration.

Lemma 4. *If the IMMUNE Algorithm changes c_{ij} at each iteration by a discrete amount such as $c_{ij} = \max\{\lfloor \overline{f_{ij}} - \Delta f_{ij}^{\max} \rfloor, \overline{f_{ij}} - \widehat{\Delta f_{ij}}\}$, then it terminates in at most $O(\sum_{(i,j)\in E}\lceil \widehat{\Delta f_{ij}} \rceil)$ iterations.*

Corollary 1. *If generators' cost functions are linear and $\mathcal{F}(n)$ indicates the running time of the LP solver of choice with n variables (e.g., simplex or ellipsoid algorithms), the IMMUNE Algorithm terminates in $O(m\mathcal{F}(n)(\sum_{(i,j)\in E}\lceil \widehat{\Delta f_{ij}} \rceil))$.*

Following a similar idea, one can decrease the running time of the IMMUNE algorithm by applying more aggressive update rules for the capacities in line 11 of the algorithm. For example, line 11 can be replaced by $c_{ij} = 0.9(\overline{f_{ij}} - \Delta f_{ij}^{\max})$ or $c_{ij} = 0.95(\overline{f_{ij}} - \Delta f_{ij}^{\max})$. We call these variations of the IMMUNE Algorithm, IMMUNE-0.9, and IMMUNE-0.95. In Section 8.2, we numerically evaluate and compare the performance of these algorithms and demonstrate that more aggressive update rules result in faster convergence.

One favorable property of the IMMUNE Algorithm is that it can be easily parallelized. This parallelization can be used to simultaneously compute $\Delta f_{ij}^{\max}$ for all the lines at each iteration in order to expedite the algorithm.

If the OPF problem becomes infeasible in any iteration of the IMMUNE Algorithm, there are two ways to circumvent the issue: (i) By considering higher temporary limits for the lines (e.g., $1.1\overline{f_{ij}}$) which is a common practice in power systems operation, but the operator needs to ensure that

[3] Notice that for computing the maximum power flow changes on the lines, the $S_{\Delta p_d} < 0$ case should also be considered separately to see if it results in a larger power flow change than the one obtained from (13). However, as we mentioned at the beginning of the subsection, here we only consider $S_{\Delta p_d} \geq 0$ for the brevity of presentation.

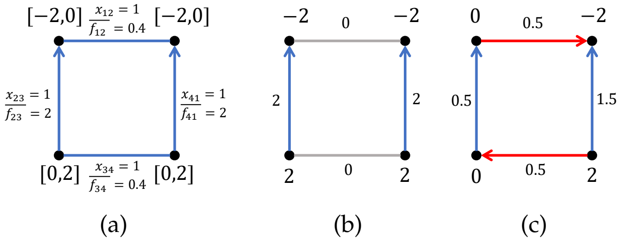
[−2,0]
$\frac{x_{12}=1}{f_{12}=0.4}$
[−2,0]
$\frac{x_{23}=1}{f_{23}=2}$
$\frac{x_{41}=1}{f_{41}=2}$
[0,2]
$\frac{x_{34}=1}{f_{34}=0.4}$
[0,2]
(a)
−2
0
−2
2
2
2
0
2
(b)
0
0.5
−2
0.5
1.5
0
0.5
2
(c)

Fig. 6: Complexity of secondary controller problem. (a) Secondary controller problem setting, (b) an attack that maximizes the demand, and (c) an attack that minimizes the demand at one node and maximizes the demand at another node.

line overloads can be cleared during the secondary control, or (ii) by returning to the unit commitment problem and change the list of committed generators to make sure (12) is feasible. We will address the first approach in the next section in detail. However, the second approach is beyond the scope of this paper and is part of our future work.

6 Power Flows: Secondary Control

In cases that primary control cannot prevent line overloads, the system operator has to clear these overloads during the secondary control instead. In such cases, the operator needs to make sure in advance that after the primary control's response to a MAD attack, there are operating points for the generators such that the demand can be supplied with no line overloads (i.e., the secondary controller can clear the overloads). Assuming that the maximum and minimum reachable demands at node i by an adversary are $\overline{p_{di}}$ and $\underline{p_{di}}$, respectively, this problem can be defined as the *secondary controller problem*:

Secondary controller problem: For any $p_{d1}, p_{d2}, \ldots, p_{dn}$ that $\forall 1 \leq i \leq n : \underline{p_{di}} \leq p_{di} \leq \overline{p_{di}}$, are there operating points $p_{g1}, \ldots, p_{gn}$ for the generators such that $\forall 1 \leq i \leq n : \underline{p_{gi}} \leq p_{gi} \leq \overline{p_{gi}}$, $\vec{1}^T(\vec{p_g} - \vec{p_d}) = 0$, and no lines are overloaded?

Definition 1. *A grid is called* secondary controllable *if the answer to the secondary controller problem is yes.*

Notice that *operating cost of the generators are not important during the secondary control* since the secondary controller activates only after a potential attack and the operator needs to bring back the grid to its normal state as soon as possible at any cost. Fig. 6 provides an example of the secondary controller problem. As can be seen in Fig. 6(b), when the demands are all equal to their maximum level after a MAD attack, the demand can be supplied by generators with no line overloads. However, as presented in Fig. 6(c), when the demand is increased to its maximum level at one node and decreased to its minimum at another one, there is no possible way to supply the demand such that no lines are overloaded. This example clearly evinces that the secondary controller problem is not intuitive.

In the following subsections, we study the secondary controller problem in detail and provide efficient algorithms to verify the secondary controllability of a power system.

6.1 Maxmin Formulation

One way of verifying the secondary controllability of a power system is by exploiting *linear bilevel programs* [37], [38]. The secondary controller problem can be written in the form of a max-min linear problem which is a special form of *linear bilevel programs* as follows:

$$\begin{array}{lll} \max\limits_{\vec{p_d}} \; \min\limits_{\vec{p_g}, \vec{q}, \vec{f}, \vec{\theta}} & \vec{1}^T \vec{q} & (16) \\ \text{s.t.} & (1), (2), (3), (4), (5), & \\ & \vec{p} = \vec{p_g} - \vec{p_d} + \vec{q}, & \\ & q_i \geq 0, \quad 1 \leq i \leq n & \\ & \underline{p_{di}} \leq p_{di} \leq \overline{p_{di}}, \quad 1 \leq i \leq n. & \end{array}$$

In optimization problem (16), vector $\vec{p_d}$ should be selected such that for the best possible selection of vector $\vec{p_g}$ and positive auxiliary vector $\vec{q}$, the objective value is maximized. The following proposition relates the solution of (16) to the secondary controller problem.

Proposition 1. *The optimal solution of (16) is 0 if, and only if, the grid is secondary controllable.*

Proof: If the optimal solution to (16) is 0, then for any demand vector $\vec{p_d}$, the vector of generation values $\vec{p_g}$ can be selected such that $\vec{1}^T(\vec{p_g} - \vec{p_d}) = 0$ and no lines are overloaded. Hence, the grid is secondary controllable. Now if the grid is secondary controllable, then for all demand vectors $\vec{p_d}$, there exists a vector of generation $\vec{p_g}$ such that $\vec{1}^T(\vec{p_d} - \vec{p_d}) = 0$ and no lines are overloaded. Hence, the auxiliary vector $\vec{q}$ can be selected to be equal to 0 by the minimization part of (16) for any vector $\vec{p_d}$. Therefore, the optimal solution to (16) would be 0. □

Proposition 1 clearly demonstrates that solving (16) can determine secondary controllability of a power system. Moreover, when the optimal solution of (16) is greater than 0, the nonzero entries of the optimal vector $\vec{q}$ can reveal the minimum extra generation required to ensure secondary controllability of the system.

Despite many advantages of the formulation (16), the max-min linear program is nonconvex [39] and proved to be NP-hard [40]. Therefore existing efficient algorithms for solving (16) only obtain local optimal solutions [38]. However, a local optimal solution of (16) with value 0 does not guarantee the secondary controllability of the system since the optimal solution may not be zero.

