

Shedding Light on Inconsistencies in Grid Cybersecurity: Disconnects and Recommendations

Brian Singer

Carnegie Mellon University
briansin@andrew.cmu.edu

Amritanshu Pandey

Carnegie Mellon University
amritanp@andrew.cmu.edu

Shimiao Li

Carnegie Mellon University
shimiaol@andrew.cmu.edu

Lujo Bauer

Carnegie Mellon University
lbauer@andrew.cmu.edu

Craig Miller

Carnegie Mellon University
craigmil@andrew.cmu.edu

Lawrence Pileggi

Carnegie Mellon University
pileggi@andrew.cmu.edu

Vyas Sekar

Carnegie Mellon University
vsekar@andrew.cmu.edu

Abstract—The operational, academic, and policy communities disagree on which threats against the power grid are likely and what damage would ensue. For instance, the feasibility and impact of MadIoT-style attacks is being actively debated. By surveying grid experts (N=18) we find that disagreements are not unique to MadIoT attacks but occur across multiple well-studied grid threats. Based on prior work and our survey, we hypothesize that the disagreements stem from inconsistencies in how grid threats are modeled. We identify five likely causes of modeling inconsistencies: 1) using unrealistic grid topologies, 2) assuming unrealistic capabilities for attackers, 3) exploring too few grid scenarios, 4) using incomplete simulators that omit relevant grid processes, and 5) using simulators that incorrectly model key grid processes. To check these hypotheses, we create a modeling framework and examine how these factors change our understanding of the feasibility and impact of grid threats. We use four diverse grid threats as case studies: MadIoT, False Data Injection Attacks, Substation Circuit Breaker Takeover, and Power Plant Takeover. We find that each of our hypothesized causes of modeling inconsistencies has a significant effect on modeling the outcomes of attacks. For example, we find that MadIoT attacks are much less feasible and require significantly more high-wattage IoT devices on realistic topologies than on topologies previously used to model them. In contrast, we find that Substation Circuit Breaker Takeover attacks are much more feasible in emergency scenarios and may require significantly fewer substations for failure than previous modeling suggested. We conclude with actionable recommendations for accurately assessing the impact of threats against the grid.

Index Terms—power grid, cybersecurity, grid security, computer security, industrial control systems

1. Introduction

Over the last decade, the frequency and impact of cyber attacks on the electric grid infrastructure have increased. For instance, the Ukrainian power grid had major outages in 2015 and 2016 due to cyber attacks [1], resulting in electric

supply disruption to hundreds of thousands of people. In addition, the U.S. electric grid has had several instances of cyber-intrusions as well [1]. A 2015 report [2] estimated that a large-scale cyber-attack on the U.S. grid “could leave 15 states and 93 million people [...] without power.” The report projected the economic impact of such an outage to be between 243 billion and 1 trillion USD, resulting from “direct damage to assets and infrastructure, decline in sales revenue to electricity supply companies, loss of sales revenue to business and disruption to the supply chain.”

Despite the importance of grid security, there are significant disconnects across operational, academic security, academic grid, and policy communities on which threats against the grid are dangerous. For example, in 2018, Soltan et al., proposed a novel threat against the power grid, MadIoT [3], in which an attacker compromises a set of high-wattage IoT devices (e.g., smart thermostats, refrigerators, solar panels) and synchronously turns on or off the compromised devices to drastically change the grid’s load demand and cause a blackout. Unfortunately, it is not clear how much change in demand (from high-wattage IoT devices) an attack would require. While the original study estimated that a 2% demand increase could induce a blackout [3], another study estimated that the attack would require at least a 10% demand increase [4], while a third suggested that in some scenarios the attack is not possible [5].

While it is tempting to dismiss this inconsistency in assessing risk as unique to this threat and specific studies, the reality could be far worse. We surveyed grid security experts and asked them about the likelihood and impact (see Sec. 3) of four threats proposed in the literature [1], [3], [6], [7]; participants disagreed on the likelihood and impact of all four threats.

Our goal in this paper is to provide actionable recommendations for avoiding such confusion going forward. We do so by revisiting, from first principles, how prior work has modeled various grid threats. Specifically, any attempt to model the effectiveness of a grid threat consists of three components: a topology (a description of the grid connectivity and component properties; see Sec. 2), an

attack, and a simulator. Inconsistencies or errors in these components can lead to different outcomes in likelihood and impact estimates. For example, even a simulator with perfect precision can yield unrealistic attack impact estimates if used with unrealistic grid topologies.

We identify five (non-exhaustive) likely causes of modeling inconsistencies in prior literature:

- 1) attacks modeled on unrealistic grid topologies (e.g., that did not meet standard resiliency requirements [8]);
- 2) attacks assumed unrealistic attacker capabilities (e.g., attackers required control of all sensors [6]);
- 3) attacks studied on too few grid scenarios (e.g., attacks not evaluated on overall high demand scenarios);
- 4) simulators did not model relevant grid processes (e.g., did not simulate reserve generation [9]); and
- 5) simulators modeled grid processes incorrectly (e.g., simulated droop reserves incorrectly [3]).

To quantitatively analyze these factors and validate the potential sources of inconsistencies, we implement a framework to evaluate how various design and dataset choices influence findings on threat likelihood and impact. The framework contains 1) a topology-verification module, 2) a steady-state power-flow module to model attacks that target physical grid elements, and 3) a state-estimation module that models attacks that target grid operators.

Using the framework we find that:

- 1) Simulations on realistic topologies show that successful MadIoT attacks require a significantly greater change in demand from high-wattage IoT devices than suggested by prior work that analyzed less realistic topologies [3].
- 2) Prior work assumes that during a False Data Injection Attack attackers have control over all grid measurement devices. If attackers control only a subset of measurement devices, which we argue is overwhelmingly likely, attacks become detectable and less impactful.
- 3) Substation Circuit Breaker Takeover attacks on grids in emergency scenarios need to compromise significantly fewer substations than previously suggested to cause grid failure.
- 4) When droop reserves are correctly modeled, attackers can induce grid failure by compromising fewer high-power generators than previously suggested.

Based on these findings we recommend (1) creating standardized realistic public topologies for evaluation; (2) standardizing simulation tools, techniques, and procedures; (3) using evaluation methodology that considers more comprehensive scenarios, including emergency scenarios; and (4) developing assumed-failure cybersecurity reliability standards that consider the grid as a single entity.

Reproducibility: We released all of the topologies used in this study, in standard formats so that any steady-state power flow solver can simulate them [10]. Additionally, we released the extensions [10] we have built on top of a state-of-art power-flow solver [11].

Contributions and roadmap: In summary, this paper makes the following contributions:

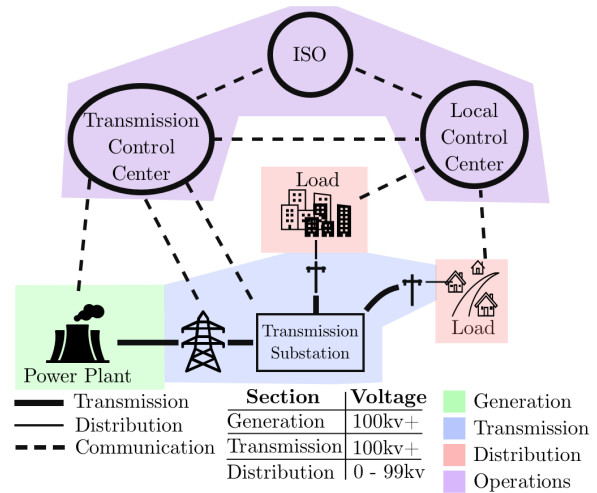


Figure 1. The grid consists of three distinct sections: generation, transmission, and distribution. The generation and transmission sections operate at $>100\text{kV}$ while the distribution section operates at $0-99\text{kV}$ [13]. Operators manage these three sections. Transmission control centers monitor and manage generation power plants and transmission substations. Local control centers monitor and manage substations, distribution loads, and distribution generation. Independent system operators (ISOs) are non-profit organizations that manage the generation and transmission sections of the grid

- 1) Identifying key factors that lead to modeling inconsistencies in prior work.
- 2) Evaluating the variation in threat outcomes due to modeling inconsistencies on four grid threats: MadIoT, False Data Injection Attack, Substation Circuit Breaker Takeover, and Power Plant Takeover.
- 3) Recommendations for the grid security community to avoid inconsistencies in the future.

The next section provides background on the power grid and grid threats. Sec. 3 uses anecdotal evidence from prior work and an expert survey to motivate a first-principles approach to identifying grid modeling inconsistencies. Sec. 4 identifies five key (but non-exhaustive) factors that likely contribute to modeling disparities. Sec. 5 describes how we validate that the identified factors change the understanding of the feasibility and impact of grid threats. These findings inform our recommendations, in Sec. 7. Sec. 8 reviews other relevant related work, and we conclude in Sec. 9.

2. Background: Grid and Grid Threats

In this section, we first provide a background on the components of the power grid and how operators manage the grid. A detailed tutorial is outside the scope of this paper, and we refer readers to reference guides on power systems [12]. Our focus here is on explaining the relevant components and management strategies as they pertain to grid threats. We also provide a taxonomy of grid threats to scope our work.

2.1. Overview of grid

Grid devices are categorized into three sections (Fig. 1).

- 1) The *generation* section includes large-scale power plants that generate hundreds of megawatts (MW).
- 2) The *transmission* section consists of high-capacity power lines and transformers carrying high-voltage ($> 100kV$) [13] generator power over long distances. For example, they transfer electricity from Cleveland to Boston.
- 3) The *distribution* section has low voltage equipment ($< 99kV$ [13]), which includes consumer and industrial loads, low voltage transmission lines, and distributed energy resources. For example, Boston distributes electricity to homes through the distribution section.

The generation and transmission sections together comprise the *bulk energy system*, where equipment operates at high voltage ($> 100kV$) [13]. Power from the transmission section transfers to the distribution section through substations, which step-up or step-down voltages through transformers. These physical sections are managed by the operations section, as shown in Fig. 1.

2.2. Grid operation

The operations section monitors electrical properties by creating a *power grid topology* (GT), which describes the energized section of the grid at an instance in time. A grid topology $GT = G(V, E)$ comprises a collection of vertices V and a collection of edges E . In the grid community, a vertex is called a *bus*. A bus is a connection point that aggregates power from devices such as generators and loads. A *generator* is a device that produces electricity and a *load* is a device that consumes electricity. The edges in a grid topology are the transmission lines and transformers that carry power between the buses in the grid. The edges and vertices in a grid topology have attributes that describe device properties, such as the maximum generation power of a generator or the power capacity of a transmission line.

Operators use the topology to simulate the grid to estimate a *grid state*. A grid state describes all of the electrical properties of the grid at a single instance of time, such as power flow along transmission lines and voltages of buses. Operators will take corrective action if they notice electrical properties of concern. For instance, if a disturbance such as a grid attack occurs, grid electrical states will begin to deviate from their nominal values, requiring corrective actions.

Threats against the grid typically involve causing an imbalance between supply (generation) and demand (load) [3], [9], [14]. We next give a brief overview of the techniques operators use to compensate for such imbalances. For descriptions of other control mechanisms, we refer the reader to reference textbooks [12], [15], [16].

