Tolerating Compound Threats in Critical Infrastructure Control Systems

Sahiti Bommareddy[‡*], Maher Khan[†*], Huzaifah Nadeem[†*], Benjamin Gilby[†], Imes Chiu[§],
John W. van de Lindt[¶], Omar Nofal[‖], Mathaios Panteli[††], Linton Wells II[‡‡], Yair Amir[‡], Amy Babay[†]
†University of Pittsburgh, ‡Johns Hopkins University, §Defense Logistics Agency, ¶Colorado State University,
‖Florida International University, ††University of Cyprus, ‡‡George Mason University

***Abstract*—Compound threats, in which cyberattacks are targeted in the aftermath of a natural hazard, pose an important emerging threat for critical infrastructure. In this paper, we analyze the system design implications of compound threats for power grid SCADA systems for the first time. We introduce a novel compound threat model and develop a tool for analyzing resilience under this threat model. Using our tool, we compare the resilience of existing fault- and intrusion-tolerant SCADA system architectures in case studies based on two power utilities: Hawaiian Electric (HECO) and Florida Power & Light (FPL).**

We show that no existing system architecture adequately addresses compound threats, but that it is possible to improve resilience to such threats by explicitly considering natural hazards in the system design and by employing a new out-of-band reconfiguration mechanism for intrusion-tolerant systems. However, an important outcome of our work is that compound threats remain a challenging problem, with no complete solution.

I. Introduction

Compound threats, in which cyberattacks compound the damage caused by natural hazards, are an increasing concern for critical infrastructure [1]–[3] and are attracting government attention. During the 2021 Multi-State Information Sharing and Analysis Center annual exercises, ten scenarios involved planning for the dual impact of a cyberattack and a natural disaster [1]. After Hurricane Harvey in 2017, the U.S. Army Cyber Institute conducted a three-day drill simulating a cyberattack during a hurricane [1]. Cyberattacks on hospitals are also growing [2], including attacks during the COVID-19 pandemic [3]. There is an emerging trend of sophisticated cyberattacks targeting infrastructure that is already damaged or stressed by natural threats.

Power grids are an important potential target for compound threats, as other critical infrastructures rely on power to function. While the distributed systems research community has designed *intrusion-tolerant* Supervisory Control and Data Acquisition (SCADA) systems based on Byzantine Fault Tolerant (BFT) replication to withstand cyberattacks [4]–[7], the threat model these systems target does not consider natural hazards that may induce geographically dependent correlated failures. Understanding and mitigating the impacts of compound threats requires interdisciplinary expertise to model the effects of both natural hazards and cyberattacks on SCADA systems, and to design systems capable of withstanding such threats.

[*]These authors contributed equally to this work.

To address this gap, we introduce a new compound threat model that jointly captures natural hazard and cyberattack impacts. We develop a tool, the Compound-Threat Analyzer [8], to analyze system resilience under this model, and investigate design tradeoffs in the context of power grid SCADA systems. The Analyzer combines natural hazard realizations based on real hazard data and civil engineering expertise with cyberattack models from the distributed systems research community to compare their effects on different SCADA architectures.

Using the Compound-Threat Analyzer, we investigate power-grid SCADA resilience to compound threats involving a hurricane and cyberattack in case studies based on Hawaiian Electric (HECO) and Florida Power & Light (FPL). Our results show that no existing SCADA architecture provides an adequate solution to compound threats. Employing intrusion-tolerant architectures that can withstand cyberattacks and selecting site locations that provide effective fault tolerance under natural hazards are necessary, but not sufficient. To improve resilience to compound threats, we propose a new out-of-band reconfiguration mechanism. We implement this mechanism in the Spire intrusion-tolerant SCADA system [9] and demonstrate that it can recover from threat scenarios that would otherwise lead to extended service outages. However, we show that there is still no complete solution for compound threats, and further investigation and investment in infrastructure is needed. The contributions of our work are:

- We introduce a novel compound threat model that jointly considers the effects of natural hazards and cyberattacks.
- We develop the Compound-Threat Analyzer for analyzing system resilience to compound threats.
- We analyze SCADA resilience to compound threats in the context of two power utilities: Hawaiian Electric and Florida Power & Light.
- We demonstrate that no existing SCADA architecture, including state-of-the-art intrusion-tolerant architectures, adequately addresses the compound threat problem.
- We introduce out-of-band reconfiguration to mitigate compound threats, and implement it in the Spire intrusion-tolerant SCADA system.

