Analysis of Load-Altering Attacks Against Power Grids: A Rare-Event Sampling Approach

Maldon Patrice Goodridge*, Subhash Lakshminarayana† and Christopher Few‡
*Global Development Initiatives, Queen Mary University of London
†School of Engineering, University of Warwick, UK
‡National Grid, UK
Emails: *m.p.goodridge@qmul.ac.uk, †subhash.lakshminarayana@warwick.ac.uk, ‡christopher.few@nationalgrid.com

***Abstract*—By manipulating tens of thousands of internet-of-things (IoT) enabled high-wattage electrical appliances (e.g., WiFi-controlled air-conditioners), large-scale load-altering attacks (LAAs) can cause severe disruptions to power grid operations. In this work, we present a *rare-event sampling* approach to identify LAAs that lead to critical network failure events (defined by the activation of a power grid emergency response (ER)). The proposed sampler is designed to 'skip' over LAA instances that are of little interest (i.e., those that do not trigger network failure), thus significantly reducing the computational complexity in identifying the impactful LAAs. We perform extensive simulations of LAAs using the Kundur two-area system (KTAS) power network while employing the rare-event sampler. The results help us identify the victim nodes from which the attacker can launch the most impactful attacks and provide insights into how the spatial distribution of LAAs triggers the activation of ERs.**

I. Introduction

Cyber attacks against power system grids can have significant social and economic consequences. The threats can be broadly divided into two categories – (i) attacks that directly target the power grid's supervisory control and data acquisition (SCADA) system, and (ii) attacks that indirectly target the power grid's control loops by manipulating end-user internet-of-things (IoT) enabled electrical appliances. Direct attacks against the SCADA system, such as false data injection attacks and/or coordinated cyber-physical attacks against power grid state estimation have received significant attention [1]–[3]. In contrast, indirect attacks that target a large number of demand-side appliances in a Botnet-type attack have been studied only recently [4], [5]. Unlike the SCADA assets, these devices cannot be monitored continuously due to their large numbers.

The focus of this work is on *load-altering attacks* (LAAs), which refer to a sudden and abrupt change in the power grid demand by synchronously turning on/off a large number of IoT-enabled high-wattage appliances [4]–[8]. LAAs pose a major threat to power grid operations since they can potentially disrupt the balance between supply and demand. It has been shown [4], [5] that such attacks can lead to unsafe frequency excursions, line outages, and/or increase the grid's operational costs. Moreover, dynamic LAAs, in which the attacker injects a series of load perturbations over time, can also destabilize the power grid's frequency control loop [7]. Subsequent work has also focused on detecting LAAs using data-driven approaches based on the data gathered from phasor measurement units [9], [10].

Understanding the impact of LAAs is an important component in risk analysis. Existing work on quantifying the impact of LAAs can be categorized into two approaches – (i) simulation-based approach [4]–[7], and (ii) analytical approach [8]. Under the former, the attack impact is assessed by simulating the power grid's control loops (e.g., frequency dynamics) [4]–[6] or computing the system's eigenvalues under LAAs [7]. However, the results presented in these works correspond to only a few specific LAA scenarios (i.e., specific load perturbations injected at the victim nodes). Evaluating the attack impact under all possible spatial distributions of LAAs (over the victim nodes) requires performing extensive simulations considering different combinations of the victim nodes and attack magnitudes, which can be computationally prohibitive. To overcome these issues, an analytical approach based on the theory of second-order dynamical systems was proposed in [8]. The closed-form analytical functions to evaluate the attack impact (in terms of the dynamic response and eigensolutions of the power grid's frequency control loop) only need to be computed once, thus avoiding the requirement for repeated simulations. However, the analytical approach is restricted to a second-order model with a direct current (DC) power flow model. Extending these results to higher-order models (e.g., one that considers both frequency and voltage dynamics) and involving the non-linear alternating current (AC) power flow model is non-trivial. This is important to obtain a realistic assessment of the LAA impact.

To overcome the aforementioned limitations, we apply a sampling approach to map LAA magnitudes and their spatial distributions (across the different victim nodes) to their attack impact. In particular, the focus of this work is on LAAs that lead to the activation of power system emergency responses (ERs) which disconnect critical power system components (e.g., generators/load/transmission lines). A common sampling approach is *Monte Carlo* simulations, which would apply randomly generated realisations of LAAs from an underlying distribution to a simulated model of the physical system. However, power grid design philosophies, such as $N-1$ scheduling, make the grid resilient to various contingencies (including cyber attacks). This implies component disconnections induced by LAAs can be extremely rare. Consequently, Monte Carlo sampling can be computationally expensive, as the majority of sampled LAAs will not result in component disconnections, requiring a large number of realisations applied to the power

system model to generate a sample of the rare event (i.e., the activation of an ER).