Droop control (*primary control*) serves as the first safeguard against a change in power demand. As power demand increases, the grid's electrical frequency decreases. Governor systems in generators monitor grid frequency,

and if the frequency starts rising or falling, a generator's governor control changes its power output to stabilize the frequency [17]. Droop control occurs within seconds of the change in power demand. Almost all generators in the bulk energy system participate in droop control. Furthermore, regulations require droop reserves to be sufficient compensation for the failure of any one grid component [18].

Droop can stabilize the grid frequency, but the stabilized frequency will be different than the nominal frequency (50Hz or 60Hz) [19]. To correct the frequency to a nominal value, operators use *secondary control* to alter the power output for a select group of generators. For example, if a US grid stabilizes to a frequency of 59.90 Hz, the secondary control signal (also known as *automatic generation control* (AGC)) will order several generators to increase their power output to get the frequency closer to the nominal 60Hz. In the US, secondary control signals are sent once every minute or two [18].

If both primary and secondary control fail to compensate for insufficient generation and restore the grid state to nominal values, operators will resort to using *responsive reserves* or *tertiary control* to increase grid generation. Responsive reserves are a type of generation that can turn on within minutes of a disturbance [18]. For example, the Texas grid regulations require that at least a total of 3000MW of generation is available within three minutes of a disturbance [20]. Operators use responsive reserves in conjunction with primary and secondary control to stabilize the grid. Tertiary control refers to operators purchasing additional energy through a process called economic dispatch. Tertiary control takes at least 10 minutes because of generation-ramping constraints and the slow speed of economic dispatch algorithms [18].

If operators use all the aforementioned tools and generation still cannot meet demand, operators will use *under-frequency load shedding* to maintain grid stability. In under-frequency load shedding, devices disconnect from the grid because their protective relays trip when they sense the frequency is below a set threshold. Operators set a relay's frequency threshold based on the importance of the device: thresholds for non-critical devices are set to higher than for critical devices. As a result, if a disturbance occurs, non-critical devices will disconnect and stabilize the grid before critical devices shut off. Grid operators use under frequency-load shedding to shut off power to non-critical sections of the grid as a last-ditch effort to preserve power delivery to critical sections (e.g., hospitals, transportation) and prevent uncontrolled grid failure (like the 2003 Northeast Blackout [24]). For instance, in the winter of 2020 when the demand spiked, the Texas grid had to use load shedding (e.g., rolling blackouts) to prevent a total blackout.

During extreme events, operators may be unable to prevent the grid from failing and a *blackout* will occur. A blackout is when a significant portion of the grid loses electricity (e.g., a hurricane causing a city to lose power) [25]. Blackouts take at least a few hours to recover from but can last much longer [26]. Temporary reductions in electrical quality (e.g., lights flickering) are brownouts [25].

TABLE 1. TAXONOMY OF POWER GRID THREATS, WITH EXAMPLES.

Threat	Description	Target	Goals	Resource	Capabilities
			Damage		Action
MadIoT [3]	Attackers compromise high-wattage IoT devices to destabilize the grid by synchronously powering devices.	Distribution	Power delivery Economic	High-wattage IoT devices	Toggle power
False Data Injection Attack [6]	Attackers compromise grid measurement devices and modify the data sent back to operators.	Generation Transmission Distribution	Power delivery Economic	Measurement devices Topology/state information	Send false measurements Craft data to avoid bad data detection
Substation Circuit Breaker Takeover [1]	Attackers open substation circuit breakers to disconnect devices and prevent power delivery.	Transmission Distribution	Power delivery	Substations' circuit breakers Grid topology information	Open circuit breaker to disconnect devices Strategically compromise substations
Power Plant Takeover [21]	Attackers compromise power plant control systems to turn off power delivery or safety protocols.	Generation	Power delivery Human endangerment	Power plant control system Grid topology information	Shut off power Disable safety equipment Strategically target power plants
Measurement Jitter [22]	Attackers compromise grid measurement devices and delay the data sent back to operators.	Generation Transmission Distribution	Power delivery Economic	Measurement devices	Delay measurement data
Ransomware [23]	Attackers encrypt computer systems data to prevent power delivery or hold the data ransom.	Generation Transmission Distribution	Power delivery Economic	IT computer systems	Disk encryption

2.3. Grid threats

Before we describe the specific threats we study, we define key terms. We define a *Threat* as a tuple $\langle \text{Goals}, \text{Capabilities} \rangle$; the *threat* will attempt to achieve its *goals* by using its *capabilities*. For example, for a MadIoT threat, the goal is to cause a power-grid blackout and the capability is to synchronously turn on or off high-wattage IoT devices. We define an *Attack* as a specific instance of a threat and a tuple of: $\langle \text{Goals}, \text{Instance of Capabilities}, \text{Strategy} \rangle$. An attack's goals are a subset of the threat's goals; an attack's capabilities are an instantiation of the threat's capabilities. For example, for a MadIoT attack, the capabilities are the set of high-wattage IoT devices that the attacker can control. An attack's strategy is how the attacker utilizes their capabilities to achieve their goals. For example, one MadIoT attack strategy is to briefly power on all compromised IoT devices and subsequently power off all those devices.

To describe and categorize grid threats, we created a non-exhaustive taxonomy based on a threat's goals and capabilities. We define a *Goal* as a tuple $\langle \text{Target}, \text{Damage} \rangle$, where the *Target* is the section of the grid targeted and the *Damage* is the type of impact the *threat* causes. A *Capability* is a tuple $\langle \text{Resource}, \text{Actions} \rangle$. A capability maps a *Resource* to the set of *Actions* the attacker can perform if they compromise the resource. For example, if an attacker compromises a generator (*Resource*), the attacker gains access to the following actions: shutting off the

generator, raising the generator's power output set-point, or even destabilizing the generator causing it to explode [27]. Table 1 gives an overview of well-studied grid threats.

Scope: Our paper considers only threats that directly impact power delivery. Other threats, such as market manipulation [28] and wearing down devices [27], are out of scope. Within threats that directly stop power delivery, we examine those that 1) target different sections of the grid, 2) are a mix of theoretical and real-world threats, and 3) are relevant to the grid, security, and operational communities. As a result, we focus on MadIoT, False Data Injection Attacks (FDIA), Substation Circuit Breaker Takeover (SCBT), and Power Plant Takeover (PPT) threats. MadIoT targets the distribution section, FDIA targets the operations section, SCBT targets the transmission section, and PPT targets the generation section. To the best of our knowledge, MadIoT and FDIA attacks have not been observed in practice. In contrast, SCBT and PPT have been documented and observed in practice [1], [21]. Finally, papers discussing each of these threats are highly cited [3], [6], [9], [14]. We discuss these four threats in more detail below.

MadIoT: An attacker compromises a set of high-wattage IoT devices (e.g., smart thermostats, refrigerators, solar panels). Next, the attacker synchronously toggles power for all of the compromised IoT devices to cause a drastic change in the grid's load demand. This change in load can cause a blackout.

False Data Injection Attack (FDIA): An attacker com-

promises grid measurement devices and uses them to send fake data to grid operators. To remain stealthy, the attacker uses their knowledge of grid topology structure (i.e., nodes and edge connections), topology attributes (e.g., edge impedance, resistance), and grid state (i.e., voltages of all nodes) to generate adversarial data. Operators process raw measurement data through *AC state estimation*, which outputs a grid state and a *J-value*, a measure of how likely it is that some measurement data is incorrect. An FDIA attack succeeds if the adversarial data does not raise the J-value enough to cause an alarm, thereby deceiving operators into believing the grid is in a different state than reality.¹

Substation Circuit Breaker Takeover (SCBT): An attacker compromises substations and disconnects substations' loads, generators, and transmission lines from the grid. Substations are critical to the stability of the grid; a 2014 estimate stated that if any nine of the US's 30 critical substations are taken offline, all US grids will experience blackouts [29]. In 2015 and 2016, hackers used an SCBT attack to take down portions of the Ukrainian grid [1].

Power Plant Takeover (PPT): An attacker compromises multiple high-power generators. The attacker can then turn off power generation, increase power output, or even disable safety equipment to endanger human lives [21]. A sudden loss or increase in power output can cause a supply-demand imbalance and other grid stability issues [7], [14].

3. Motivation

Although there are many known grid threats, there is lack of consensus on how important or impactful a particular threat can be. We begin by illustrating the inconsistencies from the debate on the MadIoT threat. We then discuss the results of an expert survey that suggest that this confusion is not restricted to MadIoT and also affects the other grid threats we study.

3.1. Debate around MadIoT

The literature is inconsistent in estimating potential the likelihood and potential impact of the MadIoT threat. When Soltan et al., first postulated the attack, they showed that using compromised IoT devices to increase demand by as little as 2% was enough to collapse the (synthetic) Polish grid [3]. Follow-up work, however, came to different conclusions. For example, Huang et al., simulated MadIoT attacks on a confidential North America Regional grid topology and found that a 1% demand increase attack had no negative impact, but a 10% demand increase caused *minor* amounts of load shedding [4]. Ospine et al., modeled the impact of a MadIoT attacks on four New York grid scenarios and found in one scenario that "there are no time periods an attacker could effectively compromise the frequency [to cause a grid failure]" [5]. Taken together, it is unclear to what extent an attacker would need to modify demand (and therefore the

1. Both FDIA and MadIoT have secondary goals of manipulating power grid markets, but this is out of scope for our analyses.

number of compromised high-wattage IoT devices) to cause grid failure or if the attack is even feasible.

3.2. Expert survey

The confusion surrounding the likelihood and impact of MadIoT threats motivated us to ask a broader question: Is this confusion specific to MadIoT? Or is this lack of consistency restricted to a specific set of academic papers? To this end, we conducted an expert survey (N=18) to understand if threat evaluation inconsistencies are common across other threats as well. The survey is qualitative and illustrative; it elicits experts' *perception* of threat likelihood and impact and does not empirically evaluate the likelihood and impact. Specifically, we asked experts for their perception of four grid threats: MadIoT, FDIA, SCBT, and PPT. We found that experts were not confident in their answers and diverged on their content. Next, we discuss the methodology and insights from the survey results. The survey questions are in App. A and additional results in App. B.

Expert recruitment: The survey was distributed from May to August 2021. Participants were grid security experts from academia and industry. To trial the survey, we recruited a participants by reaching out to colleagues working in grid security from academia and industry (N=2). After the trial, we distributed the survey on PowerGlobe [30] (N=10), a professional power-systems forum, and CRED-C [31] (N=6), an academic and industry grid-resilience consortium. We chose PowerGlobe because of its high engagement with power-systems experts and CRED-C because its members have professional experience with power-systems cybersecurity. Of the 18 participants, seven specialized in both computer security and power systems, five just in power systems, two in computer security, and four preferred not to state their specialization.

Survey design: Experts were told to evaluate threats based on the most dangerous attacks that instantiate each threat. We asked experts the likelihood of a grid threat causing a *significant disruption on the grid*, where a significant disruption was defined as either 1) 300 MW of load dropped in a dispatch cycle, 2) economic loss greater than ten million dollars, or 3) endangering any number of human lives (definition based on NERC's event reporting requirements [32]).

Threat impact was elicited by asking about power outage size and duration: for power outage size, participants estimated the largest amount of load dropped from the grid because of the attack; for power outage duration, experts estimated the amount of time required to restore power if the threat caused a significant disruption.