One way of solving (16) optimally, albeit in exponential running time, is through brute force search. Following lemma demonstrates that to solve the secondary controller problem, one needs to check only the extreme demand points due to the convexity of the space of all possible demand values and linearity of power flow equations.

Lemma 5. *The grid is secondary controllable, if and only if for all $p_{d1}, \ldots, p_{dn}$ such that $p_{di} \in \{\overline{p_{di}}, \underline{p_{di}}\}$ there exist operating points $p_{g1}, \ldots, p_{gn}$ for the generators such that $\forall 1 \leq i \leq n$: $\underline{p_{gi}} \leq p_{gi} \leq \overline{p_{gi}}$, $\vec{1}^T(\vec{p_g} - \vec{p_d}) = 0$, and no lines are overloaded.*

On the other hand, for a given demand vector $\vec{p_d}$, it can be verified in polynomial time whether there exist operating points for the generators that satisfy the secondary controller problem by solving the minimization part of (16) using LP:

$$\begin{aligned} \min_{\vec{p_g},\vec{q},\vec{f},\vec{\theta}} \quad & \vec{1}^T\vec{q} \\ \text{s.t.} \quad & (1),(2),(3),(4),(5), \\ & \vec{p} = \vec{p_g} - \vec{p_d} + \vec{q} \\ & q_i \geq 0, \quad 1 \leq i \leq n. \end{aligned} \tag{17}$$

If the optimum solution to (17) is not 0, then the optimal vector $\vec{q}$ can be used by the operator to make more generators online for controllability of the grid. Hence by solving (17) for all extreme demand vectors, one can verify secondary controllability of a system in *exponential running time* and also find how to make it controllable–if it is not–based on obtained vectors $\vec{q}$.

By focusing only on nodes with the largest demands, one can approximately verify if for a subset of extreme points there exist operating points for the generators satisfying the secondary controller problem. In general, however, such an approach may not be able to guarantee the secondary controllability of a grid. Hence, in the next subsection, we provide *sufficient conditions* to ensure secondary controllability of a grid in polynomial time.

6.2 Predetermined Secondary Controllers

Despite the difficulty in exact determination of secondary controllability of a grid, in this subsection, we introduce and exploit suboptimal predetermined controllers to verify controllability of a grid *with no false positives* (i.e., presented methods cannot determine *uncontrollability* of a system).

In order to verify secondary controllability of the grid, one can find *the best* predetermined way to set the generation values given a demand vector $\vec{p_d}$ such that the maximum power flows over all demand vectors is minimized. In particular, we define the $\vec{\beta}$-determined controller as follows.

Definition 2 ($\vec{\beta}$-determined controller)**.** *For any demand vector $\vec{p_d}$, set $\vec{p_g} = (\sum_{i=1}^n p_{di}) \times \vec{\beta}$, for a vector $\vec{\beta}$ satisfying:*

(i) $\vec{\beta} \geq 0$, *(ii)* $\vec{1}^T\vec{\beta} = 1$, *(iii)* $(\sum_{i=1}^n \overline{p_{di}}) \times \vec{\beta} \leq \overline{p_g}$,
(iv) $(\sum_{i=1}^n \underline{p_{di}}) \times \vec{\beta} \geq \underline{p_g}$.

Definition 3. *A controller is called* reliable, *if for all feasible demand vectors $\vec{p_d}$, it provides a vector of operating points for the generators like $\vec{p_g}$ such that $|\vec{f}| = |\boldsymbol{B}(\vec{p_g} - \vec{p_d})| \leq \overline{f}$.*

Proposition 2. *If there exists a vector $\vec{\beta}$ such that the $\vec{\beta}$-determined controller is reliable, then the grid is secondary controllable.*

For a vector $\vec{\beta}$ satisfying conditions (i-iv) in Definition 2, define vectors $\vec{w_i}^{(\beta)} := -\vec{e_i} + \vec{\beta}$ for $1 \leq i \leq n$ (as in Section 5.2). The following lemma proves that maximum flow on the lines over all feasible demand vectors, given a $\vec{\beta}$-determined controller, can deterministically be computed.

Lemma 6. *Given a $\vec{\beta}$-determined controller, the maximum power flow on each line e_k over all possible demand vectors is:*

$$\begin{aligned} \max_{\underline{p_d} \leq \vec{p_d} \leq \overline{p_d}} |f_k| = & \left| \sum_{i=1}^n \frac{(\overline{p_{di}} + \underline{p_{di}})}{2} \boldsymbol{B}_k \vec{w_i}^{(\beta)} \right| \\ & + \sum_{i=1}^n \frac{(\overline{p_{di}} - \underline{p_{di}})}{2} |\boldsymbol{B}_k \vec{w_i}^{(\beta)}|. \end{aligned} \tag{18}$$

The main question is now whether there exists a vector $\vec{\beta}$ such that the maximum power flows as determined in (18) are less than their capacities? We prove that one can examine this efficiently and in polynomial time by solving the following optimization:

$$\begin{aligned} \min_{\eta,\vec{\beta},\vec{f}} \quad & \eta \\ \text{s.t.} \quad & \text{(i-iv) in Definition 2}, \\ & \vec{f} = |\mathbf{BW}^{(\beta)}(\overline{p_d} + \underline{p_d})/2| + |\mathbf{BW}^{(\beta)}|(\overline{p_d} - \underline{p_d})/2, \\ & \vec{f} \leq \eta\overline{f}, \end{aligned} \tag{19}$$

in which matrix $\mathbf{W}^{(\beta)} := [\vec{w_1}^{(\beta)}, \ldots, \vec{w_n}^{(\beta)}]$. The following proposition demonstrates that (19) can be solved using LP in polynomial time. Moreover, it indicates that the optimal solution to (19) can provide the best vector $\vec{\beta}$ for deterministically controlling the grid and its optimal value demonstrates if the corresponding $\vec{\beta}$-determined controller is reliable.

Proposition 3. *Optimization (19) can be solved using LP. Moreover, if the optimal value η^* to (19) is less than or equal to 1, then the $\vec{\beta}^*$-determined controller obtained from the corresponding solution is reliable, and therefore the grid is secondary controllable.*

From (18), it can be seen that the formula for computing maximum flow on the lines consists of two separate sums which can be controlled by different vectors and obtained a better controller. Hence, one can define the $(\vec{\gamma}, \vec{\beta})$-determined controller as follows.

Definition 4 ($(\vec{\gamma}, \vec{\beta})$-determined controller)**.** *For any demand vector $\vec{p_d}$, set $\vec{p_g} = (\sum_{i=1}^n(\overline{p_{di}} + \underline{p_{di}})/2) \times \vec{\gamma} + (\sum_{i=1}^n(p_{di} - \overline{p_{di}}/2 - \underline{p_{di}}/2)) \times \vec{\beta}$, for vectors $\vec{\gamma}$ and $\vec{\beta}$ satisfying:*

(i) $\vec{\beta}, \vec{\gamma} \geq 0$, *(ii)* $\vec{1}^T\vec{\gamma} = \vec{1}^T\vec{\beta} = 1$,
(iii) $(\sum_{i=1}^n(\overline{p_{di}} + \underline{p_{di}})/2) \times \vec{\gamma} + (\sum_{i=1}^n(\overline{p_{di}} - \underline{p_{di}})/2) \times \vec{\beta} \leq \overline{p_g}$,
(iv) $(\sum_{i=1}^n(\overline{p_{di}} + \underline{p_{di}})/2) \times \vec{\gamma} + (\sum_{i=1}^n(-\overline{p_{di}} + \underline{p_{di}})/2) \times \vec{\beta} \geq \underline{p_g}$.