To avoid an exhaustive search to locate potential LAAs that lead to network failures, we instead employ a novel methodology in the context of LAAs based on a Markov chain Monte Carlo (MCMC) approach for rare-event sampling, known as the *skipping sampler*. The proposed approach is designed to reduce the time and computational effort spent on evaluating LAAs that are of little interest (i.e. those that do not result in an emergency response), allowing it to more efficiently construct a sample of a rare event. The skipping sampler has been applied to draw samples of low probability, high impact events in power networks in the literature (see [12] and [13]). Section III-B provides a detailed discussion on the skipping sampler.

We evaluate the framework by performing extensive simulations using the Kundur two-area system (KTAS) [11]. We simulate the power grid's transient dynamics by a third-order model, which accounts for both the frequency as well as the voltage dynamics [12]. The LAAs that perturb these dynamics are modelled according to the *Log-normal* distribution (since real-world cyber attack magnitudes are well modelled by this distribution [13]). If the local frequency metrics at any node exceed pre-set tolerances, appropriate ERs, such as generation/load shedding, are activated. Our results show that in the KTAS network, most instances of ERs are activated for two specific spatial distributions of LAAs, namely (i) when the attacker increases the load throughout the system or (ii) when the attacker decreases the load in the over-provisioned area (with excess generation) and increases the load in the under-provisioned area (with excess load). In particular, attacks that exacerbate the power imbalance in the system (i.e., type (ii)) can trigger inter-connector line disconnections, and lead to other network failures.

The rest of the paper is organised as follows. Section II introduces the system model; Section III presents the statistical model and details of the rare event sampling approach. Section IV describes the simulation results and Section V concludes. The simulation parameters are provided in an online Appendix found in [14].

II. System model

A. Power Grid Model

Using the third-order model for the generator, the power system model for rare-event sampling simulations will also include a model for governor action, automatic voltage regulation and a model of protection system ERs if the frequency or RoCoF exceeds pre-defined thresholds. We consider a power grid represented by $\mathcal{G} = \{\mathcal{N}, \mathcal{W}\}$, where $\mathcal{N}$ is the set of buses and $\mathcal{W}$ is the set of transmission lines. The set of buses $\mathcal{N}$ consists of N generation buses and L load buses, with $|\mathcal{N}| = N + L$. At each generation bus $i = 1, \dots N$, the dynamics for the phase angle δ_i, voltage magnitude E_i and governor action ρ_i are given respectively by:

$$\begin{cases} M(\psi)\ddot{\delta}_i + D\dot{\delta}_i = \psi_i \chi_i^G - \chi_i^L(\mathcal{R}_i) - \\ \qquad\qquad E_i \sum_{j=1}^{N+L} B_{ij}(\Omega_{ij}) E_j \sin(\delta_{ij}) & \text{(1a)} \\ S_i \dot{E}_i = \psi_i (E_{f,i} - \mathrm{v}_i) - E_i + X_i \sum_{j=1}^{N+L} B_{ij}(\Omega_{ij}) E_j \cos(\delta_{ij}) & \text{(1b)} \\ \dot{\rho}_i = -A_i \dot{\delta}_i (1 - 1_{\mathcal{W}}[\dot{\delta}_i]). & \text{(1c)} \end{cases}$$

In a similar manner, the dynamics for δ_i and E_i at each load bus $i = N+1, \dots, N+L$ are given by:

$$\begin{cases} M(\psi)\ddot{\delta}_i + D\dot{\delta}_i = -\chi_i^L(\mathcal{R}_i) - E_i \sum_{j=1}^{N+L} B_{ij}(\Omega_{ij}) E_j \sin(\delta_{ij}) & \text{(2a)} \\ S_i \dot{E}_i = \psi_i E_{f,i} - E_i + X_i \sum_{j=1}^{N+L} B_{ij}(\Omega_{ij}) E_j \cos(\delta_{ij}.) & \text{(2b)} \end{cases}$$

In equations (1) and (2), ψ_i, Ω_{ij} and $\mathcal{R}_i$ are indicator variables associated with generator, line and load disconnections respectively, explained in Section II-C. The system angular momentum, $M(\psi) = \sum_{j=1}^{N} \psi_j H_j$ is given by the sum of each generator's inertia constant, H_i (see online Appendix for parameter values).