Beyond our generic system design contributions, our work can directly inform the construction of compound-threat-resilient infrastructure in Hawaii and Florida, two areas highly susceptible to such threats [10], [11].

II. Background and Related Work

Intrusion-Tolerant SCADA. Because of SCADA systems' critical role in grid operations, and the fact that they are increasingly targeted by attackers [12], the computer systems research community has designed *intrusion-tolerant* SCADA systems that can withstand intrusions of one or more SCADA control servers [4]–[7]. These works employ Byzantine Fault Tolerant (BFT) State Machine Replication (SMR) [13] of the SCADA control server (SCADA master) to create SCADA systems that can operate correctly even if up to f of the SCADA master replicas are compromised by an attacker. Typical BFT SMR protocols require $3f + 1$ total replicas to overcome f compromises (e.g. 4 replicas for $f = 1$). To support *proactive recovery*, where replicas are periodically restarted from a known good state to remove (potentially undetected) compromises, systems with $3f + 1$ replicas can be adapted to use $3f + 2k + 1$ replicas, where k is the maximum number of replicas that may be simultaneously unavailable due to proactive recovery (e.g. 6 replicas for $f = 1, k = 1$) [14].

Prior work has tackled the practical challenges of integrating BFT SMR into real SCADA systems. Zhao et al. use the PBFT protocol [13] to overcome one compromised SCADA master [4]; Kirsch et al. use the Prime BFT replication engine [15] to create an intrusion-tolerant prototype of a Siemens SCADA product [5]; Nogueira et al. integrate the BFT-SMaRt replication library [16] with the open-source EclipseNeoSCADA [7]; and Babay et al. focus on meeting strict SCADA latency requirements in a system designed from the ground up to support intrusion tolerance [6].

The work in [6] extended the BFT threat model to include network denial-of-service (DoS) attacks, and analyzed the resilience of industry-standard and intrusion-tolerant SCADA architectures under a threat model in which an attacker can compromise f servers and isolate one site via DoS. To improve resilience under this threat model, the work introduced new multi-site network-attack-resilient SCADA architectures (e.g. the "6+6+6" configuration we discuss in Section IV-A) [6]. We build on this work; our threat model includes system compromises and network attacks (Section III), and we analyze a subset of the architectures considered in [6] (Section IV-A). However, all prior work on intrusion-tolerant SCADA considered cyberattacks only; we tackle a compound threat model that includes both cyberattacks and natural hazards. While failures due to natural hazards may appear to be subsumed by the BFT model (since it tolerates arbitrary behavior), they induce geographically dependent correlated failures that have not been considered in the BFT literature.

Power Grid Threat Modeling and Risk Analysis. The risk assessment literature examines natural hazard impacts on power grids [17], [18], as well as cybersecurity risks to their control systems [19], [20], but, to the best of our knowledge, only Avraam et. al. [21] jointly consider both types of risks. Some work analyzes failures affecting multiple dependent infrastructures (e.g. SCADA and communications networks), which are sometimes referred to as compound threats, but these differ from the compound threats we consider [22].

Risk assessment typically aims to determine the likelihood of undesirable events and quantify their impact (e.g. in terms of monetary loss or unserved load for the power grid). To quantify risk, such assessments typically focus on detailed modeling of a specific grid's full infrastructure (e.g. SCADA, substations, generators, transmission lines) to analyze power outages resulting from specific events (e.g. [18], [21]). In contrast, we focus on comparing the resilience of different SCADA architectures, since our goal is to gain insight into how to design SCADA systems to withstand compound threats, and into the new challenges this threat model brings compared with withstanding cyberattacks alone.