The $(\vec{\gamma}, \vec{\beta})$-determined controller generalizes the $\vec{\beta}$-determined controller (just set $\vec{\gamma} = \vec{\beta}$) and it is easy to see that the maximum power flow on the lines over all demand vectors, given a $(\vec{\gamma}, \vec{\beta})$-determined controller can be computed similarly to (18) as follows:

$$\begin{aligned} \max_{\underline{p_d} \leq \vec{p_d} \leq \overline{p_d}} |f_k| = & \left| \sum_{i=1}^n \frac{(\overline{p_{di}} + \underline{p_{di}})}{2} \mathbf{B}_k \vec{w_i}^{(\gamma)} \right| \\ & + \sum_{i=1}^n \frac{(\overline{p_{di}} - \underline{p_{di}})}{2} |\mathbf{B}_k \vec{w_i}^{(\beta)}|. \end{aligned} \tag{20}$$

Optimal $(\vec{\gamma}, \vec{\beta})$-determined controller can be found similar to the optimal $\vec{\beta}$-determined controller using an optimization similar to (19) with a few small changes:

$$\begin{aligned} \min_{\eta,\vec{\gamma},\vec{\beta},\vec{f}} \quad & \eta && (21)\\ \text{s.t.} \quad & \text{(i-iv) in Definition 4,} \\ & \vec{f} = |\mathbf{BW}^{(\gamma)}(\overline{\vec{p_d}} + \underline{\vec{p_d}})/2| + |\mathbf{BW}^{(\beta)}|(\overline{\vec{p_d}} - \underline{\vec{p_d}})/2, \\ & \vec{f} \leq \eta \overline{f}. \end{aligned}$$

Again, as in the $\vec{\beta}$-determined controller case, the optimal value of (21) determines if the optimal $(\vec{\gamma}, \vec{\beta})$-determined controller is reliable or not. Hence, the grid operator can use (21) to efficiently determine the secondary controllability of the grid, albeit obtaining false negatives in some cases.

In Section 8, we numerically evaluate the performance of the controllers introduced in this section. Before that, however, we demonstrate that these controllers can be used to efficiently provide lower bounds on the maximum scale of a MAD attack for which the grid remains secondary controllable.

7 αD-ROBUSTNESS

Power grids are required to withstand single equipment failures (e.g., lines, generators, and transformers) with no interruptions in their operation (a.k.a. $N-1$ standard) [27]. Following $N-1$ standard, we define a new standard for the grid operation to ensure its robustness against MAD attacks called αD standard. It requires grid operators to either prevent line overloads (as in Section 5) or be able to clear them (as in Section 6) after a MAD attack by an adversary that can change the demands by at most α fraction at each node.[4] We call a grid that conforms with this standard, αD-robust.

In this section, for a given grid, we are interested in finding the maximum α such that the grid is αD-robust. We denote this value by $\alpha^{\max}$. Since ensuring that line overloads can be cleared during the secondary control is less restrictive than preventing them after the primary control, *we mainly focus on finding the maximum α such that the grid is αD-robust based on its ability to clear line overloads after the secondary control (i.e., grid's secondary controllability).*

As we described in the previous section, verifying the secondary controllability of the grid for a given upper and lower limits on the demands is hard. Hence, we cannot expect to find the $\alpha^{\max}$ efficiently. Nevertheless, in the next two subsections, we develop efficient methods for obtaining upper and lower bounds on $\alpha^{\max}$.

7.1 Upper Bound

Assume $\vec{p_d}^\dagger$ denotes the vector of predicted demand values. For a given α, the demand vector $\vec{p_d}$ resulted by a MAD attack will be bounded by $(1-\alpha)\vec{p_d}^\dagger \leq \vec{p_d} \leq (1+\alpha)\vec{p_d}^\dagger$. Now if a grid is αD-robust, it should particularly be robust against the maximum demand attack. Hence, finding the maximum α for which the grid can handle the maximum demand attack provides an upper bound for $\alpha^{\max}$. Such α can be found efficiently by an LP:

$$\begin{aligned} \max_{\alpha,\vec{p_d},\vec{p_g},\vec{f},\vec{\theta}} \quad & \alpha && (22)\\ \text{s.t.} \quad & (1), (2), (3), (4), (5), \\ & \vec{p_d} = (1+\alpha)p_d^\dagger, \\ & \vec{p} = \vec{p_g} - \vec{p_d}. \end{aligned}$$

Proposition 4. *Assume $\hat{\alpha}$ denotes the optimal value of (22), then $\alpha^{\max} \leq \hat{\alpha}$.*

The optimal value of (22) provides a good upper bound for $\alpha^{\max}$ and can be computed efficiently. One can also consider $\vec{p_d} = (1-\alpha)p_d^\dagger$ to obtain another upper bound. However, if we set $\vec{p_d} = (1-\alpha)p_d^\dagger$ in (22) instead of $\vec{p_d} = (1+\alpha)p_d^\dagger$, it is easy to see that its optimal solution will be $\alpha = 1$. Hence, the case of $\vec{p_d} = (1-\alpha)p_d^\dagger$ only provides a trivial upper bound of $\alpha^{\max} \leq 1$ (assuming $\underline{p_g} = 0$).

In the next subsection, we provide algorithms to find lower bounds for α based on the controllers developed in Section 6.2.

7.2 Lower Bound

To find a lower bound for $\alpha^{\max}$, we use the controllers in Section 6.2 to limit the secondary controller's ability to change the generators' operating points. Limiting the secondary controller's ability allows us to efficiently approximate the maximum α, but because of this limitation, we only obtain lower bounds for $\alpha^{\max}$.

First, assume that we limit the secondary controller to the $\vec{\beta}$-controller for a fixed $\vec{\beta}$. We show that in this case the maximum α can be found by solving a single LP. Assume $\vec{p_g}^*$ is the optimal solution to (22) with value $\hat{\alpha}$ and set $\vec{\beta} = \vec{p_g}^*/\|\vec{p_g}^*\|_1$ (i.e., the controller only scales down the generation compared to the maximum demand case). Using (18), we show that the optimal value of the following LP gives a lower bound for $\alpha^{\max}$:

$$\begin{aligned} \max_{\alpha,\vec{f}} \quad & \alpha && (23)\\ \text{s.t.} \quad & (1+\alpha)(\sum_{i=1}^{n} p_{di}^\dagger) \times \vec{\beta} \leq \overline{p_g}, \\ & (1-\alpha)(\sum_{i=1}^{n} p_{di}^\dagger) \times \vec{\beta} \geq \underline{p_g}, \\ & \vec{\beta} = \vec{p_g}^*/\|\vec{p_g}^*\|_1, \\ & \vec{f} = |\mathbf{BW}^{(\beta)}\vec{p_d}^\dagger| + |\mathbf{BW}^{(\beta)}|(\alpha\vec{p_d}^\dagger), \\ & |f_{ij}| \leq \overline{f_{ij}}, \quad \forall (i,j) \in E. \end{aligned}$$

Proposition 5. *The optimal solution α^* of (23) can be found in polynomial time using LP. Moreover, $\alpha^* \leq \alpha^{\max}$.*

Optimization (23) allows us to efficiently compute a lower bound for $\alpha^{\max}$. However, similar to Section 6.2, instead of fixing $\vec{\beta}$, we can compute a $\vec{\beta}$ that results in the largest possible lower bound. Due to the nonlinearity of the problem, however, we cannot optimize $\vec{\beta}$ and found maximum α in (23) simultaneously. The idea is to fix α, compute the optimal $\vec{\beta}$ and η using (19), then update α

[4] This is based on the assumption that the IoT bots are uniformly distributed in an area. Therefore, an adversary's ability to change the demands is determined by the initial demand at each node.