Experts were asked about likelihood and impact of attacks for each of three grid states, based on the North American Electric Reliability Corporation standard [33]:

- 1) *Normal:* Physical Responsive Capability (PRC) is at normal capacity and grid frequency is stable at 60Hz.
- 2) *Energy Emergency Alert Level 2 (EEA2):* Operators begin utilizing load management procedures such as reserve generation.




Figure 2. Expert opinions on the likelihood of grid threats causing a significant disruption. Experts did not agree about the likelihood of any grid threat.

3) *Energy Emergency Alert Level 3 (EEA3):* Operators begin to use load interruption techniques (e.g., load shedding) to stabilize the grid.

For example, experts were asked “If the power grid is in EEA2, what would be the likelihood of a MadIoT attack causing a significant disruption?” For each question, participants were asked about their confidence in their response.²

In addition, throughout the survey we provided options for free-form responses for participants to explain the reasoning behind their answers and to provide comments in the context of specific questions. While these free-form responses were optional, they provided useful hints on the underlying reasons on why participants responded to the questions the way they did.

Findings: Next, we report on key findings from our survey.

Finding 1: For all grid states and threats, experts gave conflicting answers for likelihood (Fig. 2), outage size, and outage duration (Fig. 3).

At a high level, experts had conflicting perceptions about both the likelihood and the impact of attacks. As shown in Fig. 2, experts disagreed on the likelihood of attacks; in several scenarios, predictions had ranges greater than 95%. Additionally, Fig. 3 illustrates that experts disagreed about threat impact. When the grid is in normal state, the greatest consensus for outage size was for the Substation Circuit Breaker Takeover threat, for which 36% of experts believed it would cause an outage of 1–9MW and 36% believed 100–999MW, which are drastically different outage sizes. Consensus for outage size was higher if the grid is in EEA3. For example, 53% of experts agreed that a power plant takeover threat would cause more than 1000MW of load dropped if the grid was in EEA3. In addition, for all grid states experts disagreed about recovery time after an attack. The largest consensus for threat recovery time was for MadIoT when the grid is in EEA2 with 50% of experts believing it would take 1–23 hours for the grid to recover.

Finding 2: For all grid states and threats, experts report low confidence in their answers for likelihood, outage size, and outage duration (Fig. 4).

2. To limit survey time, we did not ask the likelihood and impact questions for EEA1, when all non-reserve generation is at capacity.

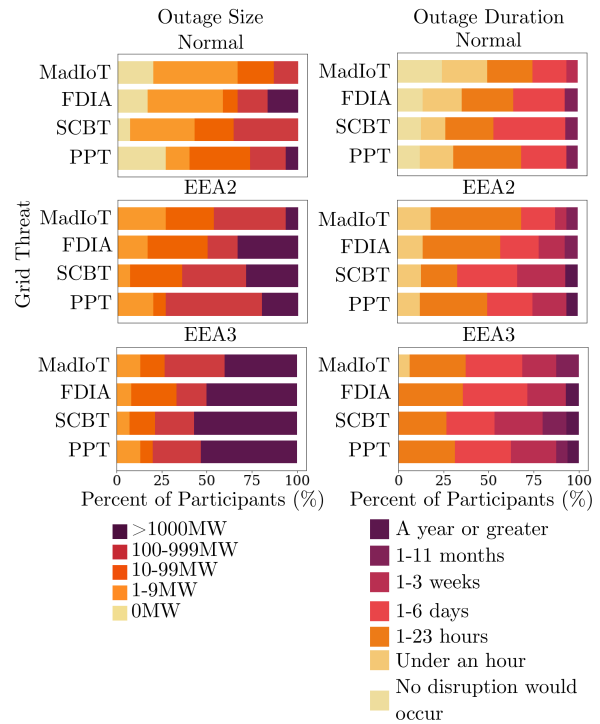


Figure 3. Expert opinions on the outage size and duration of grid threats. Experts did not agree about the outage size and duration for any grid threat.

Fig. 4 shows that experts were not confident in their answers. We considered an expert to have high confidence if they reported they were *very confident* or *confident* in their answers. For all threat questions, the average confidence was only 24% reporting high confidence in their answers. Experts were most confident about the outage duration of a PPT attack, for which 45% reported high confidence. They were least confident about their predictions for the likelihood of a MadIoT attack causing a significant disruption, with 85% reporting low confidence.

Free-form inputs from experts: We manually analyzed the free-form responses to see if there was some emerging pattern. Interestingly, we found that many experts indicated that inconsistencies may be due to shortcomings in modeling. For example, one expert claimed FDIA attacks were infeasible because of a modeling error “Other sorts of attack are inherently much more likely and plausible and that the conception and modeling for this attack [FDIA] is likely in error”. In addition, experts had conflicting views on the state of the art in MadIoT modeling. One expert cited work saying the MadIoT threat was not possible “Granger Morgan’s work has shown that this sort of attack [MadIoT] is very unlikely to cause a significant disruption,” but another expert thought the field was poorly studied, “Black start [starting a collapsed grid from a MadIoT attack] with extensive DER [Distributed Energy Reserves] is poorly studied.”³

3. To protect anonymity, we do not share the full text of free-form responses.

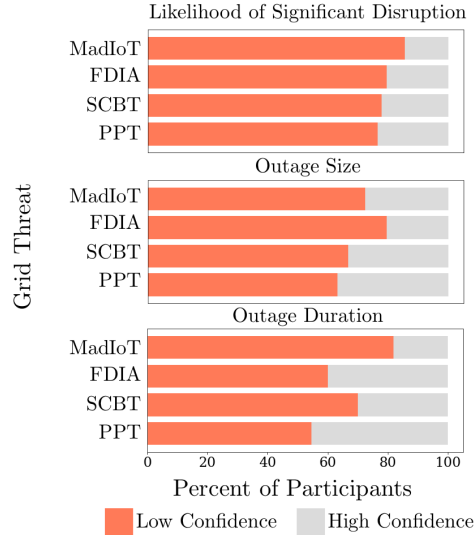


Figure 4. Expert confidence in their predictions for likelihood, outage size, and outage duration. Low confidence is defined as an expert reporting not at all confident, not confident, or somewhat confident. High confidence is defined as an expert reporting confident or very confident. The majority of experts reported low confidence in their predictions for all threats.

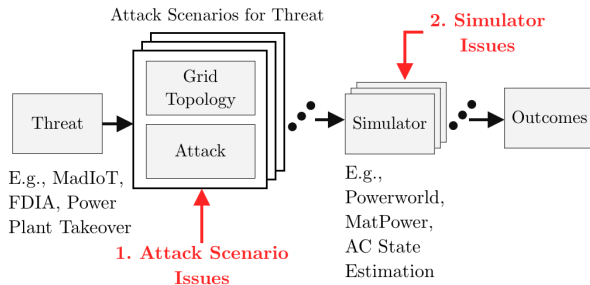


Figure 5. An abstract view of grid threat modeling. Previous literature has had challenges with 1) creating a sufficient set of attack scenarios and 2) simulating the attack scenario with sufficient accuracy.

4. Factors Likely Causing Inconsistencies

Based on the MadIoT debate and the survey responses, we hypothesize that many of the disagreements in literature stem from inconsistencies in modeling grid attacks. To qualitatively understand these inconsistencies, we create an abstract view of existing efforts in grid threat simulation. This abstract view then helps identify candidate factors that cause modeling and outcome inconsistencies.

4.1. Abstract view of grid threat simulation

Fig. 5 shows an abstract view of the grid modeling procedure. First, users generate a set of *attack scenarios* to evaluate a grid threat’s impact. An *attack scenario* is an attack occurring on a grid topology, defined by a tuple $(Attack, Topology)$. A threat is typically analyzed through a set of attack scenarios. For example, to study a MadIoT

threat on the Texas grid, a set of attack scenarios describing various Texas grid topologies and MadIoT attacks will be analyzed. Next, each of the attack scenarios is modeled and evaluated with a simulator (e.g., MATPOWER [34], PowerWorld [35], SUGAR [11]). The simulator evaluates the impact of the threat by modeling grid operation processes and estimating the grid state for each attack scenario. A *power-grid process* is any control or procedure used to operate the grid reliably and securely. Examples include droop control and economic dispatch (see Sec. 2).

4.2. Likely sources of modeling inconsistencies

Given this abstraction, we revisit the literature on these various threats [3], [6], [7], [9], and manually identify factors in this modeling process that potentially lead to inconsistent estimates of likelihood and impact.

We cluster these potential factors into two primary groups: (1) issues with the set of evaluated attack scenarios (labeled as **Scen**); and (2) issues with the choice of simulator and the processes it captures (labeled as **Sim**). Within each group, we identify several sub-factors, which we highlight next and then quantitatively analyze in Sec. 5.

The first set of factors relates to the simulator input in Fig. 5.

Scen 1. Attacks evaluated on unrealistic grid topologies.

Grid attacks have been studied on a variety of topologies, such as small bus systems (<300 buses)⁴ [3], [5], [6], [37], synthetic topologies [3], [9],⁵ and confidential real-world topologies [4]. It is often unclear how realistic these topologies are or whether any lack of realism has a significant impact on the evaluation of a threat. For example, in many cases of prior work [3], [6], [9], [38], it is unclear if the grid topologies satisfy N-1, which is a regulatory requirement for operators [8].

Scen 2. Attacks assume unrealistic attacker capabilities.

In addition, previous work has studied threats using attacks that assume unrealistic attacker capabilities. For example, previous work evaluating FDIA assumes the attacker has control over all or a majority of grid measurement devices [6]. In reality, simultaneously compromising all grid measurement devices is a monumental task for attackers.

Scen 3. Low coverage of grid topology scenarios.

Finally, previous work has evaluated the impact of threats by simulating attacks on few grid scenarios [4], [37]. By only sampling a small number of grid scenarios, the analyses may miss critical or extreme grid situations. For instance, to our knowledge, there is little understanding of the impact of SCBT attacks during various extreme weather conditions (e.g., hot or winter weather).

The second group of factors relates to the simulation engine. We identify two candidate sub-factors:

4. For comparison, the eastern US grid has over 85K buses [36]

5. “Synthetic” refers to topologies that have similar power demand profiles as their realistic-grid counterparts but do not divulge proprietary grid information.

Sim 1. Simulators do not model all relevant grid processes.

The FDIA threat has repeatedly been evaluated by simulating state-estimation algorithms not used on real grids [6], [39]: the original analysis [6] uses DC state-estimation, while operator control rooms use AC-state estimation [37]). In another case, the MadIoT threat has been evaluated without simulating responsive reserves [3]–[5], a process that would likely occur during a real MadIoT attack.

Sim 2. Simulators model grid processes incorrectly.

Generators take time to change power output during droop control. The impact of a MadIoT threat has been evaluated using droop control [3], [4]; however, the rate at which generators can physically ramp up or down was sometimes not considered [3]. Most generators can only change their power output by 5–10% of operating power within ten minutes [18].

Our goal in this paper is to be illustrative rather than comprehensive. While we identify some factors that can lead to inconsistencies and illustrate their impact (Sec. 5.2), we do not claim that this set of factors is exhaustive.

5. Quantitative Assessment of Factors

After enumerating candidate factors, we next quantitatively validate that these factors indeed cause inconsistencies in modeling outcomes. Since the errors are intrinsic to the custom simulation approaches in prior work, we revisit the simulation process and develop a framework to avoid reintroducing the same biases. We describe the framework (Sec. 5.1) and then discuss our key findings (Sec. 5.2).