BFT SMR Reconfiguration. Prior work has studied *in-band SMR reconfiguration*, where an operational system is reconfigured with no loss of state [23]. In-band reconfiguration requires a quorum of the old configuration to remain up and correct until the logical time to change configurations is agreed on (typically by ordering a reconfiguration command through the SMR protocol), and the new configuration retrieves any necessary state and takes over. Some BFT SMR systems use manual reconfiguration, where a trusted administrator issues commands to add or remove replicas [16], [24], and some introduce automatic reconfiguration triggered by replicas [25]–[27], clients [28], [29], or a trusted control plane [24], [30], [31]. However, in automatic reconfiguration, compromised entities may try to exploit the reconfiguration mechanism, e.g., to cause denial of service by forcing frequent reconfiguration, or to violate safety by reconfiguring to a set of replicas that the attacker (stealthily) controls. Systems that consider this issue require a quorum of replicas to vote to remove a replica [27], [28] (although the voting criteria is left undefined) or restrict which replicas a given replica can suspect and how often they can do so [26]. We show that the compound threat model raises the need for a new type of *out-of-band reconfiguration*. Natural hazards may make a large part of the system unavailable all at once, making it impossible to perform in-band reconfiguration, since no quorum exists. Out-of-band reconfiguration can be used to recover from such outages.

III. System and Threat Model

A. System Model

SCADA systems are responsible for the monitoring and control of power grid infrastructure. Figure 1 shows a basic SCADA architecture. Control servers called *SCADA masters* (SMs) are located in *control centers*. SMs collect data from and issue commands to equipment located in *remote sites* (e.g. power plants, substations). These remote sites are geographically distributed and communicate with the SMs over a wide-area network. SMs present the state of the system to human operators through a *human machine interface* (HMI).

B. Compound Threat Model

We define a compound threat model with two stages. In the first stage, a natural hazard may make one or more SCADA control sites unavailable. For example, a hurricane

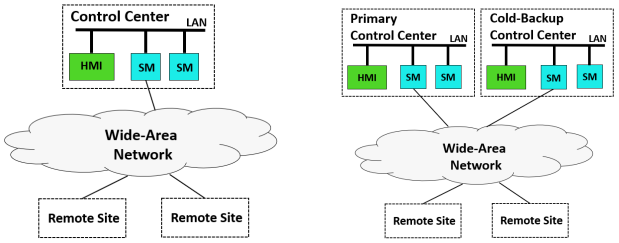

(a) Configuration "2" (b) Configuration "2-2"

Fig. 1: Current Industry-Standard SCADA Architectures

may cause flooding that prevents a control center from operating. In the second stage, a cyberattack occurs. In our analysis, we consider system-level intrusions and network-level attacks (similar to [6]). Specifically, we consider: *server intrusions* and *site isolations*. In a server intrusion, the attacker compromises a SCADA master, gaining control of it and causing it to behave arbitrarily (Byzantine failure). In a site isolation attack, the attacker isolates a control site from the rest of the network, preventing it from communicating with the other sites. This can be achieved through sophisticated resource-intensive denial of service attacks [32], [33]. Our full compound threat model includes four threat scenarios:

- **Natural Hazard:** This is a baseline scenario where some control sites may be rendered non-operational due to the natural hazard, but there is no cyberattack.
- **Natural Hazard + Server Intrusion:** The attacker can compromise f SCADA masters after the natural hazard.
- **Natural Hazard + Site Isolation:** The attacker can isolate s control sites after the natural hazard.
- **Natural Hazard + Server Intrusion + Site Isolation:** The attacker can compromise f SCADA masters and isolate s control sites after the natural hazard.