Module 1: Lower Bound on $\alpha^{\max}$ using $(\vec{\gamma}, \vec{\beta})$-determined Controllers

```
Input: $G$, $\lambda$
$\alpha^{(0)} = \hat{\alpha}$
flag = 1
i = 0
while flag do
    flag = 0
    Compute the optimal value $\eta$, $\vec{\gamma}$, and $\vec{\beta}$ of (21) for
    $\overline{p_d} = (1 + \alpha^{(i)})\vec{p_d}^{\dagger}$ and $\underline{p_d} = (1 - \alpha^{(i)})\vec{p_d}^{\dagger}$
    Set $\alpha^{(i+1)} = \alpha^{(i)} + \lambda(1 - \eta)$
    if $|\alpha^{(i+1)} - \alpha^{(i)}| > 0.001$ then
        flag = 1
        $i = i + 1$
return $\alpha^{(\gamma,\beta)} := \alpha^{(i)}$, $\vec{\gamma}$, and $\vec{\beta}$
```

using η and repeat the process until α does not change by much. As in Section 6.2, we can use the $(\vec{\gamma}, \vec{\beta})$-determined controller instead of the $\vec{\beta}$-determined controller to improve the obtained lower bound. The method is summarized in Module 1. When $\gamma = \beta$, Module 1 provides a lower bound on $\alpha^{\max}$ like $\alpha^{(\beta)}$ based on $\vec{\beta}$-determined controllers.

Notice that λ in Module 1 should be set such that updates to α at each iteration are neither too large that the module falls into a loop, nor are too small that it takes a long time to converge.

Proposition 6. *When $\gamma = \beta$, for a good λ, Module 1 converges to an $\alpha^{(\beta)}$ value such that $\alpha^{(\beta)} \leq \alpha^{\max}$. Moreover, $\alpha^* \leq \alpha^{(\beta)}$. (Recall that α^* is the optimal solution of (23).)*

Proposition 7. *For a good λ, Module 1 converges to an $\alpha^{(\gamma,\beta)}$ value such that $\alpha^{(\gamma,\beta)} \leq \alpha^{\max}$. Moreover, $\alpha^{(\beta)} \leq \alpha^{(\gamma,\beta)}$.*

In the next section, we numerically compare the upper bound $\hat{\alpha}$, and lower bounds α^*, $\alpha^{(\beta)}$, and $\alpha^{(\gamma,\beta)}$ with $\alpha^{\max}$ in order to demonstrate the tightness of these bounds in approximating $\alpha^{\max}$.

8 Numerical Results

In this section, we first numerically evaluate the performance of SAFE and IMMUNE Algorithms developed in Section 5. Then, we numerically evaluate the accuracy of the upper and lower bounds developed in Section 7 in approximating the maximum α such that the grid is αD-robust (i.e., $\alpha^{\max}$).

8.1 Simulations Setup

For solving LP, we use *CVX*, a package for specifying and solving convex programs [41], [42]. For computing the optimal power flow part of the IMMUNE Algorithm, we use *MATPOWER* [43] which is a MATLAB based library for computing the power flows. We also exploit the power system test cases available with this library for our simulations. In particular, we use the IEEE 14-bus, 30-bus, and 57-bus test systems, and the New England 39-bus system.

TABLE 1: Performance Evaluation of SAFE and IMMUNE Algorithms on the New England 39-bus system. Cost values are in $\$/hr$. Numbers in parenthesis indicate the number of iterations took the IMMUNE Algorithm to converge.

α	OPF	SAFE	IMMUNE	IMMUNE-0.95	IMMUNE-0.9
0.09	41264	-	43434 (7)	43805 (4)	43859 (3)
0.08	41264	43628	42394 (8)	42431 (3)	42982 (3)
0.07	41264	42665	41773 (5)	41991 (3)	42405 (3)
0.06	41264	42050	41492 (4)	41698 (3)	41534 (2)
0.05	41264	41668	41339 (10)	41421 (3)	41419 (2)

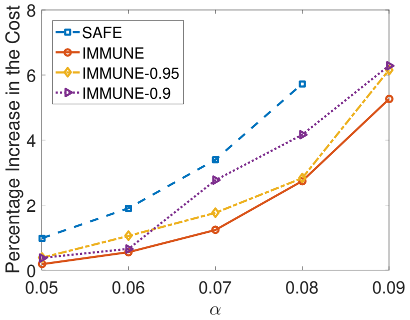

Fig. 7: Percentage increase in operating cost of the grid in order to make it robust against MAD attacks obtained by SAFE and IMMUNE Algorithms versus the magnitude of the attack (α) in New England 39-bus system.

The line capacities are only provided for the IEEE 30-bus and New England 39-bus systems. Hence, for the other two systems, we set the capacities ourselves in two-different ways: (i) following [9] for each line we set $\overline{f_i} = \max\{1.2|f_i^{\dagger}|, \text{median}(|\vec{f}^{\dagger}|)\}$, and (ii) set $\overline{f_i} = 1.1\max(|\vec{f}^{\dagger}|)$, in which $\vec{f}^{\dagger}$ are the power flows given the default supply and demand values in the test systems. When the first method is used for determining the capacities, it is indicated by (f) in front of the grid name, and when the second method is used, it is indicated by (u) (e.g., see Table 3).

8.2 Primary Control

In this subsection, we evaluate the performance of SAFE and IMMUNE Algorithms on NEW England 39-bus and IEEE 30-bus systems. We assume that $(1-\alpha)p_{di}^{\dagger} \leq p_{di} \leq (1+\alpha)p_{di}^{\dagger}$ and consider different α values to capture attacks with different magnitudes (which depends on the number of controlled bots by an adversary).

Table 1 compares the performance of SAFE and three variations of the IMMUNE Algorithm for different α values. Recall from Section 5.2 that IMMUNE-0.95 and IMMUNE-0.9 are similar to the IMMUNE Algorithm but apply more aggressive updates on the capacities in each iteration of the algorithm. This, as mentioned in Section 5.2 and demonstrated numerically here in Table 1, results in faster convergence of the algorithm. Since the OPF problem does not consider the robustness of the grid against MAD attacks, its value is independent of the magnitude of an expected attack (α).

As can be seen in Table 1 and as we expected, the grid needs to be operated in a non-optimal operating point in order to be robust against MAD attacks. The required percentage increase in the operating cost of the grid obtained by the SAFE and IMMUNE Algorithms versus α are presented in Fig. 7. IMMUNE Algorithm results in the least amount of increase in the operating cost. However, since as demonstrated in Table 1, IMMUNE Algorithm takes longer that IMMUNE-0.95 and IMMUNE-0.9 Algorithms to

TABLE 2: Performance Evaluation of SAFE and IMMUNE Algorithms on the IEEE 30-bus system. Cost values are in $\$/hr$. Numbers in parenthesis indicate the number of iterations took the algorithm to converge.

α	OPF	SAFE	IMMUNE
0.31	565.2	-	- (3)
0.3	565.2	614.8	- (4)
0.28	565.2	571.6	569.6 (3)
0.26	565.2	565.32	565.22 (2)
0.22	565.2	565.2	565.2 (1)

converge, the system operator may prefer to use IMMUNE-0.95 which performs approximately as well as the IMMUNE Algorithm but converges faster. Notice that due to nonconvexity of the problem, a more aggressive update rule may not necessarily result in a costlier operating point, as we see here that IMMUNE-0.9 results in a lower operating cost than IMMUNE-0.95 for $\alpha = 0.06$.

It can also be seen that SAFE Algorithm performs relatively well in finding a robust operating point of the grid much faster than all variations of IMMUNE Algorithm (recall from Section 5.3 that SAFE Algorithm requires only to solve a single LP). However, it may become infeasible for higher magnitude attacks (in this case for $\alpha = 0.09$).

We repeated the simulations on the IEEE 30-bus system. The results are presented in Table 2. First, it can be seen that the IEEE 30-bus system can be protected against much stronger attacks ($\alpha = 0.3$) which demonstrates that different grids may have different levels of robustness against MAD attacks (we will make a similar observation in the secondary control case in the next subsection). Unlike the New England 39-bus case, here the IMMUNE Algorithm does not converge for the strongest attack ($\alpha = 0.3$) rather than the SAFE Algorithm. This demonstrates that each of these algorithms may be useful in finding a robust operating point for the grid in different scenarios–besides their running time and optimality.