5.1. Approach

Our framework has three modules: 1) a topology-verification module, 2) a steady-state power-flow module to model attacks that target the physical sections of the grid (i.e., generation, transmission, and distribution), and 3) a state-estimation module to simulate attacks that target the operations section. We use the topology-verification module to evaluate whether grid topologies are realistic. We use the steady-state module to study the impact of unrealistic topologies on the attack outcome and to evaluate whether emergency grid scenarios are more susceptible to attacks. We also use the steady-state module to evaluate how simulating an incorrect set of processes or simulating processes incorrectly affects attack outcomes. Finally, we use the state-estimation module to evaluate FDIA attacks with varying levels of unrealistic capabilities. Next, we discuss the design of each module in more detail.

Topology verification: To ensure that the topologies used for threat assessment are realistic, we implement two checks:

- 1) *N-1 verification* iterates over the components in the power grid topology and runs a power-flow simulation with the component in a failure state. One device failure can alter the electrical properties of the grid and cause other devices to (attempt to) operate outside of their capabilities. For example, if a high-capacity transmission

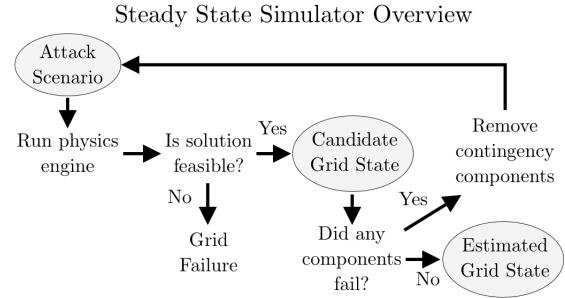


Figure 6. An overview of the framework’s steady-state simulator. The simulator models the impact of attacks that target the physical sections of the grid (i.e., Generation, Transmission, and Distribution).

line fails, electricity may be rerouted to low-capacity transmission lines and may cause them to exceed their capacity (causing them to also fail). If any simulation does not converge or results in device failures, we consider the power-grid topology unrealistic, because real grids are regulated to satisfy N-1 [8].

- 2) *Topology verification* checks that device parameters are realistic, e.g., that all line flows (for nominal and all N-1 contingent states) are below each line’s capacity.

Steady-state power flow: Power-flow solvers simulate the physical sections of the grid by solving the equations that describe a grid scenario and estimating the grid state. SUGAR [11], MATPOWER [34], and PowerWorld [35] are examples of power-flow solvers. They use different underlying algorithms, but should provide the same results if they converge. We extend SUGAR to analyze threats that affect the physical sections of the grid, such as MadIoT, SCBT, and PPT. Our choice of SUGAR is driven both by the solver’s robustness of convergence (i.e., ability to solve topologies that other power-flow tools cannot) and its ability to distinguish blackouts from failures due to initial conditions being far from the solution (topologies with initial conditions far from the solution are harder to solve and may cause some power-flow solvers to incorrectly report a blackout) [40]. Our extensions output RAW files [41], a standard grid-simulation file format. RAW files can be solved using any power-flow solver, but we leave integrating the extensions with other solvers for future work.⁶

First, we extend SUGAR to simulate droop, AGC, and responsive reserve generation to create a *physics engine* with realistic ten-minute power reserves. We simulate droop control by having generators increase or decrease their power generation proportionally to operational power. Most generators can only change their power output by 5–10% of operating power within ten minutes [18]. To be conservative, we constrain the change to be within $\pm 5\%$ of the generator’s power output. In regards to AGC, operators pay a select group of large generators to have additional reserves [18]. We implement AGC by increasing power reserves of generators with operating power greater than

6. Integrating with other solvers requires additional engineering effort due to different programming languages and solver configurations.

100MW and bounding the change to $\pm 10\%$ of the operating power. The engine adds AGC until the topology has realistic AGC reserves [42]. We simulate responsive reserve generation by adding generators in diverse locations with a total power output matching real-world estimates [42]. The total ten-minute power reserves of US grids are 6–15% of total operating demand [42]. In our physics engine, of the overall reserves (6–15% of total operating demand), droop contributes 5% of total operating demand.

Next, we use our physics engine to implement a cascading analysis simulator following the procedure shown in Fig. 6: The simulator runs the attack scenario on a physics engine. The simulation ends if the physics engine does not find a feasible solution. If the physics engine finds a feasible solution (i.e., a candidate grid state), but components violate their thermal or voltage limits, the simulator removes those components and simulates the attack again with the modified topology. The process repeats until a solution is found or the physics engine fails to find a feasible solution. A *grid failure* occurs if the physics engine fails to find a feasible solution; this is consistent with operational and prior security literature (e.g., [3], [43]).

State-estimation module: State-estimation modeling is for analyzing threats that target the operations section of the grid, such as FDIA. We extend an open-source AC state-estimation simulator [34] to accurately evaluate the outcome of FDIA attacks. AC state-estimation outputs an estimated grid state and a J-value by solving a weighted least-squares optimization problem. The J-value is the residual sum of squares and signifies how likely it is that some measurement data is incorrect. The greater the J-value, the more likely it is that operators will identify the attack through *bad data detection* [44], [45], a process in which operators run hypothesis tests on J-values. Several bad-data-detection algorithms have been proposed [44], [45]; we simulate the original bad-data detection [44] because control rooms still utilize this algorithm [6].

We extend the simulator to model FDIA attacks by generating adversarial measurements based on the attacker’s goal. AC state-estimation interprets the measurements and outputs J-values and an erroneous grid state. FDIA attacks that cause large errors in grid-state estimates but insignificant changes in J-values are successful because operators may think the grid is in a different state than it actually is.

Our simulator is more accurate and realistic than prior work along several dimensions. We use state-of-the-art, realistic topologies [36] (unlike [3], [5]). The simulator has realistic FDIA attacker capability assumptions (unlike [6], [37], [46]). Threats are evaluated on a variety of emergency scenarios (unlike [3], [4], [6]). Additionally, we simulate AGC, Responsive Reserve Generation, and realistic power limits to accurately evaluate MadIoT attacks (unlike [3]). We increase the accuracy of our analysis by using robust solvers and realistic grid scenarios (based on grid regulations [17], [47]). Additionally, we verify our key simulator results with PowerWorld [35] (a standard, commercial modeling tool) to confirm the correctness of our simulator implementation.

5.2. Findings

Next, we use this framework to validate that the factors illustrated in Sec. 4 change our understanding of the feasibility and impact of grid threats. For each identified factor, we choose one of the four grid threats from our survey as an example on which to illustrate the impact of this factor.

Grid topologies: For attacks targeting the physical sections of the grid, we use the following public topologies: the 3120-bus Synthetic Polish Summer Peak 2008 [38], the 2k-bus Synthetic Texas [36], and the 10k-bus Synthetic Western Interconnection [36]. Due to scalability limitations of the state-estimation tool, we only analyze FDIA attacks on the 500-bus Synthetic South Carolina [36] topology. We use these topologies because they are commonly used in prior work [3], [11]. To our knowledge, these US topologies are the most realistic publicly available US grid topologies.

Finding 3 (Scen 1): Many publicly available grid topologies are unrealistic on two fronts: 1) they do not satisfy N-1 criteria (mainly due to line and transformer flow violations) and 2) they have insufficient power reserves.

We analyze topologies used in prior literature to determine if they meet minimum operational standards. Specifically, we evaluate whether they meet the real-world N-1 property for resilience [8] and whether they have sufficient 10-minute power reserves [18].

The N-1 property states that the grid must withstand any single device failure. US regulators require grids to satisfy the N-1 property at all times [8] and operators often strive for even higher resiliency [47]. However, in practice the grid frequently does not satisfy even the N-1 property (i.e., is not secure against the loss of even a single element). In 2019 and 2020, the US grids had 19 and 17 EEA3 events, respectively; by definition, when those events occur the grid is not N-1 secure [26]. As an aside, we note that the N-1 property is insufficient to defend against many practical threats, including cyber attacks. For example, a substation failure should be considered an N- x event, where x is the number of devices attached to the substation; under a cyber attack, for example, it is feasible to compromise all of the devices attached to a substation [1].

Another notion of resilience is N-2, which requires the grid to be robust to the concurrent failure of the two devices that are currently carrying or producing the greatest amount of power [47]. The N-2 property does not require that the grid tolerate the failure of *any* two devices: it only needs to tolerate the failure of the two devices currently carrying or producing the most power. In theory, a grid could satisfy the N-2 property without satisfying the N-1 property: if a *single* device fails, and that device isn’t one of the two currently carrying or producing the most power, this could cause a grid failure that wouldn’t violate the N-2 property. US power grids attempt to satisfy N-2 [47], particularly in times of stability. However, just as with N-1, under EEA3 events the grid does not satisfy the N-2 property. In our experiments, we make the conservative assumption that the grid is N-1 secure, except when we simulate emergency scenarios.

TABLE 2. PROPERTIES OF A SELECT SET OF PUBLICLY AVAILABLE GRID TOPOLOGIES.

Topology	# of single device failures that cause a blackout	# line limits doubled to satisfy N-1 criteria	Topology 10-min. power reserves (% of demand)	Real-world 10-min. power reserves (% of demand)
Syn. Polish Summer Peak 2008 [38]	4195/4195	358/3693	4.4	9.0 [48]
Syn. Texas [36]	134/3749	121/3206	2.0	12.0 [18]
Syn. Western [36]	17/15190	7/12706	0.6	12.2 [18]

In the second column, the denominator is the total components in the topology. If a topology satisfies the N-1 property, the numerator should be 0. For example, on the Synthetic Polish Summer Peak 2008 topology the failure of any device will cause a blackout. The 10-minute power reserves are the reserve generation that operators can deploy within 10 minutes of an attack (i.e., droop, AGC, and responsive reserves). We compare the 10-minute power reserves of each topology to their real-world counterpart (gathered from public information [18], [48]).

Table 2 illustrates that three often-cited grid topologies (Synthetic Polish Summer Peak 2008 [38], Synthetic Texas [36], and Synthetic Western Interconnection [36]) are unrealistic because of insufficient line capacities and 10-minute power reserves. The unmodified Synthetic Polish Summer Peak 2008 topology has 21 lines exceeding their capacity ratings, causing any device failure to cause a blackout. We validated these results with PowerWorld, another commercial power-flow simulator. Additionally, although the Texas Synthetic and Western Synthetic topologies had fewer device failures cause a blackout, neither topology satisfied the N-1 property. Furthermore, all three topologies had lower power reserves than their real-world counterparts.

To make these three datasets more realistic, we modify them by adjusting line capacity ratings and increasing reserve generation. We double any line capacities that cause a blackout until the topology satisfies the N-1 property, making the topology more representative of real-world topologies. Additionally, we modify the topologies to have realistic, 10-minute power reserves. We change generator limits so that all generators have at least 5% droop reserves. Similarly, we include AGC and responsive reserve generation to match real values [42], [48]. The resulting topologies are both N-1 resilient and contain realistic reserve generation that meets common standards [8]. We use these extended topologies for the following analyses, unless otherwise specified.

Finding 4 (Scen 1): Simulations on more realistic topologies show that successful MadIoT attacks require significantly more high-wattage IoT devices to cause blackouts than suggested by prior work [3], which analyzed less realistic topologies (Fig. 7).