We start our analysis with the worst-case cyberattack model normally assumed in BFT systems, where the attacker can select the f servers and s sites that cause the most damage to the system. We assume the attacker can observe or forecast the natural hazard outcome, e.g. which control sites become unavailable, and can execute their cyberattacks to maximize damage to the remaining system.

However, the compound threat model raises an important new consideration regarding the attacker model. Natural hazards can create asymmetry where certain sites become more valuable than others: attacking sites that have already failed is not useful, and the system becomes more dependent on the surviving sites. In this context, a worst-case analysis gives the attacker additional power: being able to isolate one of two power-grid control centers is a much weaker capability than the power to successfully target *the* remaining control center in the aftermath of a natural hazard. For example, connectivity for different sites may be provided by different sets of ISPs, where an attacker may be able to successfully attack certain ISPs but not others. Or, due to the complex, long-running nature of the system, servers may run different software versions, or have configuration differences that affect their vulnerability.

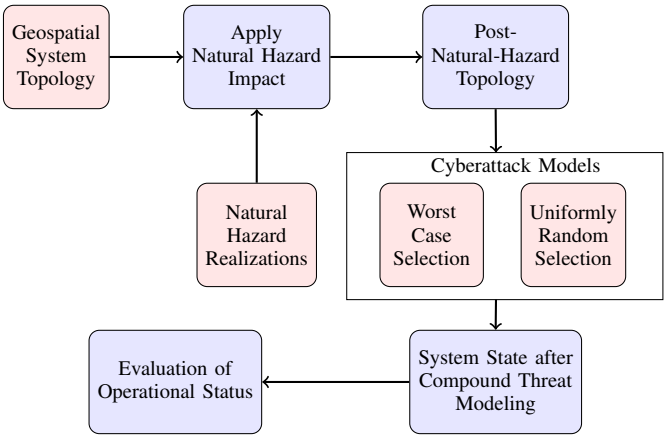

Fig. 2: Compound Threat Analyzer workflow

It is not always true that an attacker has the power to compromise/isolate the worst-case system components.

Therefore, we propose a variant of the traditional threat model, where the f servers and s sites are selected uniformly at random. Arguably, a realistic attacker power will fall somewhere on the spectrum between the worst-case and uniformly random selection, and new insights can be gained by examining the difference between them under the compound threat model. Thus, our full compound threat model considers cyberattacks with both **worst-case selection** and **uniformly random selection** of the target components.

IV. Compound-Threat Analyzer

Our Compound-Threat Analyzer integrates realistic models of natural hazard effects with cyberattack models, using the workflow in Figure 2. The Analyzer takes in a "Geospatial System Topology" specifying the locations of relevant SCADA assets, and a set of "Natural Hazard Realizations" detailing the impact of a natural hazard on those assets. Each natural hazard realization represents a specific hazard outcome, in which some sites may be rendered non-operational. This abstraction allows us to use realistic hazard models to generate the needed realizations (see Section V-B). Based on these inputs, the Analyzer determines the "Post-Natural-Hazard Topology" for each hazard realization, which specifies which sites and servers remain operational after the hazard impact.

Next, the Analyzer determines the impact of a cyberattack for each post-hazard topology instance. We consider all four threat scenarios described in Section III: Natural Hazard, Natural Hazard + Server Intrusion, Natural Hazard + Site Isolation, and Natural Hazard + Server Intrusion + Site Isolation. For each threat scenario, the Analyzer applies both cyberattack models described in Section III: "Worst Case Selection" and "Uniformly Random Selection". For each combination of threat scenario and cyberattack model, the output is the "System State after Compound Threat Modeling", a representation of which sites and servers remain available, and which of the available servers (if any) are compromised. Finally, based on this system state, the Analyzer evaluates the "Operational Status" of the system, i.e. whether it remains correct (the