As can be seen in Table 2, in this case also, if the IMMUNE Algorithm converges, it converges to a lower cost operating point than the one obtained by the SAFE Algorithm. Here, the IMMUNE Algorithm converged within a few iterations. Therefore, there was no need to consider a faster variation of the IMMUNE Algorithm as in the New England 39-bus case.

Finally, it can be seen that for $\alpha = 0.31$, none of the algorithms can obtain a robust operating point for the grid. We show in the next subsection that this case can be handled by the secondary controller instead (assuming that lines can handle temporary overloads).

8.3 Secondary Control

In order to evaluate the performance of the controllers developed in Section 6.2, in this subsection, we focus on their performance in approximating $\alpha^{\max}$ as described in Section 7.

TABLE 3: Lower and upper bounds for $\alpha^{\max}$.

Test case	α^*	$\alpha^{(\beta)}$	$\alpha^{(\gamma,\beta)}$	$\alpha^{\max}$	$\hat{\alpha}$
`IEEE 14-bus(f)`	0.058	0.1649	0.1906	0.2117	0.2117
`IEEE 14-bus(u)`	0.950	1.0243	1.1454	1.1479	1.1479
`IEEE 30-bus`	0.214	0.2851	0.3126	0.37	0.3717
`NE 39-bus`	0.039	0.0796	0.0962	0.0962	0.0962
`IEEE 57-bus(f)`	0.024	0.0307	0.0311	< 0.09	0.2
`IEEE 57-bus(u)`	0.128	0.2396	0.2864	-	0.3468

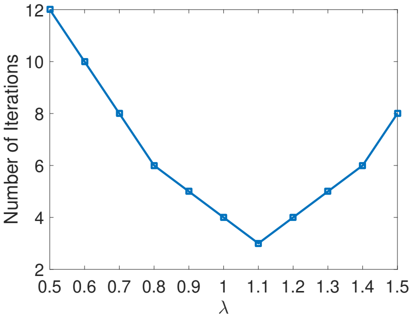

Fig. 8: Number of iterations in Module 1 before it converges versus its update step size λ in the IEEE 30-bus system.

Table 3 compares the maximum α obtained by different methods in several test cases. As can be seen and proved in Section 7, in all cases, $\alpha^* \leq \alpha^{(\beta)} \leq \alpha^{(\gamma,\beta)} \leq \alpha^{\max} \leq \hat{\alpha}$. Notice that for the IEEE 57-bus system, since the brute force search algorithm needs to solve (17) about 2^{42} times for each given α to determine the secondary controllability of the grid, we could not exactly determine $\alpha^{\max}$. However, in the case of IEEE 57-bus (f), after initial iterations of the brute force search algorithm, we could determine that the grid is not secondary controllable for $0.09 \leq \alpha$ as presented in the table.

It can be seen that $\hat{\alpha}$ provides a very close upper bound for $\alpha^{\max}$ most of the time (except in IEEE 57-bus (f)). And since it can be computed by a single LP, the numerical results suggest that it is an efficient and reliable way to find an upper bound for $\alpha^{\max}$. On the other hand, α^* that can also be computed efficiently by a single LP does not provide a very close lower bound in the test systems that we studied here. However, $\alpha^{(\beta)}$ and $\alpha^{(\gamma,\beta)}$ that require more time to be computed, provide much better lower bounds. In particular, in the case of New England 39-bus system $\alpha^{(\gamma,\beta)} = \hat{\alpha}$ which implies that $\alpha^{\max} = \alpha^{(\gamma,\beta)} = \hat{\alpha}$.

Although finding $\alpha^{(\beta)}$ and $\alpha^{(\gamma,\beta)}$ requires solving an LP in several iterations (as summarized in Module 1), the number of iterations can be minimized by selecting a good step size λ. For example, the number of iterations of Module 1 versus λ is presented in Fig. 8 in the IEEE 30-bus system. As can be seen, for the optimal λ (in this case $\lambda = 1.1$), the module converges in 3 iterations. Hence, it can find a good lower bound for α, as shown in Table 3, very efficiently and in polynomial time (since it solves a single LP at each iteration). A good λ can be found in practice heuristically after the first few iterations and observing the rate of changes.

Finally, as mentioned in Section 6, the secondary controllability becomes more important when the primary controller cannot prevent line overloads, but the overloads can be tolerated for a short period of time. An example of such scenario happens in IEEE 30-bus system and when $\alpha = 0.31$. As can be seen in Table 2, none of the SAFE and IMMUNE Algorithms can find a robust operating point for the grid in this case. However, as can be seen in Table 3, since this value is less that $\alpha^{\max} = 0.37$, any line overloads can be cleared by the secondary controller.

8.4 Open Questions

As we observed in the previous two subsections, different test systems demonstrate different levels of robustness against MAD attacks. For example, as can be seen in Table 3, the $\alpha^{\max}$ for the IEEE 30-bus system is 0.37, whereas this value for the New England 39-bus system is only 0.0962. This difference in robustness can be due to the structure of the network as well as the location of the generators and loads. Analytically studying such features and developing efficient algorithms to improve grid robustness by adding extra lines to a system or build future generators at certain locations would be interesting future research directions.

Another important observation from the numerical results is that the performance of the proposed algorithms varies in different test systems. For example, in the New England 39-bus system, the IMMUNE Algorithm successfully finds robust operating points for the generators for different α values, whereas in the IEEE 30-bus system the IMMUNE Algorithm may not converge for $\alpha = 0.3$. Moreover, as can be seen in Table 3, the approximation algorithms for estimating $\alpha^{\max}$ provide tight bounds for the New England 39-bus system, whereas the bounds are not tight for the IEEE 30-bus system. Hence, finding sufficient conditions on the structure and properties of a test case under which the approximation bounds are tight and the IMMUNE Algorithm is guaranteed to converge to a locally optimal solution would be important future research directions.

9 Conclusions

In this paper, we analyzed the effect of MAD attacks on power flows in detail and presented SAFE and IMMUNE algorithms for finding robust operating points for the generators during economic dispatch such that no lines are overloaded after automatic primary control response to any MAD attacks. Moreover, we demonstrated that in cases that temporary overloads can be tolerated, the system operator can approximately but efficiently verify in advance if line overloads can be cleared during the secondary control after any MAD attacks. Based on these two forms of defenses, we defined αD-robustness notion and demonstrated that upper and lower bounds on the maximum α for which the grid is αD-robust can be found efficiently and in polynomial time. We finally evaluated the performance of the developed algorithms and methods, and showed that they perform very well in practical test cases.

We believe that with universality and growth in the number of high-wattage IoT devices and smart thermostats, the probability of MAD attacks is increasing and there is an urgent need for more studies on the potential effects of these attacks and developing tools for grid protection. Our work provides the first methods for protecting the grid against potential line failures caused by newly discovered MAD attacks via IoT devices. However, our work can be extended in several directions. A natural direction is to extend the developed results to the AC power flow model. A more challenging research direction is to extend the methods to unit commitment phase of the grid operation. Since regular unit commitment problem is already a combinatorial problem, incorporating security constraints into that problem will be a challenging task and part of our future work.

In the worst-case scenario that the scale of a MAD attack is greater than grid robustness (i.e., adversary manipulates the demands by greater than $\alpha^{\max}$ factor), the grid operator may not be able to clear the possible line overloads in a timely manner. This can consequently force the overloaded lines to trip leading to more line overloads and a cascading failure in the system [3]. To prevent cascading failures in such scenarios, the grid operator may apply common control algorithms such as *optimal load-shedding* [44] or *power grid intentional islanding* [45]. However, since an adversary can suddenly decrease the demands after an initial increase in the demands, these control algorithms may not be effective in their classical form (e.g., sudden decrease in the demands after load-shedding may result in a critical increase in the frequency of the system). Hence, investigating ways to improve these control algorithms to protect the grid against MAD attacks in the worst-case scenarios is also part of our future work.