As an example of how unrealistic topologies impact the outcome, we focus specifically on MadIoT attacks. We assume that attackers compromise high-wattage IoT devices

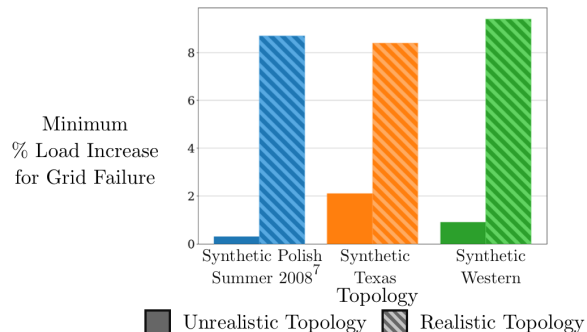


Figure 7. We simulate MadIoT attacks on the original unrealistic and the adjusted, more realistic topologies. MadIoT attacks on unrealistic topologies require a lower amount of load increase to cause a grid failure.

proportional to the power consumption of each load. For example, during a 2% demand increase attack, all load consumption increases by 2%. Furthermore, we assume that operators have 10-minute generation reserve capabilities to counteract the attack.

As shown in Fig. 7, the unrealistic Synthetic Polish topology⁷ failed after only a 0.3% increase in demand, while its more realistic counterpart requires an attacker to increase demand by 8.9%. Similar results hold for the Texas Synthetic and Western Synthetic topologies: the original and realistic Texas topologies require load increases of 2.1% and 8.3%, respectively; while the original and realistic Western topologies require load increases of 0.8% and 9.4%, respectively. Our results reconfirm the prior findings that sufficient load increase for successful MadIoT attacks in practice is likely challenging for an adversary [4]. Furthermore, we go beyond prior analyses in identifying an additional source of inconsistency: the effect of using unrealistic topologies when evaluating MadIoT attacks.

Finding 5 (Scen 2): FDIA attacks become more detectable and less impactful when the attackers' capabilities are more realistic (Fig. 8).

To show how unrealistic attack capabilities affect perceived threat feasibility, we simulate FDIA attacks on the Synthetic South Carolina topology [36].⁸ Previous work assumes that attackers have control over all system measurement devices [37], [49], [50]. As prior work noted, to bypass AC-state estimation, even for less complex attacks, the attackers would require control of almost all measurement devices [37]. There are no documented examples of real attackers conducting large-scale attacks on measurement devices. Given the diversity of operators and devices and

7. The unmodified Synthetic Polish Summer 2008 topology cascades into failure even without a MadIoT attack. Therefore, the unrealistic Synthetic Polish Summer 2008 topology refers to the unmodified topology with 21 line limits doubled.

8. The Synthetic South Carolina topology does not pass the topology-verification checks as it has insufficient 10-minute power reserves. We could not use a realistic topology because of scalability limitations of the MATPOWER state-estimation tool. Since power reserves are not relevant to FDIA attacks, we run this analysis in spite of this known limitation.

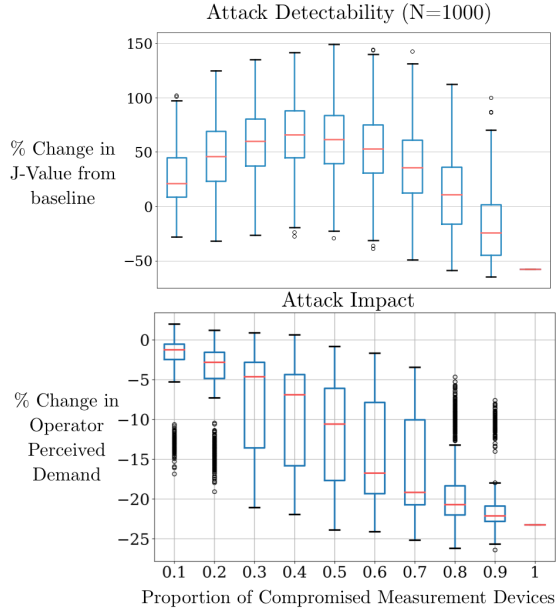


Figure 8. FDIA attacks on the Synthetic South Carolina Topology [36] with varying number of compromised measurement devices. We average the J-value of 1000 simulations without an attack occurring to calculate a baseline J-value. For each proportion of compromised devices, we randomly selected a set of compromised devices and perform 1000 simulations. J-values increase as the attacker controls fewer devices, increasing the likelihood of detection. In addition, when the attacker controls fewer devices they can exert only limited influence on the perceived state of the system (in this case, lowering the perceived load).

the grid’s distributed nature, we posit that compromising all or almost all measurement devices is unlikely.

To understand the sensitivity to attackers’ capabilities, we model FDIA attacks while varying the proportion of compromised sensors the attacker controls. For each proportion of compromised sensors, we ran $N=1000$ FDIA simulations where we randomly selected the set of compromised sensors. The attacker creates fake measurement data to trick the operators into believing that the overall system demand is 20% less than in reality. We conservatively assume the attacker has perfect system knowledge, which is in practice unlikely because the grid is constantly changing [51].

Fig. 8 shows that when the attacker controls fewer devices, operators see a greater change in J-value. Thus, *bad-data detection* algorithms are more likely to raise an alarm [44]. The change in J-value occurs because the attacker is unable to craft sufficiently realistic fake data and benign devices report data that conflicts with the attacker’s. The median change in J-value is 41.3% when the attacker controls 70% of grid measurement devices. The median J-value rises to 68.2% when the attacker controls 50% of measurement devices. Attacks cause slightly lower relative changes in J-values when the attacker controls less than 30% of devices because there is less data to manipulate and operators see more correct data. In this case, the attacker does not meet their goal of tricking the grid operator into believing the grid demand is low. When the attacker controls 50%

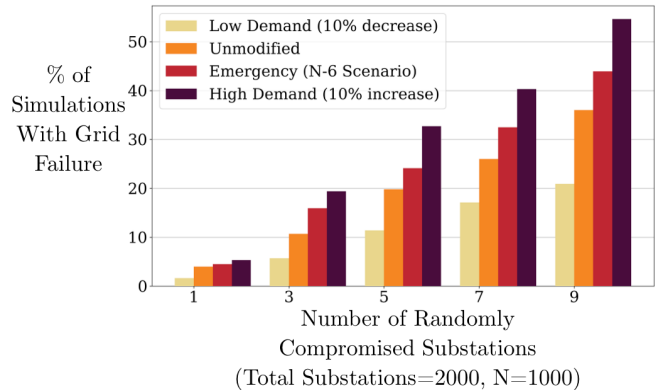


Figure 9. For each of the four scenarios, we randomly compromise substations and disconnect all their attached devices from the grid. The high-demand and emergency scenarios had a greater number of grid failures, indicating they were more susceptible to substations attacks.

of measurement devices, the median perceived decrease in load is 10.6% rather than the intended 20%.

Finding 6 (Scen 3): Substation Circuit Breaker Takeover attacks during emergency and high-demand grid scenarios require compromising significantly fewer substations for grid failure, making the attack more feasible (Fig. 9).

To confirm that different grid scenarios impact levels of susceptibility to grid threats, we focus on the Substation Circuit Breaker Takeover (SCBT) threat. Recall that during an SCBT attack, an attacker compromises a set of substations and uses the substations to disconnect loads, generators, and transmission lines from the grid [9]. In the 2015 and 2016 Ukraine attacks, adversaries compromised entire substations and opened circuit breakers (disconnecting consumer power) [1], demonstrating that compromising substations and disconnecting devices is realistic.

We model the SCBT threat as follows: We run N simulations, each randomly selecting a different set of K_{subs} compromised substations. We assume that an attacker who has compromised a substation disconnects all devices connected to the substation. We simulate SCBT attacks on the realistic Texas Synthetic topology for four scenarios: 1) unmodified, 2) emergency state (N-6), 3) high demand (10% demand increase), and 4) low demand (10% demand decrease). We create the emergency state scenario by removing six edges from the 40 highest power flow edges, motivated by North American Electric Reliability Corporation’s (NERC) emergency simulation requirements [47]. We report the fraction of these simulations that result in grid failure.

Fig. 9 illustrates the impact of SCBT attacks on the Texas synthetic topology. All four scenarios had substations that were single points of failure. In the high-demand scenario, 9.8% of the single-substation attacks caused grid failure. During the nine-substation attacks, 36.0% of simulations under the original scenario led to grid failure; under the emergency and high-demand scenarios, 43.9% and 54.6%, respectively, of simulations resulted in grid failure. This strongly suggests that grids under emergency and high-

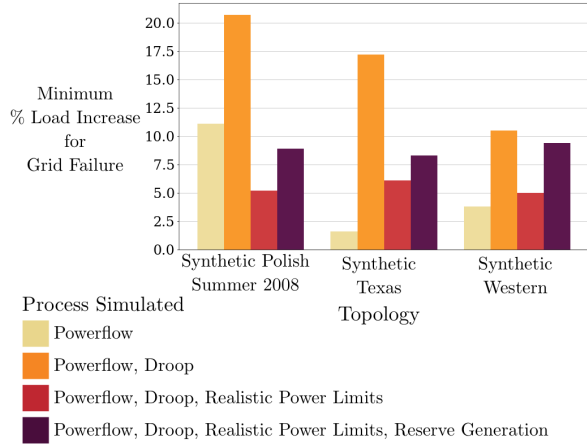


Figure 10. We simulate MadIoT with varying sets of simulated grid processes. When only simulating power flow and droop processes, the engine overestimates the demand increase required for grid failure. On the contrary, when simulating power flow, droop, and realistic power limits, the engine underestimates the demand increase required for grid failure.

demand scenarios are more susceptible to SCBT attacks.

Finding 7 (Sim 1): MadIoT attack feasibility and impact cannot be accurately evaluated unless we include all relevant processes in the simulator (Fig. 10).

To show that omitting relevant grid processes within the simulator changes our understanding of attack feasibility, we return to the MadIoT attack. We observed that different studies of MadIoT attacks considered very different sets of grid processes. For example, the original MadIoT paper simulated droop reserves but did not simulate realistic power limits (droop could change the power by any amount) or responsive reserve generation [3]. Other studies assumed operators cannot deploy responsive reserve generation because attackers would simultaneously turn on compromised high-wattage devices [4], [5]. One MadIoT follow-up study illustrates one source of modeling inconsistency with the original MadIoT paper [4]: modeling protection equipment (i.e., under-frequency load shedding). In contrast, our work identifies and models grid processes that the original paper [3] missed and that follow-up work [4] did not identify as reasons for modeling discrepancies.

We consider four sets of processes to understand whether and how modeling choices impact outcomes. We start with a baseline of power flow, and progressively add droop reserves, realistic power limits, and reserve generation (AGC and responsive reserve generation) to the simulation. For instance, droop with realistic power limits mimics the droop response of generators, which change power as a function of demand but only by up to 5% of operating power [18].