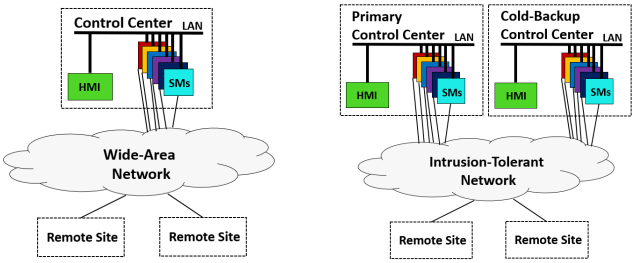

(a) Configuration "6" (b) Configuration "6-6"

Fig. 3: Intrusion Tolerant Architectures

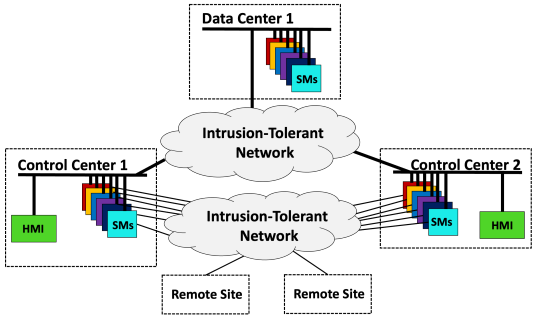

Fig. 4: Network-Attack-Resilient Intrusion-Tolerant Architecture with Configuration "6+6+6"

attacker cannot force incorrect actions) and available (the attacker has not taken the system down).

Note that while the Compound-Threat Analyzer applies the cyberattack after the natural hazard, the outcome does not change if the attack occurs before or during the hazard, as the final system state will be the same.

A. Supported SCADA Architectures

Prior work analyzed the resilience of a range of SCADA architectures to cyberattacks [6]. Our Compound-Threat Analyzer supports the three architecture classes from that work ("Industry Standard", "Intrusion-Tolerant", and "Network-Attack-Resilient Intrusion-Tolerant"); our analysis considers the representative examples from each class described below.

Industry Standard Architectures: The simplest configuration we consider is a single-control-center architecture with primary and hot-backup SCADA masters, as shown in Figure 1a. Following the notation in [6], we label this as configuration "2" (indicating a single site with 2 SMs). This architecture withstands crashes of the primary SM (by activating the hot-backup), but is not designed to withstand natural disasters or cyberattacks.

To address control center failures, many modern SCADA systems use a primary-backup architecture as shown in Figure 1b. We label this as configuration "2-2" (indicating two sites, each with 2 SMs). This configuration consists of a pair of SMs (primary and hot-backup) running in the primary site, and a second pair of SMs in a backup site. If the primary control center fails or becomes unavailable, then the backup can be brought online, although there is a delay (on the order of minutes) to activate the backup site.

Intrusion-Tolerant Architectures: As discussed in Section II, the research community has introduced SCADA architectures that use BFT SMR to tolerate server intrusions [5], [7]. We consider the configuration "6" (Figure 3a), which uses 6 SM replicas in a single control center to guarantee correctness in the presence of one server intrusion while supporting one SM undergoing proactive recovery [14].

To recover from a control center failure while supporting intrusion tolerance, the "6-6" configuration (Figure 3b) uses intrusion-tolerant replication with 6 replicas in the primary site, and has a backup control center with 6 additional replicas [6]. Here, the primary site can tolerate a server intrusion just like configuration "6", and if the primary site becomes unavailable, then the backup site, with its own 6 replicas, can be activated (after a delay).

Network-Attack-Resilient Intrusion-Tolerant Architecture: Finally, we consider the configuration "6+6+6" (Figure 4) from [6]. This configuration is designed to simultaneously withstand a site isolation and a server intrusion, without incurring downtime to activate a backup site. It includes 6 SM replicas in each of the two control centers, and an additional 6 replicas in a data center; all replicas actively participate in the BFT SMR protocol. The additional site ensures that even if one site becomes unavailable, there are enough remaining replicas to operate without interruption.[1] This architecture uses a data center because it is not typically feasible to build and operate more than two control centers with full capabilities for controlling devices in remote sites due to cost [6].