10 Omitted Proofs

Proof of Lemma 1: First, notice that $1/R_i$ is the rate with which generator i increases its generation to compensate for the extra demand. Hence, t_i denotes the time that generator i reaches its maximum capacity if the total supply does not meet the demand before t_i. Accordingly, generators reach their maximum capacity in the order of their t_i values from smallest to largest. Using this, it is easy to see that S_i is the total change in the generation at time t_i. Therefore, if $S_i < S_{\Delta p_d}$, then generators 1 to i will reach their maximum capacities before supply meets the total demand. Moreover, since $S_{\Delta p_d} \leq S_{i+1}$, generators $i+1, \dots, n$ do not reach their capacities and each contribute according to their droop characteristic to compensate for the remaining $S_{\Delta p_d} - \sum_{l=1}^{i} (\overline{p_{gl}} - p_{gl})$. □

Proof of Lemma 3: First, notice that for each line $(i, j) \in E$ and in each iteration of the IMMUNE Algorithm, c_{ij} is not increasing. To see this, assume c_{ij} changes in the l^{th} iteration, and c_{ij}^{old} and c_{ij}^{new} denote the value of c_{ij} before and after the change, respectively. Since c_{ij} is changed, it means that $\overline{f_{ij}} < |f_{ij}| + \Delta f_{ij}^{\max}$. On the other hand, $|f_{ij}| \leq c_{ij}^{\text{old}}$. Hence, $\overline{f_{ij}} < c_{ij}^{\text{old}} + \Delta f_{ij}^{\max}$ or $\overline{f_{ij}} - \Delta f_{ij}^{\max} < c_{ij}^{\text{old}}$. Since $c_{ij}^{\text{new}} = \overline{f_{ij}} - \Delta f_{ij}^{\max}$, therefore $c_{ij}^{\text{new}} < c_{ij}^{\text{old}}$.

On the other hand, from (11), it is easy to verify that after each iteration $\overline{f_{ij}} - \widehat{\Delta f_{ij}} \leq c_{ij}$. Hence, c_{ij}s cannot get smaller than the fixed values $\overline{f_{ij}} - \widehat{\Delta f_{ij}}$ and since (12) is feasible, the OPF problem remains feasible after each iteration of the IMMUNE algorithm. Now since c_{ij}s are non-increasing and limited by lower bounds, the algorithm is guaranteed to remain feasible and converge to a local optimum solution. □

Proof of Lemma 4: In each iteration of the IMMUNE algorithm, at least for a single line (i, j), the c_{ij} will be updated. Otherwise, the algorithm should terminate (either converges or become infeasible). On the other hand, since $\widehat{\Delta f_{ij}}$ is the maximum possible flow change on line (i, j), the c_{ij} cannot get smaller than $\overline{f_{ij}} - \widehat{\Delta f_{ij}}$. Hence, since the updates are discrete, in the worst case that only a single capacity is updated by a single unit at each iteration, the algorithm can take at most $\sum_{(i,j)\in E}\lceil\widehat{\Delta f_{ij}}\rceil$ iterations to terminate. □

Proof of Lemma 5: Assume $\vec{p_d}^{(1)}, \vec{p_d}^{(2)}, \ldots, \vec{p_d}^{(2^n)}$ denote all possible extreme demand vectors. Now assume that for each extreme demand vector $\vec{p_d}^{(i)}$, there exists an operating vector $\vec{p_g}^{(i)}$ for generators that satisfies the secondary control conditions. We prove that for all demand vectors $\vec{p_d}$ within the upper and lower limits also there exists an operating vector $\vec{p_g}$ that satisfies all the secondary controller conditions. Since the space of all the demand vectors is convex, each demand vector $\vec{p_d}$ within the upper and lower limits can be written as a convex combination of the extreme points such as $\vec{p_d} = \sum_{i=1}^{2^n} \beta_i \vec{p_d}^{(i)}$ in which $\forall i : \beta_i \geq 0$ and $\sum_{i=1}^{2^n} \beta_i = 1$. We show that $\vec{p_g} = \sum_{i=1}^{2^n} \beta_i \vec{p_g}^{(i)}$ satisfies all the secondary controller conditions. First, since $\vec{p_g}$ is a convex combination of $\vec{p_g}^{(i)}$s and they are within generators upper and lower limits, so is $\vec{p_g}$. Second, it is easy to see that $\vec{1}^T(\vec{p_g} - \vec{p_d}) = \sum_{i=1}^{2^n} \beta_i \vec{1}^T(\vec{p_g}^{(i)} - \vec{p_d}^{(i)}) = \sum_{i=1}^{2^n} \beta_i 0 = 0$. Finally, based on our assumptions, for each i: $-\overline{f} \leq \mathbf{B}(\vec{p_g}^{(i)} - \vec{p_d}^{(i)}) \leq \overline{f}$. Hence, $\mathbf{B}(\vec{p_g} - \vec{p_d}) = \sum_{i=1}^{2^n} \beta_i \mathbf{B}(\vec{p_g}^{(i)} - \vec{p_d}^{(i)}) \leq \sum_{i=1}^{2^n} \beta_i \overline{f} = \overline{f}$. Similarly, $-\overline{f} \leq \mathbf{B}(\vec{p_g} - \vec{p_d})$. Therefore, $\vec{p_g}$ satisfies all the constraints of the secondary controller problem. The reverse can also be similarly proved using contradiction method. □

Proof of Proposition 2: If there exists a vector $\vec{\beta}$ that the $\vec{\beta}$-determined controller is reliable, then for any feasible demand vector $\vec{p_d}$, vector of operating points $\vec{p_g} = (\sum_{i=1}^{n} p_{di}) \times \vec{\beta}$ satisfies the demands (i.e., $\vec{1}^T(\vec{p_g} - \vec{p_d}) = 0$) and $|\vec{f}| = |\mathbf{B}(\vec{p_g} - \vec{p_d})| \leq \overline{f}$. Therefore, the grid is secondary controllable. □

Proof of Lemma 6: From the definition of $\vec{w_i}^{(\beta)}$ vectors, it is easy to verify that for a demand vector $\vec{p_d}$, the power flow on line e_k can be computed as $f_k = \sum_{i=1}^{n} p_{di} \mathbf{B}_k \vec{w_i}^{(\beta)}$. For $|f_k|$ to be maximized, each p_{id} should be either equal to $\underline{p_{di}}$ or $\overline{p_{di}}$ based on signs of $\mathbf{B}_k \vec{w_i}^{(\beta)}$ and f_k. On the other hand, it is easy to see that $\underline{p_{di}} = \frac{(\overline{p_{di}}+\underline{p_{di}})}{2} - \frac{(\overline{p_{di}}-\underline{p_{di}})}{2}$ and $\overline{p_{di}} = \frac{(\overline{p_{di}}+\underline{p_{di}})}{2} + \frac{(\overline{p_{di}}-\underline{p_{di}})}{2}$. So by considering only $p_{di} \in \{\overline{p_{di}}, \underline{p_{di}}\}$, f_k can be computed as follows:

$$
\begin{aligned}
f_k &= \sum_{i=1}^{n} p_{di} \mathbf{B}_k \vec{w_i}^{(\beta)} = \sum_{i=1}^{n} \Big(\frac{(\overline{p_{di}} + \underline{p_{di}})}{2} \pm \frac{(\overline{p_{di}} - \underline{p_{di}})}{2}\Big) \mathbf{B}_k \vec{w_i}^{(\beta)} \\
&= \sum_{i=1}^{n} \frac{(\overline{p_{di}} + \underline{p_{di}})}{2} \mathbf{B}_k \vec{w_i}^{(\beta)} + \sum_{i=1}^{n} \Big(\pm \frac{(\overline{p_{di}} - \underline{p_{di}})}{2}\Big) \mathbf{B}_k \vec{w_i}^{(\beta)}.
\end{aligned}
$$