Fig. 10 shows that droop and realistic power limits are important for accurately simulating MadIoT attacks. When simulating only power flow, predictions for minimum load increase vary greatly per topology. Simulating just power flow is unrealistic because no reserves are simulated and the slack generator (an unrestricted generator used in

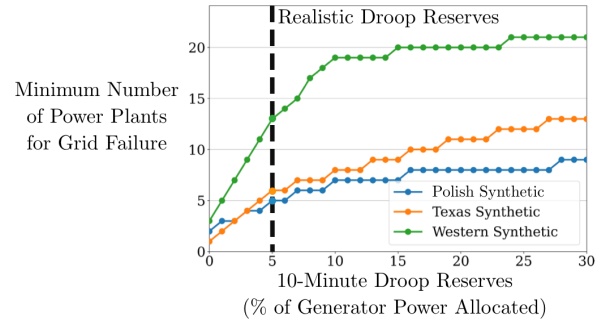


Figure 11. For each simulation, generators with the highest power output are sequentially shut off until the grid fails (we conservatively omit the five largest). A realistic amount of 10-minute droop reserves is 5% of generator operating power, but since prior work has both under- and over-estimated 10-minute droop reserves, we vary the reserves to demonstrate how these under- and over-estimations can alter the minimum number of generator failures required for grid failure.

power-flow simulation) can exceed its power limit. Once we include the droop process in the simulator, both topologies require a much higher load increase for grid collapse. Modeling the droop process makes the simulation more realistic, but results in a large overestimate of the required demand increase because simulated generators can supply more power than their limit. Next, we constrain how quickly generators, including the slack generator, can change power. As a result, the simulations show that a far lower increase in load is sufficient during MadIoT attacks to cause grid failure. Finally, we simulate responsive reserve generation, which increases the reserves, and thereby requires a higher increase in load during MadIoT attack for grid failure.

Finding 8 (Sim 2): Incorrect modeling of droop reserves makes Power Plant Takeover attacks less feasible. This is due to the unrealistic additional amount of droop power reserves during the attack (Fig. 11).

To validate that simulating processes incorrectly results in an inconsistent threat impact evaluation, we focus on the Power Plant Takeover (PPT) threat [7], [14]. During a PPT attack, an attacker compromises multiple generators. We simulate PPT by shutting off the highest-producing generators starting from the sixth largest generator. To be conservative, we assume attackers cannot compromise the five highest power-producing generators.

Prior work incorrectly modeled droop when analyzing grid threats [3], [7]. For example, one PPT study did not simulate any droop reserves [7], while another, MadIoT, study had unlimited droop reserves [3]. In the real world, droop reserve generation supplies limited amounts of power because it takes time for generators to increase power output. Generators only carry a limited amount of droop reserve for technical, regulatory, and economic reasons [18].

We simulate PPT attacks with varying amounts of droop reserves and find that if droop reserves are over-estimated, PPT attacks appear less feasible because attackers must compromise more generation plants to cause a grid failure (Fig. 11). For example, with realistic droop reserves of

5% of generators' operating power, the Synthetic Polish topology requires shutting down five generators for grid failure. When we overestimate reserves by reserving 30% of the generator operating power for droop, the Synthetic Polish topology requires shutting down ten generation plants for grid failure.

6. Limitations

We acknowledge three key limitations: 1) we only use steady-state simulations, 2) we do not explore all factors that may cause modeling inconsistencies, and 3) we only simulate threats on four emergency scenarios.

Steady-state simulations: One limitation is that we use steady-state simulations rather than dynamic models. Dynamic models analyze transient phenomena, while steady-state models are typically used in grid control rooms to evaluate the impact of contingencies (our type of analysis). While dynamic analysis might be more accurate with access to precise network topologies, data that precise is not publicly available. The state-of-the-art public grid topologies we leverage are intended to be used for steady-state simulations.

Additional modeling factors: Although we used state-of-the-art topologies and modeling techniques, there are further modeling factors that could change the outcomes of attack simulations. Specifically, using the modeling abstraction (Fig. 5), we can identify many other possible sources of inconsistency in each of the following dimensions: attack scenarios or simulators. One unexplored attack scenario factor is the realism of attack strategies. For example, it is unclear how realistic MadIoT strategies are (e.g., how quickly attackers can adjust the load of high-wattage IoT devices). Similarly, prior work only analyzes threats using transmission level topologies. The granularity of the topology may impact results too. Specifically, transmission topologies aggregate distribution buses, and meaningful information may be lost. For threats that target the distribution section, such as MadIoT, a topology containing distribution information may have more accurate results.

Comprehensive emergency scenarios: As shown by Finding 6, emergency scenarios make the grid more susceptible to cyber attacks. We evaluate threats on only four illustrative scenarios. Operators have lists of additional likely grid emergency scenarios [47]; we recommend operators simulate threats on those scenarios. Additionally, we recommend further collaboration between operators to create other emergency scenarios that pertain to cyber threats. Evaluating attacks on a comprehensive list of realistic emergency scenarios will help avoid underestimating the feasibility and impact of threats.

7. Recommendations and Open Questions

Sec. 5 validates that the factors in Sec. 4 affect our understanding of the feasibility and impact of threats. We found that each community appears to make assumptions

realistic from their own perspective but potentially unrealistic from the perspective of other fields. For example, the security community has repeatedly made unrealistic grid operational assumptions (Findings 4, 7, and 8). In contrast, the grid community often makes unrealistic security assumptions (Findings 5 and 6). Better communication and assumption standards across communities could prevent these issues. Building on these insights, we provide several actionable recommendations to avoid inconsistent threat modeling. We also advance open research questions for future work.

Evaluate threats using realistic topologies: As shown by Findings 3 and 4, current publicly available grid topologies are inadequate to evaluate the impact of grid attacks. We reiterate a similar recommendation from by Soltan et al., that operators or industry entities should consider releasing accurate topologies in a privacy-preserving manner [3]. We ensure that topologies satisfy N-1 and device parameters are realistic. As future work, we recommend the community build topologies that have other realistic properties. For example, N-2, ensuring the largest two contingencies do not cause failure, is commonly implemented on US power grids [47].

Standardize simulation tools and techniques: As shown by Findings 7 and 8, simulator issues cause many of the challenges we observed in evaluating grid threats. To overcome simulator inconsistencies we recommend two approaches: 1) create a single universal simulator capable of accurately simulating the effect of any threat or 2) create standard simulators for different types of threats.

Ideally, a universal simulator would require understanding how different sections (physical vs. operations) of the grid interact. For instance, to study the impact of FDIA on physical sections of the grid, we need to simulate 1) how operators interpret measurement data (i.e., state-estimation), 2) how operators control physical sections of the grid in response to the data, and 3) the physical sections of the grid (i.e., steady-state power flow). Current grid simulators model individual sections of the grid but not their interactions.

A universal simulator may be challenging to construct and there may exist fundamental scalability vs. fidelity tradeoffs. In the absence of such a universal approach, we need to create a more systematic way to identify the *relevant* grid processes that need to be modeled for a type of threat to get a high-confidence estimate of impact and likelihood. As Finding 7 illustrated, the set of processes modeled can impact outcomes. Less clear is how to systematically identify the set of processes that need to be simulated for a given threat to balance these scalability-fidelity tradeoffs.

Comprehensive threat evaluation: We only examined a few attack scenarios under each threat. We encourage the community to develop techniques that comprehensively evaluate grid threats. In particular, current modeling approaches and tools should be extended to be able to evaluate threats based on attacker capabilities rather than on attackers' specific strategies.

Assumed-failure cybersecurity reliability standards: Current grid cybersecurity standards prioritize the security of individual entities such as power plants, or organizations. Although these standards are important, we recommend creating standards that consider the grid as a single entity. For example, Finding 6 shows that compromising relatively few substations can lead to grid failure. However, protection measures against such events are not enforced by grid regulators. The aforementioned standard simulation tools would enable the research community to propose reasonable mitigating measures for regulatory bodies to enforce.

8. Related Work

Grid threat modeling: Prior work has modeled a variety of grid threats (e.g., [3], [6], [7], [9]). At a high level, our work takes a broader and more systematic approach across multiple threats and multiple dimensions of grid modeling (i.e., Findings 4, 5, 6, and 8). MadIoT threats in particular often have inconsistencies in threat outcomes [3], [4]. A key observation in MadIoT follow-up work [4] is that the original MadIoT paper [3] did not correctly simulate protection equipment (e.g., under-frequency load shedding). We extend this by: 1) demonstrating why topologies in [3] are unrealistic (Finding 3), 2) adjusting the topologies to be realistic (Finding 3), 3) illustrating the effects of using unrealistic topologies when evaluating MadIoT attacks (Finding 4), and 4) identifying and modeling grid processes the original paper [3] missed and follow-up work did not identify as reasons for modeling discrepancies [4].

Grid anomaly detection: Prior work has studied detecting anomalous data (e.g., FDIA) in grid control rooms [44]. Reliable grid operation depends on measurements from Supervisory Control and Data Acquisition systems, which collect: i) continuous grid-state measurements from sensors such as remote terminal units and phasor measurement units and ii) the discrete status of switching devices and circuit breakers. The Energy Management Systems within control rooms detect anomalous grid state or sensor status data. AC state-estimation with bad-data detection algorithms is used for detecting anomalous data, via hypothesis tests on the state-estimation output residuals. Additionally, topology changes are detected by the network topology processor, which uses Kirchhoff's Current Laws checks to detect faulty status readings. For more information on these processes, please refer to [44], [51]–[53].

Grid resiliency: The power grid is designed to be resilient to natural disasters, random failures, and cyber attacks. Regulatory bodies (e.g., North American Electric Reliability Corporation (NERC)) enforce reliability requirements on the bulk grid. For instance, NERC planning reliability standards (TPL standards [8]) ensure that the grid operates reliably over a broad spectrum of system conditions under various contingencies. Similarly, NERC operating standards (TOL standards [8]) ensure that grid cascading failure does not occur under the most severe single contingency or a predefined set of multiple contingencies. More recently, standards have

been created to safeguard against cyber threats [8]. However, these are limited and do not protect against complex threats. Nor are the grids protected against the loss of multiple components, which are more likely during a cyber threat.

Cyber physical systems (CPS) security: Operators rely on CPS to manage all sections of the grid. There is a large body of work studying approaches to securing CPS [54]. Techniques such as attack graphs [55] and physical simulations [56] have been applied. In addition, many are attempting to identify anomalous sensor data [57]. Researchers have also focussed on securing programmable logic controllers, a standard type of industrial computer [58].

IoT security: IoT devices have been shown to contain numerous vulnerabilities [59], [60]. Many techniques have been proposed to automatically identify vulnerabilities [59], [61]. Machine-learning techniques have been proposed to secure IoT devices at the network layer [62], [63]. Several machine-learning techniques automatically identify devices and create security policies for each device [62], [63].

9. Conclusions

This paper was inspired by the lack of consensus in results surrounding MadIoT and the debate among experts on the feasibility of this threat. We surveyed experts and found, in hindsight perhaps unsurprisingly, that confusion is not limited to MadIoT; there is a general lack of agreement among grid practitioners and academic experts on the likelihood and impact of four threats. We qualitatively observed that inconsistencies occur because of ad hoc simulation and modeling methodologies, as well as because of dataset errors. We quantitatively validated the influence of these factors by using a custom analysis framework. We posit that many of these inconsistencies stem from a lack of standardized, public toolkits and datasets, and we recommend several ways to increase the accuracy of evaluations. By shedding light on these inconsistencies and potential causes, we hope our work inspires more rigorous foundations for future grid cybersecurity research.