TABLE I: Variables used in (1) and (2).

Symbol	Meaning	Units
A_i	Governor's droop response	MW/rad
$B_{ij}(\Omega_{ij})$	Susceptance matrix	p.u.
χ_i^G	Net generation at node i	p.u.
$\chi_i^L(\mathcal{R})$	Net loads at node i	p.u.
D	System damping	%
δ_i	Phase angle	p.u
δ_{ij}	$\delta_i - \delta_j$	p.u.
$\dot{\delta}_i$	Frequency	p.u
$\ddot{\delta}_i$	Rate of change of frequency (RoCoF)	p.u.
E_i	Voltage	p.u.
$E_{f,i}$	Machine i rotor field voltage	p.u.
$M(\psi)$	System angular momentum	Ws^2
Ω_{ij}	Line disconnection indicator	-
ψ_i	Generator shed indicator	-
$\mathcal{R}_i$	UFLS counter	-
S_i	Machine i transient time constant	s
X_i	Machine i equivalent reactance	ohms
$\mathcal{W}$	Governor's deadband frequency range	Hz

As we assume the network to be lossless, the elements of $B_{ij}(\Omega_{ij})$ correspond to the imaginary part of the elements of the network's admittance matrix [11]. The net generation at node i is given by $\chi_i^G = \min\{P_i^{\max}, P_i^G + \rho_i\}$, where $P_i^{\max}$ is the nominal maximum power output of generator i, P_i^G is the equilibrium power of the generator and ρ_i is the power contributed by a governor unit, whose dynamics are given in (1c). The variable v_i accounts for the action of automatic voltage regulation (see online Appendix). The net load at node i, χ_i^L, is inclusive of the LAA and a load disconnection scheme, and is discussed in Sections II-B and II-C. The remaining parameters are given in Table I.

B. Load-Altering Attack Model

Several security vulnerabilities have been identified in IoT-enabled high-wattage consumer appliances (see, e.g., [5]).

These vulnerabilities can be exploited by a strategic attacker to cause security incidents such as information disclosure and privilege escalation, leading to a change in the device's operational settings (e.g., switch ON/OFF or change the mode of operation). Considering a 2 kW power rating for the ACs, and a Botnet-scale attack that potentially compromises tens of thousands of such devices, LAAs can lead to a sudden load change of several MWs of power [5].

Not all loads are expected to be susceptible to LAAs, thus we decompose P_i^L, the equilibrium load at bus i, into a vulnerable part, given by νP_i^L, where $\nu \in [0,1]$ denotes the proportion of equilibrium loads in the network vulnerable to an LAA; and a secure part (i.e., protected or non-smart loads) $(1-\nu)P_i^L$. LAAs at node i, denoted u_i, are modelled as $u_i \coloneqq \eta_i \nu P_i^L$, where $\eta_i \in [-1,1]$ is the proportional change to equilibrium vulnerable load. Thus, the load at node i, inclusive of the LAA, is given by

$$\begin{aligned}\chi_i^L &= (1-\nu)P_i^L + \nu P_i^L + u_i \\ &= (1-\nu)P_i^L + (1+\eta_i)\nu P_i^L. \end{aligned} \tag{3}$$

This constrains the authority of the attacker to alter vulnerable loads at node i to a minimum of 0 MW (i.e. $\min(u_i) = -P_i^L\nu$), or, at maximum, double vulnerable load demand (i.e. $\max(u_i) = P_i^L\nu$). This restriction to the maximum LAA reflects the finite capacity of inactive loads the attacker can activate during an LAA.

C. Emergency Responses

ERs refer to systems designed to protect sensitive power system components from excessive frequency deviations following a change in the active power balance. In this section we provide a brief discussion of ER employed, and refer the reader to [15] and [16] for a detailed mathematical description.

Generation shedding: To protect synchronous generators, we model two independent schemes intended to disconnect the generator from the network: (i) RoCoF-induced generation shedding (RIGS) - the generator is disconnected when nodal RoCoF $|\ddot{\delta}_i|$ exceeds an upper threshold; (ii) over frequency generation shedding (OFGS) - generation is shed when nodal frequency $\dot{\delta}_i$ exceeds a pre-set upper limit. The binary variable ψ_i models the activation of generation shedding at node i during the simulation: $\psi_i = 1$ initially under normal operation; however, when either threshold is met and generator i is disconnected, ψ_i is set to 0 until the end of the simulation.