B. Evaluating Compound Threat Outcomes

For its "Evaluation of Operational Status", the Compound-Threat Analyzer adopts the color-based naming scheme used in [6] for cyberattack impacts. For the architectures and threat model we consider, there are four possible operational states:

A **green** state indicates a fully operational system.

An **orange** state occurs in primary-backup architectures when the primary control center becomes unavailable, causing downtime until the backup control center is activated.

A **red** state indicates that the system is not operational and will not be able to resume operation until some system components are repaired, or an attack ends or is mitigated.

A **gray** state indicates that the attacker has compromised the safety of the system and can cause it to behave incorrectly.

For each SCADA configuration, the Analyzer evaluates which operational state a particular hazard+attack realization causes according to the rules summarized in Table I.

C. Attack Modeling

We consider an attacker with full knowledge of the SCADA system and its possible operational states. The attacker's goal is to force the system to behave incorrectly if possible and

[1]The BFT replication protocol requires that $\lfloor \frac{n+f}{2} \rfloor + 1$ replicas are up and correct to operate, where n is the total number of replicas and f is the number of tolerated intrusions. In this case, we have $\lfloor \frac{18+1}{2} \rfloor + 1 = 10$, so a single site with 6 replicas cannot operate on its own, but two sites with 12 replicas can continue to operate even if one replica is compromised and one is down for proactive recovery (see [6] for discussion).

	green	orange	red	gray
"2"	control center up, no intrusion	N/A	control center down/isolated	server intrusions ≥ 1
"2-2"	primary control center up, no intrusion	primary control center down/isolated, backup control center up, no intrusion	primary control center down/isolated, backup control center down/isolated	server intrusions ≥ 1
"6"	control center up, server intrusions ≤ 1	N/A	control center down/isolated,	server intrusions ≥ 2
"6-6"	primary control center up, server intrusions ≤ 1	primary control center down/isolated, backup control center up, server intrusions ≤ 1	primary control center down/isolated, backup control center down/isolated	server intrusions ≥ 2
"6+6+6"	at least 2 sites up, server intrusions ≤ 1	N/A	less than 2 sites up, server intrusions ≤ 1	server intrusions ≥ 2

TABLE I: Rules for the operational state for each SCADA configuration. Note that server intrusions only include intrusions in sites that are up. Intrusions in a site flooded by a hurricane do not count against the total number of tolerated intrusions.

to halt operations as long as possible otherwise. Thus, for an attacker, the outcome ranking is: gray > red > orange > green.

An attacker may choose not to execute all of the attacks in their power if doing so would lead to a lower-ranked outcome. In particular, an attacker who compromises enough servers to cause a gray state will not execute a site isolation attack, since it would only reduce their own ability to maliciously control the system. When executing a site isolation, the attacker's optimal strategy is to rank sites as: primary control center > backup control center for configurations "2-2" and "6-6", and control center > data center for configuration "6+6+6".

As discussed in Section III, we let f be the number of servers the attacker can compromise and s be the number of sites the attacker can isolate. In our analysis, we consider the cases of $f = 0, s = 0$ (no attack), $f = 1, s = 0$ (server intrusion), $f = 0, s = 1$ (site isolation), and $f = 1, s = 1$ (server intrusion + site isolation).

D. Calculating Outcome Probabilities

Because natural hazard outcomes are nondeterministic, the compound threat outcomes are also probabilistic. In our framework, this nondeterminism is captured by generating a set of realizations for each hazard scenario (in our analysis, we generate 1000 realizations per hurricane; see Section V-B). To analyze a given SCADA configuration, the Compound-Threat Analyzer applies the specified cyberattack scenario to each hazard realization and calculates the resulting operational state (i.e. green, orange, red, or gray).