Acknowledgments This research was supported in part by C3.ai Inc.; Microsoft Corporation; the Combat Capabilities Development Command Army Research Laboratory under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA); and the Defense Advanced Research Projects Agency (DARPA) under contract number W912CG-21-C-0009. The views and conclusions in this paper are those of the authors and should not be interpreted as presenting the official policies or position, either expressed or implied, of C3.ai, Microsoft Corporation, DARPA, Combat Capabilities Development Command Army Research Laboratory, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

- [1] T. C. Robert Lee, Michael Assante, “Analysis of the cyber attack on the Ukrainian power grid,” *Electricity Information Sharing and Analysis Center (E-ISAC)*, 2016.
- [2] L. Lloyds, “Business Blackout: The insurance implications of a cyber attack on the US power grid,” 2015.
- [3] S. Soltan, P. Mittal, and H. V. Poor, “BlackIoT: IoT botnet of high wattage devices can disrupt the power grid,” in *27th USENIX Security Symposium*, 2018.
- [4] B. Huang, A. A. Cardenas, and R. Baldick, “Not everything is dark and gloomy: Power grid protections against IoT demand attacks,” in *28th USENIX Security Symposium*, 2019.
- [5] J. Ospina, X. Liu, C. Konstantinou, and Y. Dvorkin, “On the feasibility of load-changing attacks in power systems during the COVID-19 pandemic,” *IEEE Access*, 2020.
- [6] Y. Liu, P. Ning, and M. K. Reiter, “False data injection attacks against state estimation in electric power grids,” *ACM Transactions on Information and System Security (TISSEC)*, 2011.
- [7] Y. Yuan, Z. Li, and K. Ren, “Modeling load redistribution attacks in power systems,” *IEEE Transactions on Smart Grid*, 2011.
- [8] NERC, “All reliability standards,” 2021. [Online]. Available: <https://www.nerc.com/pa/Stand/Pages/AllReliabilityStandards.aspx>
- [9] J. Yan, Y. Zhu, H. He, and Y. Sun, “Multi-contingency cascading analysis of smart grid based on self-organizing map,” *IEEE Transactions on Information Forensics and Security*, 2013.
- [10] B. Singer, A. Pandey, S. Li, L. Bauer, C. Miller, L. Pileggi, and V. Sekar, “Grid cybersecurity framework,” <https://github.com/bsinger98/CyberGridSim>, 2022.
- [11] A. Pandey, M. Jereminov, M. R. Wagner, D. M. Bromberg, G. Hug, and L. Pileggi, “Robust power flow and three-phase power flow analyses,” *IEEE Transactions on Power Systems*, 2018.
- [12] P. Kundur, “Power system stability,” *Power system stability and control*, 2007.
- [13] NERC, “Bulk electric system definition reference document,” Tech. Rep., August 2018.
- [14] C. Bruno, L. Guidi, A. Lorite-Espejo, and D. Pestonesi, “Assessing a Potential Cyberattack on the Italian Electric System,” *IEEE Security and Privacy*, 2015.
- [15] F. Wu, K. Moslehi, and A. Bose, “Power system control centers: Past, present, and future,” *Proceedings of the IEEE*, 2005.
- [16] V. Vittal, J. D. McCalley, P. M. Anderson, and A. Fouad, *Power system control and stability*. John Wiley & Sons, 2019.
- [17] NERC, “Balancing and frequency control,” Tech. Rep., January 2011.
- [18] E. Ela, M. Milligan, and B. Kirby, “Operating reserves and variable generation,” NREL, Tech. Rep., August 2011.
- [19] P. S. Kundur and O. P. Malik, *Power system stability and control*. McGraw-Hill Education, 2022.
- [20] ERCOT, “Annual Report of Demand Response In the ERCOT Region,” December 2020. [Online]. Available: <https://www.ercot.com/services/programs/load>
- [21] B. Johnson, D. Caban, M. Krotofil, D. Scali, N. Brubaker, and C. Glyer, “Attackers deploy new ICS attack framework “TRITON” and cause operational disruption to critical infrastructure,” *Mandiant Blog*, 2017.
- [22] A. Baiocco, C. Foglietta, and S. D. Wolthusen, “Delay and jitter attacks on hierarchical state estimation,” in *2015 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2015.
- [23] S. Leppänen, S. Ahmed, and R. Granqvist, “Cyber Security Incident Report—Norsk Hydro,” 2019.
- [24] J. Minkel, “The 2003 Northeast Blackout—Five Years Later,” *Scientific American*, 2008.
- [25] NEC co-op energy. (2021) Brownout vs. Blackout: What’s the Difference? [Online]. Available: <https://neccoopenergy.com/brownout-vs-blackout-whats-the-difference/>
- [26] NERC, “2021 state of reliability,” Tech. Rep., August 2021.
- [27] M. Zeller, “Myth or reality—Does the aurora vulnerability pose a risk to my generator?” in *2011 64th Annual Conference for Protective Relay Engineers*, 2011.
- [28] T. Shekari, C. Irvine, A. A. Cardenas, and R. Beyah, “MaMIoT: Manipulation of Energy Market Leveraging High Wattage IoT Botnets,” in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021.
- [29] R. Smith, “U.S. Risks National Blackout From Small-Scale Attack,” *The Wall Street Journal*, 2014.
- [30] IEEE Power Engineering Education Committee. (2021) PowerGlobe. [Online]. Available: <https://listserv.nodak.edu/cgi-bin/wa.exe?A0=POWER-GLOBE>
- [31] ——. (2021) Cyber resilient energy delivery consortium. [Online]. Available: <https://listserv.nodak.edu/cgi-bin/wa.exe?A0=POWER-GLOBE>
- [32] NERC, “Review of bulk electric system definition thresholds,” Tech. Rep., March 2013.
- [33] ——. “Emergency operations,” EOP-011-1, 2017. [Online]. Available: <https://www.nerc.com/pa/Stand/Reliability%20Standards/EOP-011-1.pdf>
- [34] 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 Transactions on power systems*, 2010.
- [35] PowerWorld Corporation, “PowerWorld.” [Online]. Available: <https://www.powerworld.com>
- [36] A. B. Birchfield, T. Xu, K. M. Gegner, K. S. Shetye, and T. J. Overbye, “Grid structural characteristics as validation criteria for synthetic networks,” *IEEE Transactions on Power Systems*, 2017.
- [37] J. Zhao, L. Mili, and M. Wang, “A generalized false data injection attacks against power system nonlinear state estimator and countermeasures,” *IEEE Transactions on Power Systems*, 2018.
- [38] R. Korab, “Polish system during morning peak conditions in summer of 2008,” <https://eGRIDdata.org/dataset/polish-system-during-morning-peak-conditions-summer-2008>, 2018.
- [39] O. Kosut, L. Jia, R. J. Thomas, and L. Tong, “Malicious data attacks on the smart grid,” *IEEE Transactions on Smart Grid*, 2011.
- [40] M. Jereminov, D. M. Bromberg, A. Pandey, M. R. Wagner, and L. Pileggi, “Evaluating feasibility within power flow,” *IEEE Transactions on Smart Grid*, 2020.
- [41] Siemens Energy, “Pss 32.0 program operation manual,” 2021. [Online]. Available: <https://new.siemens.com/global/en/products/energy/energy-automation-and-smart-grid/pss-software/pss-e.html>
- [42] P. L. Denholm, Y. Sun, and T. T. Mai, “An introduction to grid services: Concepts, technical requirements, and provision from wind,” National Renewable Energy Lab, Tech. Rep., 2019.
- [43] A. Pandey, A. Agarwal, M. Jereminov, M. R. Wagner, D. M. Bromberg, and L. Pileggi, “Robust sequential steady-state analysis of cascading outages,” in *2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)*. IEEE, 2019.
- [44] E. Handschin, F. C. Schweppe, J. Kohlas, and A. Fiechter, “Bad data analysis for power system state estimation,” *IEEE Transactions on Power Apparatus and Systems*, 1975.
- [45] S. Li, A. Pandey, and L. Pileggi, “A WLAV-based Robust Hybrid State Estimation using Circuit-theoretic Approach,” in *2021 IEEE Power & Energy Society General Meeting (PESGM)*, 2021.

- [46] G. Hug and J. A. Giampapa, "Vulnerability assessment of AC state estimation with respect to false data injection cyber-attacks," *IEEE Transactions on Smart Grid*, 2012.
- [47] NERC, "System Performance Following Extreme BES Events," TPL-004-1. [Online]. Available: <https://www.nerc.com/files/TPL-004-1.pdf>
- [48] RAP (2018), "Report on the Polish Power System. Version 2.0 Study commissioned by Agora Energie-wende."
- [49] S. Bi and Y. J. Zhang, "Defending mechanisms against false-data injection attacks in the power system state estimation," in *2011 IEEE GLOBECOM Workshops*, 2011.
- [50] G. Liang, S. R. Weller, F. Luo, J. Zhao, and Z. Y. Dong, "Generalized FDIA-based cyber topology attack with application to the Australian electricity market trading mechanism," *IEEE Transactions on Smart Grid*, 2017.
- [51] S. Li, A. Pandey, B. Hooi, C. Faloutsos, and L. Pileggi, "Dynamic graph-based anomaly detection in the electrical grid," *IEEE Transactions on Power Systems*, 2021.
- [52] F. F. Wu and W.-H. Liu, "Detection of topology errors by state estimation," *IEEE Transactions on Power Systems*, 1989.
- [53] S. Li, A. Pandey, and L. Pileggi, "A convex method of generalized state estimation using circuit-theoretic node-breaker model," *arXiv preprint arXiv:2109.14742*, 2021.
- [54] D. Ding, Q.-L. Han, Y. Xiang, X. Ge, and X.-M. Zhang, "A survey on security control and attack detection for industrial cyber-physical systems," *Neurocomputing*, 2018.
- [55] L. Chen, D. Yue, C. Dou, J. Chen, and Z. Cheng, "Study on attack paths of cyber attack in cyber-physical power systems," *IET Generation, Transmission and Distribution*, 2020. [Online]. Available: <https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/iet-gtd.2019.1330>
- [56] R. V. Yohanandhan, R. M. Elavarasan, P. Manoharan, and L. Mihet-Popa, "Cyber-Physical Power System (CPPS): A Review on Modeling, Simulation, and Analysis With Cyber Security Applications," *IEEE Access*, 2020.
- [57] J. Goh, S. Adepu, M. Tan, and Z. S. Lee, "Anomaly detection in cyber physical systems using recurrent neural networks," in *2017 IEEE 18th International Symposium on High Assurance Systems Engineering (HASE)*, 2017.
- [58] R. Sun, A. Mera, L. Lu, and D. Choffnes, "SoK: Attacks on Industrial Control Logic and Formal Verification-Based Defenses," in *2021 IEEE European Symposium on Security and Privacy*, 2021.
- [59] E. Fernandes, J. Jung, and A. Prakash, "Security analysis of emerging smart home applications," in *2016 IEEE symposium on security and privacy (SP)*, 2016.
- [60] O. Alrawi, C. Lever, M. Antonakakis, and F. Monrose, "SoK: Security Evaluation of Home-Based IoT Deployments," in *2019 IEEE Symposium on Security and Privacy (SP)*, 2019.
- [61] Z. B. Celik, P. McDaniel, and G. Tan, "Soteria: Automated IoT safety and security analysis," in *2018 USENIX Annual Technical Conference (USENIX ATC18)*, 2018.
- [62] M. Miettinen, S. Marchal, I. Hafeez, N. Asokan, A.-R. Sadeghi, and S. Tarkoma, "IoT sentinel: Automated device-type identification for security enforcement in IoT," in *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, 2017.
- [63] T. Yu, T. Li, Y. Sun, S. Nanda, V. Smith, V. Sekar, and S. Seshan, "Learning context-aware policies from multiple smart homes via federated multi-task learning," in *2020 IEEE/ACM Fifth International Conference on Internet-of-Things Design and Implementation (IoTDI)*, 2020.