Under-frequency load shedding (UFLS): We model UFLS schemes as a progressive disconnection of loads when the frequency $\dot{\delta}_i$ falls below a strictly decreasing sequence of four frequency thresholds $F^U \coloneqq \{F_1^U, \ldots, F_4^U\}$ where $F_{j-1}^U > F_j^U$. In our model, at each frequency threshold, 10% of equilibrium loads P_i^L is automatically disconnected to arrest the decline in nodal frequency. Letting $\mathcal{R}_i \in \{0,1,2,3,4\}$ count the total number of UFLS activations at node i at each time step t in the simulation, the *net load* is a dynamic variable in the power system model:

$$\chi_i^L(\mathcal{R}_i) = \Big(1 - 0.1\mathcal{R}_i\Big)\Big((1-\nu)P_i^L + (1+\eta_i)\nu P_i^L\Big) \tag{4}$$

Line disconnection: In the KTAS network, we model the disconnection of the line connecting Areas 1 and 2 when the power flow through the line, given by $\phi_{ij} \coloneqq B_{ij}E_iE_jsin(\delta_i - \delta_j)$, exceeds a pre-set power threshold P^ϕ. When excess power flow is detected through the inter-connector line, the indicator Ω_{ij} switches from 1 to 0 for the remainder of the simulation, setting the ij^{th} element of B_{ij} to 0 [16].

The ER model inspects the continuous time variables $\dot{\delta}_i(t)$, $\ddot{\delta}_i(t)$ and $\phi_{ij}(t)$ from the power system model (1) at regular time intervals. Once a criteria for activation is observed, the corresponding ER is activated. This is represented in (1) as a discontinuity, where changes to the relevant input variables (power injection, load or network topology) are applied. Subsequently, the simulation is resumed with the new network parameters.

III. Statistical Model for LAAs

In this section, we present the statistical model for the distribution of LAAs and describe the proposed rare-event sampling approach to identify the impactful LAAs.

A. Modeling the Unconditional Distribution of LAAs

There are several studies that document the frequency and magnitude of cyber breaches in enterprise networks [13]. For instance, [13] demonstrates the size of data breaches can be well-modelled by the log-normal family of distributions. Assuming nodal LAAs magnitudes are independent and follow a similar distribution, we model $U \in \mathbb{R}^{N+L} \coloneqq [|u_1|, \ldots, |u_{N+L}|]$ as

$$U \sim \prod_{i=1}^{N+L} \text{Lognormal}(\mu_i, \sigma_i^2). \tag{5}$$

We also note that our analysis is not restricted to the log-normal family of distributions, and can be extended to any underlying distribution in a straightforward manner.

B. Rare-Event Sampler for LAAs

If given an unconditional density ρ over $\mathbb{R}^{N+L}$ and a rare event of interest $C \subset \mathbb{R}^{N+L}$, *rare-event sampling* involves drawing elements from ρ conditioned on the occurrence C. The density of this conditional distribution for the element $U \in \mathbb{R}^{N+L}$ is:

$$\pi(U) = \frac{\rho(U)\mathbb{1}_C(U)}{\rho(C)}, \tag{6}$$

where $\rho(C)$ is the probability of the event C occurring, and

$$\mathbb{1}_C(U) = \begin{cases} 1 & U \in C \\ 0 & U \notin C. \end{cases} \tag{7}$$

In the context of our research, C is the set of LAA magnitudes which result in the activation of at least one ER and ρ is the distribution in (5). As C is expected to be rare, we employ the *skipping sampler* MCMC algorithm to efficiently draw samples of $U \in C$. The skipping sampler is formalised in Algorithm 1. We provide an intuitive explanation of the algorithm in the following.

Algorithm 1: Skipping sampler algorithm

Input: initial state U_1;
for i = 1 to n do:
Generate an initial proposal Z_1 distributed according to the density $q(y - U_i)dy$;
Calculate the direction $\Phi = (Z_1 - U_i) / \|Z_1 - U_i\|$;
Generate a halting index $K \sim K_\varphi$;
Set $k = 1$ and:
while $Z_k \notin C$ **and** $k < K$ **do**
  Generate a distance increment R distributed according to $q_{r|\Phi}(r|\Phi)$;
  Set $Z_{k+1} = Z_k + \Phi R$;
  k=k+1;
end
Set $Z := Z_k$;
Evaluate the acceptance probability:

$$\alpha(U_i, Z) = \begin{cases} \min\left(1, \frac{\pi(Z)}{\pi(U_i)}\right) & \text{if } \pi(U_i) \neq 0, \\ 1, & \text{otherwise,} \end{cases} \quad (8)$$