For the **worst-case selection cyberattack model**, the attack effect for each hazard realization is deterministic: the attacker evaluates every possible attack within their power and chooses the up to f servers and up to s sites that cause the most damage, according to the outcome ranking gray > red > orange > green. Thus, for each hurricane, SCADA configuration, and attack scenario (no attack, server intrusion, site isolation, server intrusion + site isolation), we have 1000 hurricane + attack outcomes. The operational state for each outcome is calculated according to the rules in Table I. We then calculate the probability of each operational state as the fraction of the 1000 outcomes in which that state occurs.

For the **uniformly random selection cyberattack model**, attack effects are randomized. For each SCADA configuration, we generate 100,000 random cyberattack realizations and apply this set of cyberattack realizations to each hurricane realization (for a total of 100 million hurricane+attack realizations). For each attack realization, we select f servers and s sites uniformly at random from the total set of servers and sites; these are the assets that the attacker has the power to compromise/isolate in that realization. Then, we apply each attack realization to each hurricane realization: first the attacker compromises all servers that they have the power to in that attack realization; if the resulting state is gray, they do not isolate any sites; otherwise, they then isolate all sites that they have the power to in that attack realization. If a server/site that the attacker has the power to compromise/isolate is non-operational due to the hurricane, the attack has no effect on that asset. For each hurricane, SCADA configuration, and attack scenario, we have 100 million hurricane + attack outcomes. The operational state for each outcome is calculated according to the rules in Table I, and we then calculate the probability of each operational state as the fraction of the 100 million outcomes in which that state occurs.

V. Analyzing Architectural Tradeoffs Under Compound Threats: Case Studies

The compound threat model is motivated by the need for extreme resilience in strategically important systems that represent high-value targets for attackers and are susceptible to natural hazards. Power grid infrastructure in Hawaii and Florida are two such systems for the US, where disruptions not only affect substantial local populations, but also impact national security interests (e.g. the base of Indo-Pacific Command for Hawaii). Both Hawaii and Florida are susceptible to hurricanes and their associated flooding. Our analysis not only provides insight into architectural tradeoffs between different SCADA system configurations, but can serve as a concrete basis for system design decisions in Hawaii and Florida.

A. Power Topologies

The combination of a SCADA configuration and its mapping to physical sites represents the "Geospatial System Topology" used in the Compound-Threat Analyzer (Figure 2).

Oahu, Hawaii. The power grid SCADA topology in Oahu currently includes a control center, power plants, and substations, as shown in Figure 5. For single-control-center configurations ("2" and "6"), we locate the control center in

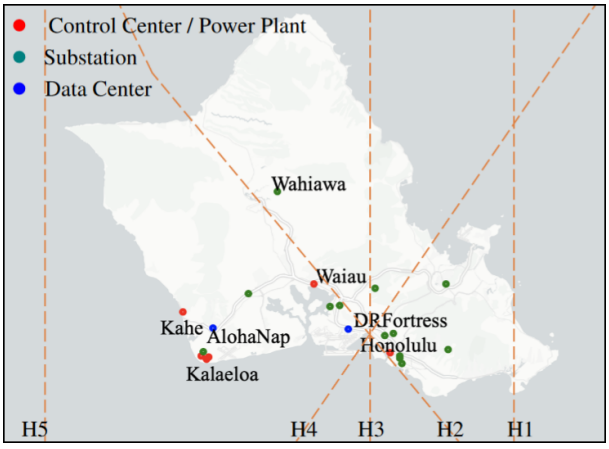

Fig. 5: Oahu, Hawaii power topology with 5 hurricane paths (H1- H5)

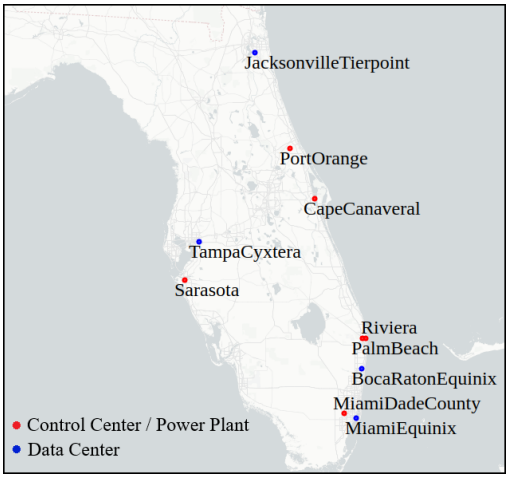

Fig. 6: Florida power topology.