Appendix A. Survey Questions

Q1 This survey will ask you cyber security questions about the bulk energy system. Which of the following interconnections would you like to answer for? Please choose the bulk grid that you work on or the bulk grid that you are most comfortable answering questions for.

- Western Interconnection
- Texas Interconnection
- Eastern Interconnection

For the following questions, a significant disruption on the grid is either 1) 300 MW of load dropped in a dispatch cycle, 2) economic loss greater than 10 million dollars, or 3) endangerment to any number of human lives.

Q2 Do you agree that it is technically possible for a cyber attack to cause a significant disruption on the [chosen grid] within the next 5 years?

- Strongly agree
- Agree
- Neither agree or disagree
- Disagree
- Strongly disagree

Q3 Please rate your confidence of your answer to the previous question

- Very confident
- Confident
- Somewhat confident
- Not confident
- Not at all confident

Q4 Roughly, how likely is it that a cyber attack will cause a significant disruption on the [chosen grid]] within the next 5 years? [Slider 0-100%]

Q5 [same as Q3]

Q6 Do you agree that it is technically possible for a cyber attack to cause a significant disruption on the [chosen grid] within the next 15 years?

- Strongly agree
- Agree
- Neither agree or disagree
- Disagree
- Strongly disagree

Q7 [same as Q3]

Q8 Roughly, how likely is it that a cyber attack will cause a significant disruption on the [chosen grid]] within the next 15 years?

[Slider 0-100%]

Q9 [same as Q3]

Q10 Please leave any additional comments below

[Free response]

Assume the worst feasible attack has taken place.

Q11 If a cyber attack caused a significant disruption on the [chosen grid], roughly how much of the grid would go down (In terms of percent of total [chosen grid])? Please give your best estimate

[Slider 0-20% $\pm 5\%$]

Q12 [same as Q3]

Q13 If a cyber attack caused a significant disruption on the [chosen grid], what is your best estimate for how long it would take to restore power to greater than 90% of the customers that lost power?

- Under an hour
- 1 - 23 hours
- 1 - 6 days
- 1 - 3 weeks
- 1 - 11 months
- A year or greater

Q14 [same as Q3]

Q15 Please leave any additional comments below

[Free response]

We asked experts the following questions for each of the following attacks: MadIoT, FDIA, PPT, and SCBT. Before asking the questions, we gave experts a brief summary of the attack.

Q16 Have you heard of this attack before?

- Yes
- No
- Not sure
- Prefer not to answer

Q17.A Is this attack against the [chosen grid] technically feasible today?

- Definitely yes
- Probably yes
- Probably no
- Definitely no

If no selected: **Q17.B** Please explain why
[Free response]

Q18 [same as Q3]

Q19 Please leave any additional comments below
[Free response]

Please assume the most dangerous feasible version of the attack. Assume the following:

Normal Condition: Physical Responsive Capability (PRC) is at normal capacity and grid frequency is stable at 60Hz.

Emergency Level 2: PRC declines and won't be recovered for at least 30 minutes. In addition, system frequency starts to change significantly.

Emergency Level 3: PRC is drastically low and system frequency is unable to stabilize after a significant delay. Load shedding occurs at this emergency level.

A *significant disruption* on the grid is either 1) 300 MW of load dropped in a dispatch cycle, 2) economic loss greater than 10 million dollars, or 3) endangerment to any number of human lives.

Q20 If the [chosen grid] is in normal condition, roughly how likely is it that this attack will cause a significant disruption?
[Slider 0-100%]

Q21 [same as Q3]

Q22 If the [chosen grid] is in Emergency Level 2, roughly how likely is it that this attack will cause a significant disruption?
[Slider 0-100%]

Q23 [same as Q3]

Q24 If the [chosen grid] is in Emergency Level 3, roughly how likely is it that this attack will cause a significant disruption?
[Slider 0-100%]

Q25 [same as Q3]

Q26 Please leave any additional comments below
[Free response]

Q27 If the attack causes a significant disruption, for each of the following grid states, what would be the maximum load dropped from the [chosen grid] as a direct result of the attack?

For each grid state in: normal condition, Emergency Level 2, and Emergency Level 3:

Outage size

- 0MW
- 1 - 10MW
- 10-99MW
- 100-999MW
- Greater than 1000MW

Confidence [same as Q3]

Q28 If the attack causes a significant disruption, for each of the following grid states, how long do you expect it to take to restore power to at least 90% of the customers that lost power?

For each grid state in: normal condition, Emergency Level 2, and Emergency Level 3:

Restoration time

- No disruption would occur
- Under an hour
- 1 - 23 hours
- 1 - 6 days
- 1 - 3 weeks
- 1 - 11 months
- A year or greater

Confidence [same as Q3]

Q29 Please leave any additional comments below

Q30 Roughly, how likely is it that this attack will cause a significant disruption on the [chosen grid] within the next 5 years? Please give your best estimate

[Slider 0-100%]

Q31 [same as Q3]

Q32 Roughly, how likely is it that this attack will cause a significant disruption on the [chosen grid] within the next 15 years? Please give your best estimate
[Slider 0-100%]

Q33 [same as Q3]

Q34 How much do you agree with the following statement "Grid operators are aware of the possibility of this kind of attack?"

- Strongly agree
- Agree
- Neither agree or disagree
- Disagree
- Strongly disagree

Q35 [same as Q3]

Q36 Are there any other grid conditions (e.g., wildfires impacting distribution, a series of generators fail) that are particularly worrisome for this type of attack?
[Free response]

Q37 If the attack were to succeed on the [chosen grid], roughly how likely is this attack to result in permanent grid device damage (i.e., a device that has to be replaced after the attack)? Please give your best estimate
[Slider 0-100% ±5%]

Q38 [same as Q3]

Q39 If you believe the attack could result in permanent grid device damage, what devices would likely be damaged?
[Free response]

Q40 Do you have any recommended changes (e.g., adding certain types of generators to the grid, implementing better state estimation, adding additional regulations, etc) to improve resiliency against this attack?
[Free response]

Q41 Please leave any additional comments below
[Free response]

End of attack questions.

Q42.A Have any of the grid organizations that you have been a part of ever been a target of a cyber attack that has impacted power delivery?

- Yes
- No
- Prefer not to answer
- Not sure

If yes to **Q42.A**: If your organization has been attacked multiple times, please apply the following questions to the most significant attack.

Q42.B Please describe the type of attack (e.g., an attacker infiltrates a corporate network, or an attacker takes control of a generator):
[Free response]

If yes to **Q42.A**: **Q42.C** How many of your customers were impacted?
[Slider 0-100% ±5%]

If yes to **Q42.A**: **Q42.D** How long did it take to restore power to greater than 90% of the customers whom lost power?

- The attack did not impact power delivery
- No disruption would occur
- Under an hour
- 1 - 23 hours
- 1 - 6 days
- 1 - 3 weeks
- 1 - 11 months
- A year or greater
- Prefer not to answer

Q43 What is your gender?

- Male
- Female
- Non-binary / third gender
- Prefer not to say

Q44 How old are you?

- 18 - 24
- 25 - 34
- 35 - 44
- 45 - 54
- 55 - 64
- 66+
- Prefer not to say

Q45 What is your highest degree?

- Less than a high school diploma
- High school degree or equivalent (e.g., GED)
- Some college, no degree
- Associate degree (e.g., AA, AS)
- Bachelor's degree (e.g., BA, BS)
- Master's degree (e.g., MA, MS, MEd)
- Professional degree (e.g., MD, DDS, DVM)
- Doctorate (e.g., PhD, EdD)
- Prefer not to answer

Q46 How many years have you worked in the grid industry or on topics related to the grid?

- 0 - 5 years
- 6 - 10 years
- 11 - 15 years
- 16 - 20 years
- 20+ years
- Prefer not to say

Q47.A Which of the following terms best describes the organization you work for now or your most recent form of employment related to the electric grid?

- Independent System Operator (ISO), Regional Transmission Organization (RTO) or Transmission System Operator (TSO)
- Utility Company
- University
- Research Organization
- Consultant Firm
- Federal Regulator or Public Utility Commission (PUC)
- Grid Product Manufacturer (e.g., Siemens, Schneider, etc)
- Other
- Prefer not to say

If other: **Q47.B** Please explain?

[Free response]

Q48 How many people does your organization employ?

- 1 - 49
- 50 - 999
- 1000 - 4999
- 5000+
- Prefer not to say

Q49 Check all of the following that apply to your current employer

- Has an internal red team
- Has an internal cybersecurity department
- Has an internal IT department
- Has an internal blue team
- Hires external cybersecurity consultants

Q50 What is the title of your job?

[Free response]

Q51 Which best describes your professional specialty?

- Cybersecurity
- Power systems
- Both cybersecurity and power systems
- Other (free response)
- Prefer not to answer

Appendix B. Survey Results

Fig. 12 is a detailed breakdown of expert confidence for likelihood, outage size, and outage duration for all grid threats. The majority of experts reported they were somewhat confident or less in their predictions.

We asked experts several general cybersecurity questions about the grid (Q2-Q14) and the results can be found in Table 3.




Figure 12. Expert confidence in their predictions for likelihood, outage size, and outage duration for all grid threats. The majority of experts reported somewhat confident or less in their predictions for all threats.

TABLE 3. GENERAL GRID CYBERSECURITY QUESTION AND ANSWERS

Question	Answers
Q2: Is it technically possible for attack to occur within 5 years?	Strongly agree: 10, Agree: 7, Neither agree or disagree: 0, Disagree: 1, Strongly disagree: 0
Q3: Confidence of Q2	Very confident: 7, Confident: 7, Somewhat confident: 7, Not confident: 0, Not at all confident: 0
Q4: Likelihood attack to occurring within 5 years	Average: 44.1, Median: 47.5, Min: 10.0, Max: 90.0
Q5: Confidence of Q4	Very confident: 5, Confident: 3, Somewhat confident: 6, Not confident: 3, Not at all confident: 0
Q6: Is it technically possible for attack to occur within 15 years?	Strongly agree: 12, Agree: 6, Neither agree or disagree: 0, Disagree: 0, Strongly disagree: 0
Q7: Confidence of Q6	Very confident: 11, Confident: 2, Somewhat confident: 4, Not confident: 0, Not at all confident: 0
Q8: Likelihood attack to occurring within 15 years	Average: 60.2, Median: 61.5, Min: 10.0, Max: 100.0
Q9: Confidence of Q8	Very confident: 6, Confident: 5, Somewhat confident: 4, Not confident: 2, Not at all confident: 0
Q11: Percent of grid to go down from cyber attack	Average: 8.4, Median: 5.0, Min: 3.0, Max: 20.0
Q12: Confidence of Q11	Very confident: 2, Confident: 7, Somewhat confident: 4, Not confident: 4, Not at all confident: 0
Q13: Grid cyber attack recovery time of grid	Under an hour: 0, 1 - 23 hours: 7, 1 - 6 days: 6, 1 - 3 weeks: 2, 1 - 11 months: 3, A year or greater: 0
Q14: Confidence of Q13	Very confident: 4, Confident: 6, Somewhat confident: 5, Not confident: 2, Not at all confident: 0