From the equation above, it can be seen that the first part is fixed but the second part can be selected based on the sign of the first part in order to maximize $|f_k|$. Hence, it is easy to see that maximum value of $|f_k|$ is:

$$
\begin{aligned}
\max_{\underline{p_d} \leq \vec{p_d} \leq \overline{p_d}} |f_k| = &\left| \sum_{i=1}^{n} \frac{(\overline{p_{di}} + \underline{p_{di}})}{2} \mathbf{B}_k \vec{w_i}^{(\gamma)} \right| \\
&+ \sum_{i=1}^{n} \frac{(\overline{p_{di}} - \underline{p_{di}})}{2} |\mathbf{B}_k \vec{w_i}^{(\beta)}|.
\end{aligned}
$$

□

Proof of Proposition 3: In order to solve (19) using LP, one can define auxiliary vector $\vec{u}$ and matrix $\mathbf{Q}$ and replace the constraint $\vec{f} = |\mathbf{BW}^{(\beta)}(\overline{p_d} + \underline{p_d})/2| + |\mathbf{BW}^{(\beta)}|(\overline{p_d} - \underline{p_d})/2$ in (19) with following set of inequalities:

$$
\begin{aligned}
&\vec{f} = \vec{u} + \mathbf{Q}(\overline{p_d} - \underline{p_d})/2, \\
&\vec{u} \geq \mathbf{BW}^{(\beta)}(\overline{p_d} + \underline{p_d})/2, \\
&\vec{u} \geq -\mathbf{BW}^{(\beta)}(\overline{p_d} + \underline{p_d})/2, \\
&\mathbf{Q} \geq \mathbf{BW}^{(\beta)}, \;\; \mathbf{Q} \geq -\mathbf{BW}^{(\beta)},
\end{aligned}
$$

in which the matrix inequalities are entry by entry. Now it is easy to verify that since the optimization minimize η and $\vec{f} \leq \eta\overline{f}$, in the optimal solution $\vec{f}$ will be minimized and therefore $\vec{u}$ and $\mathbf{Q}$ will be equal to $|\mathbf{BW}^{(\beta)}(\overline{p_d} + \underline{p_d})/2|$ and $|\mathbf{BW}^{(\beta)}|$, respectively. Hence using the above transformation, (19) can be solved using LP. It can be seen that if the optimal solution η^* to (19) is less than or equal to 1, then since $\vec{f}$ is equal to the maximum power flow on the lines over all possible demand vectors (and corresponding generation operating points obtained by the $\vec{\beta}^*$-determined controller) and $\vec{f} \leq \eta^*\overline{f} \leq \overline{f}$, the $\vec{\beta}^*$-controller is reliable. Hence, the grid is secondary controllable. □

Proof of Proposition 4: Since in optimization (22) only the maximum demand case (i.e., $\vec{p_d} = (1+\alpha)\vec{p_d}^{\dagger}$) is being verified to be satisfiable by the generators with no line overloads, the optimal solution of (22) only provides an upper bound for $\alpha^{\max}$. □

Proof of Proposition 5: Using (18), it can be verified that the maximum power flow on a line (i,j) over all the demand vectors and corresponding generation vector determined by the $\vec{\beta}$-determined controller is equal to $|\mathbf{BW}^{(\beta)}\vec{p_d}^{\dagger}| + |\mathbf{BW}^{(\beta)}|(\alpha\vec{p_d}^{\dagger})$. Hence, optimization (23) maximizes α such that the grid is αD-robust using the specified $\vec{\beta}$-determined controller. On the other hand, since the operating points of the generators are limited to the operating points obtained by the specified $\vec{\beta}$-determined controller, it is obvious that demand vectors that are controllable by this controller are a subset of all controllable vectors. Hence, α^* only provides a lower bound for $\alpha^{\max}$. Finally, it is also easy to see that similar to the technique presented in the proof of Proposition 3, optimization (23) can be solved using LP and therefore α^* can be computed in polynomial time. □

Proof of Proposition 6: At each iteration, if $\alpha^{(i)} > \alpha^{\max}$, then the solution η to (19) would be greater than 1. Hence, if λ is small enough, $0 \leq \alpha^{(i+1)} = \alpha^{(i)} + \lambda(1-\eta) \leq \alpha^{(i)}$. Similarly, it can be shown that if $\alpha^{(i)} < \alpha^{\max}$, then $\alpha^{(i+1)} > \alpha^{(i)}$. On the other hand, for $\alpha^{(i)} = \alpha^{\max}$, the solution η to (19) would be zero and $\alpha^{(i)} = \alpha^{(i+1)} = \alpha^{\max}$. Hence, $\alpha^{\max}$ is the only absorbing point for this algorithm which it converges to (if λ is small enough). □

Proof of Proposition 7: The convergence proof is similar to the proof of Proposition 6. It is also easy to see that since $\vec{\beta}$-determined controllers are a special case of $(\vec{\gamma}, \vec{\beta})$-determined controllers, $\alpha^{(\beta)} \leq \alpha^{(\gamma,\beta)}$. □

ACKNOWLEDGEMENTS

This work is supported in part by the Siebel Energy Institute, in part by the National Science Foundation under grants DMS-1736417, ECCS-1824710, and CNS-1553437, and in part by the Office of Naval Research YIP Award.

REFERENCES

[1] NERC, "Analysis of the cyber attack on the Ukrainian power grid," 2016, http://www.nerc.com/pa/CI/ESISAC/Documents/E-ISAC_SANS_Ukraine_DUC_18Mar2016.pdf. Accessed: Jan. 2018.

[2] N. S. Malik and R. Collins, "The cyberattack that crippled gas pipelines is now hitting another industry," 2018, https://www.bloomberg.com/news/articles/2018-04-04/cyberattack-bleeds-into-utility-space-as-duke-sees-billing-delay. Accessed: June. 2018.

[3] S. Soltan, P. Mittal, and H. V. Poor, "BlackIoT: IoT botnet of high wattage devices can disrupt the power grid," in *Proc. USENIX Security'18*, Aug. 2018.

[4] A. Dabrowski, J. Ullrich, and E. R. Weippl, "Grid shock: Coordinated load-changing attacks on power grids: The non-smart power grid is vulnerable to cyber attacks as well," in *Proc. ACM ACSAC'17*, Dec. 2017.

[5] "How hacked water heaters could trigger mass blackouts," *Wired magazine*, Aug. 2018, https://www.wired.com/story/water-heaters-power-grid-hack-blackout/.

[6] "Your smart air conditioner could help bring down the power grid," *CNET*, Aug. 2018, https://www.wired.com/story/water-heaters-power-grid-hack-blackout/.

[7] I. Dobson, "Cascading network failure in power grid blackouts," *Encyclopedia of Systems and Control*, pp. 105–108, 2015.

[8] S. Soltan, D. Mazauric, and G. Zussman, "Analysis of failures in power grids," *IEEE Trans. Control Netw. Syst.*, vol. 4, no. 3, pp. 288–300, 2017.

[9] H. Cetinay, S. Soltan, F. A. Kuipers, G. Zussman, and P. Van Mieghem, "Analyzing cascading failures in power grids under the AC and DC power flow models," in *Proc. IFIP Performance'17*, Nov. 2017.

[10] D. Bienstock, *Electrical Transmission System Cascades and Vulnerability: An Operations Research Viewpoint*. SIAM, 2016.

[11] B. Carreras, V. Lynch, I. Dobson, and D. Newman, "Critical points and transitions in an electric power transmission model for cascading failure blackouts," *Chaos*, vol. 12, no. 4, pp. 985–994, 2002.

[12] J. Song, E. Cotilla-Sanchez, G. Ghanavati, and P. D. Hines, "Dynamic modeling of cascading failure in power systems," *IEEE Trans. Power Syst.*, vol. 31, no. 3, pp. 2085–2095, 2016.

[13] L. Garcia, F. Brasser, M. H. Cintuglu, A.-R. Sadeghi, O. Mohammed, and S. A. Zonouz, "Hey, my malware knows physics! attacking plcs with physical model aware rootkit," in *Proc. NDSS'17*, 2017.

[14] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," *ACM Trans. Inf. Syst. Secur.*, vol. 14, no. 1, p. 13, 2011.