Honolulu (matching the current situation for HECO). For primary-backup configurations ("2-2" and "6-6"), we locate the primary control center in Honolulu, and the backup control center in Waiau due to its geographically central location with good connectivity. Our communications with HECO engineers confirm that Waiau is under consideration as a backup control center site. For the additional data center required by the "6+6+6" configuration, we can use existing commercial data centers; in our analysis, we select DRFortress.

Florida. Florida encompasses a large power topology of over 200 power plants and over 2500 substations [34]. We consider a set of six large power plants as candidate control center sites. Figure 6 shows these sites, as well as commercial data centers. In our analysis, we consider the first control center to be at Palm Beach (matching the current situation, where FPL invested in a Category-5-hurricane-resilient control center [35]). For primary-backup configurations, we locate the backup control center at Port Orange, where publicly available information shows that FPL has operated a backup control center [36]. For the "6+6+6" configuration, we additionally use the Jacksonville Tierpoint data center.

B. Hurricane Modeling

To generate "Natural Hazard Realizations" for the Compound Threat Analyzer, we model realistic hurricane impacts.

Oahu, Hawaii. We consider five hurricane paths used by emergency planners in Hawaii (dashed lines in Figure 5). For each path, we generated hurricane realizations using a wave-surge model based on a simulated hurricane using the ADvanced CIRCulation Model (ADCIRC). The ADCIRC model models the inundation as a hurricane makes landfall [37]. For each path, we modeled a Category 2 and a Category 3 strength hurricane, for a total of 10 hurricanes.

The ADCIRC model uses mesh discretization near the Oahu shoreline and then calculates the surge elevations over time as a result of a numerical model of the hurricane wind field. Because the mesh was coarse near the shoreline, we averaged the water surface elevations near the shoreline, and then extended the water surface elevation onto the shoreline to produce the inundation (following common practice). This provides what is typically felt to be a realistic representation of hurricane surge flooding from a storm [38].

The power assets shown in Figure 5 were tracked to determine the inundation levels at those sites in each hurricane realization. The peak inundation from the simulated hurricane was assumed to cause the asset to fail if inundation exceeded 0.5m (2ft), since this is the typical height for switches in power plants and substations. For each hurricane path and strength, we performed 1000 realizations.

Florida. Because detailed flood hazard data is not available for Florida (due to the large size of the state), we take a different approach than for Hawaii. For each site in Figure 6, we consider the 100-year flood hazard intensity at that site, i.e. the level of flooding that has a 1% chance of being reached or exceeded each year, based on the NOAA storm surge SLOSH Model [39], [40] and the FEMA national flood hazard layer [41], [42], along with local flood hazard maps, where available. These data sources provide an upper and lower bound for the flood depth at each site, which we use to develop a truncated normal distribution for each site (parameters are shown in Appendix B). To generate 1000 hurricane realizations, we perform Monte Carlo simulation, randomly generating the flood depth at each site based on its normal distribution parameters. As for Hawaii, we assume a site fails in a realization if its inundation exceeds 0.5m. We note that this represents a worst-case scenario, where the entire state is assumed to be affected by the same hurricane event. While this is not fully realistic, it represents a reasonable upper bound on the flooding impacts of any real Category 3 hurricane, and provides useful insights when combined with our Hawaii data. We refer to this as the Florida "Superstorm".

VI. Evaluation and Insights

A. Analyzing SCADA Architectures Under Compound Threats

We start by analyzing the resilience of different SCADA architectures under compound threats in Hawaii and Florida, using the system configurations in Section V-A. These configurations are designed to be as realistic as possible: in both cases,