[15] G. Dán and H. Sandberg, "Stealth attacks and protection schemes for state estimators in power systems," in *Proc. IEEE SmartGridComm'10*, 2010.

[16] S. Li, Y. Yılmaz, and X. Wang, "Quickest detection of false data injection attack in wide-area smart grids," *IEEE Trans. Smart Grid*, vol. 6, no. 6, pp. 2725–2735, 2015.

[17] J. Kim, L. Tong, and R. J. Thomas, "Subspace methods for data attack on state estimation: A data driven approach." *IEEE Trans. Signal Process.*, vol. 63, no. 5, pp. 1102–1114, 2015.

[18] S. Soltan, M. Yannakakis, and G. Zussman, "Joint cyber and physical attacks on power grids: Graph theoretical approaches for information recovery," in *Proc. ACM SIGMETRICS'15*, June 2015.

[19] D. Bienstock and M. Escobar, "Computing undetectable attacks on power grids," *ACM PER*, vol. 45, no. 2, pp. 115–118, 2017.

[20] Z. Li, M. Shahidehpour, A. Alabdulwahab, and A. Abusorrah, "Bilevel model for analyzing coordinated cyber-physical attacks on power systems," *IEEE Trans. Smart Grid*, vol. 7, no. 5, pp. 2260–2272, 2016.

[21] R. Deng, P. Zhuang, and H. Liang, "CCPA: Coordinated cyber-physical attacks and countermeasures in smart grid," *IEEE Trans. Smart Grid*, vol. 8, no. 5, pp. 2420–2430, 2017.

[22] J. Zhang and L. Sankar, "Physical system consequences of unobservable state-and-topology cyber-physical attacks," *IEEE Trans. Smart Grid*, vol. 7, no. 4, 2016.

[23] Y. Yuan, Z. Li, and K. Ren, "Modeling load redistribution attacks in power systems," *IEEE Trans. Smart Grid*, vol. 2, no. 2, pp. 382–390, 2011.

[24] A.-H. Mohsenian-Rad and A. Leon-Garcia, "Distributed internet-based load altering attacks against smart power grids," *IEEE Trans. Smart Grid*, vol. 2, no. 4, pp. 667–674, 2011.

[25] S. Amini, F. Pasqualetti, and H. Mohsenian-Rad, "Dynamic load altering attacks against power system stability: Attack models and protection schemes," *IEEE Trans. Smart Grid*, vol. 9, no. 4, pp. 2862–2872, 2018.

[26] Y. Dvorkin and S. Garg, "IoT-enabled distributed cyber-attacks on transmission and distribution grids," in *Proc. NAPS'17*, Sept 2017.

[27] A. J. Wood and B. F. Wollenberg, *Power generation, operation, and control*. John Wiley & Sons, 2012.

[28] S. H. Low, "Convex relaxation of optimal power flow-part I: Formulations and equivalence," *IEEE Trans. Control Netw. Syst.*, vol. 1, no. 1, pp. 15–27, 2014.

[29] A. Castillo and R. P. ONeill, "Survey of approaches to solving the ACOPF," *US FERC Tech. Rep*, 2013.

[30] A. Monticelli, M. Pereira, and S. Granville, "Security-constrained optimal power flow with post-contingency corrective rescheduling," *IEEE Trans. Power Syst.*, vol. 2, no. 1, pp. 175–180, 1987.

[31] D. Bienstock, M. Chertkov, and S. Harnett, "Chance-constrained optimal power flow: Risk-aware network control under uncertainty," *Siam Rev.*, vol. 56, no. 3, pp. 461–495, 2014.

[32] R. Bapat, *Graphs and matrices*. Springer, 2010.

[33] Federal Energy Regulatory Commission and others, *Energy Primer, a Handbook of Energy Market Basics*, 2012.

[34] European Network of Transmission System Operators for Electricity (ENTSOE), "Continental europe operation handbook," 2004, https://www.entsoe.eu/publications/system-operations-reports/operation-handbook/Pages/default.aspx. Accessed: Jan. 2018.

[35] J. Machowski, J. Bialek, and J. R. Bumby, *Power system dynamics and stability*. John Wiley & Sons, 1997.

[36] R. Hebner, J. Beno, and A. Walls, "Flywheel batteries come around again," *IEEE spectrum*, vol. 39, no. 4, pp. 46–51, 2002.

[37] J. E. Falk, "A linear max-min problem," *Mathematical Programming*, vol. 5, no. 1, pp. 169–188, 1973.

[38] J. F. Bard, *Practical bilevel optimization: algorithms and applications*. Springer, 1998.

[39] W. F. Bialas and M. H. Karwan, "Two-level linear programming," *Management science*, vol. 30, no. 8, pp. 1004–1020, 1984.

[40] P. Hansen, B. Jaumard, and G. Savard, "New branch-and-bound rules for linear bilevel programming," *SIAM J. Sci. Comput.*, vol. 13, no. 5, pp. 1194–1217, 1992.

[41] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," http://cvxr.com/cvx, Mar. 2014.

[42] ——, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, ser. Lecture Notes in Control and Information Sciences, V. Blondel, S. Boyd, and H. Kimura, Eds. Springer-Verlag Limited, 2008, pp. 95–110, http://stanford.edu/~boyd/graph_dcp.html.

[43] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, "MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education," *IEEE Trans. Power Syst.*, vol. 26, no. 1, pp. 12–19, 2011.

[44] D. Bienstock, "Optimal control of cascading power grid failures," in *Proc. IEEE CDC-ECC'11*, Dec. 2011.

[45] S. Soltan, M. Yannakakis, and G. Zussman, "Doubly balanced connected graph partitioning," in *Proc. ACM-SIAM SODA'17*, Jan. 2017.

Saleh Soltan (M'15) is a postdoctoral research associate in the department of Electrical Engineering at Princeton University. In 2017, he obtained a Ph.D. degree in Electrical Engineering from Columbia University. He received B.S. degrees in Electrical Engineering and Mathematics (double major) from Sharif University of Technology, Iran in 2011 and the M.S. degree in Electrical Engineering from Columbia University in 2012. He is the Gold Medalist of the 23rd National Mathematics Olympiad in Iran in 2005 and the recipient of Columbia University Electrical Engineering Armstrong Memorial Award in 2012 and Jury Award in 2018.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TNSE.2019.2922131, IEEE Transactions on Network Science and Engineering

Prateek Mittal is an Associate Professor in the Department of Electrical Engineering at Princeton University. He obtained his Ph.D. from the University of Illinois at Urbana-Champaign in 2012. He is the recipient of the NSF CAREER award (2016), ONR YIP award (2018), M.E. Van Valkenburg award, Google Faculty Research Award (2016, 2017), Cisco Faculty research award (2016), Intel Faculty research award (2016, 2017), and IBM Faculty award (2017). He was awarded Princeton University's E. Lawrence Keyes Award for outstanding research and teaching, and is the recipient of multiple outstanding paper awards including ACM CCS and ACM ASIACCS.

H. Vincent Poor (S'72, M'77, SM'82, F'87) received the Ph.D. degree in EECS from Princeton University in 1977. From 1977 until 1990, he was on the faculty of the University of Illinois at Urbana-Champaign. Since 1990 he has been on the faculty at Princeton, where he is currently the Michael Henry Strater University Professor of Electrical Engineering. During 2006 to 2016, he served as Dean of Princeton's School of Engineering and Applied Science. He has also held visiting appointments at several other universities, including most recently at Berkeley and Cambridge. His research interests are in the areas of information theory and signal processing, and their applications in wireless networks, energy systems and related fields. Among his publications in these areas is the forthcoming book Advanced Data Analytics for Power Systems (Cambridge University Press).

Dr. Poor is a member of the National Academy of Engineering and the National Academy of Sciences, and is a foreign member of the Chinese Academy of Sciences, the Royal Society, and other national and international academies. Recent recognition of his work includes the 2019 ASEE Benjamin Garver Lamme Award, and a D.Eng. honoris causa from the University of Waterloo, also awarded in 2019.