Generate a uniform random variable V on $(0, 1)$;
if $V \leq \alpha(U_i, Z)$ **then**
  $U_{i+1} = Z$;
else
  $U_{i+1} = U_i$;
end
return U_{i+1}.
Output: final sample $[U_1, \ldots, U_h]$

As a Metropolis-class algorithm [17], the skipping sampler can be understood as a two step procedure: (1) a *proposal step* where, starting from a state $U_n \in \mathbb{R}^{N+L}$, a potential new state Z for the final sample is generated; (2) an *acceptance/rejection* step, which determines whether the proposed state Z is included in the final sample, according to a specified acceptance probability. This ensures the distribution of the sample follows the desired target distribution π. If it is accepted, the proposal is included in the final sample and becomes the starting state for the next proposal step. This procedure is repeated a desired number of times, after which the final sample is returned.

The skipping sampler improves the sampling of C by using a specialised proposal function which 'skips' over C^c - the set of LAAs which do not lead to the activation of an ER (which are not of interest); until the rare event is sampled or the skipping process is terminated. Thus, the proposal of the skipping sampler was designed to enable efficient transitions between connected components of C. Denoting the current state U_n, if the initial proposal $Z_1 \notin C$, we update the initial proposal (or 'skip') by an adding an independent random distance increment R_2 in the direction $\Phi = \frac{Z_1 - U_n}{|Z_1 - U_n|}$, where R_k has the conditional distribution of $\|Z_1 - U_n\|$ conditioned on Φ. The proposal function continues this linear update procedure until either C is entered or the budget for skipping is exhausted [17].

IV. Simulations

A. Simulation Settings

Our case study is based on the KTAS power grid [11]. We take a Kron reduced version consisting of $N = 4$ generation buses and $L = 2$ load buses as shown in Figure 1. At $t = 0^-$, the system is modelled in equilibrium, with power flows from Area 1 to Area 2 through the line connecting nodes 5 and 6 (tie line). The system parameters are such that the system is $N - 1$ secure, in the sense that the loss of a generator (in the absence of any other disturbance) does not trigger an ER. The initial conditions of the above system of equations, denoted $\delta_i(0)$, $E_i(0)$, $\rho_i(0)$, P_i^G and P_i^L are set equal to equilibrium states which can be determined numerically, such that $\ddot{\delta}_i \approx 0$. These values along the those of the parameters of (1) and (2) can be found in the online Appendix [14].

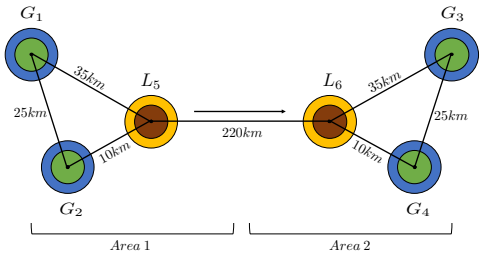

Fig. 1: Schematic drawing of the Kundur two-area 4 node network after Kron reduction. Generator buses (green circles) correspond to nodes $i = [1, \ldots, 4]$ and load buses (brown circles) correspond to nodes $i = 5, 6$. Line lengths are indicated.

To generate LAA instances, we implement the skipping sampler proposed in Section III-B. During each proposal step, we sample an $N + L-$dimension LAA vector U from Lognormal distribution as in (5). We investigated various values for $\sigma_i \in [1, 8]$ which controls the rareness of large LAAs. For this study, we present the results for $\mu_i = 0$ and $\sigma_i = 4$ for $i = 1, \ldots, N + L$, and reserve a detailed sensitivity analysis for a dedicated manuscript. We apply U as an input to the power system model (1) at $t = 0$, with frequency dynamics simulated for 15 seconds following the LAA using MATLAB. We conduct $n = 60,000$ proposals, which generated a final sample of $h \approx 10,500$ LAAs conditioned on the activation of at least one ER. This is an acceptance rate of 17.5%, within the 15 - 48% rate considered optimal for exploring of a sample [18]. Following [17], no burn-in nor thinning was required, thus all samples collected were available for analysis. Since these responses are an undesirable event (from a system operator's point of view), we label such instances as a "network failure".

B. Simulation Results

We evaluate the susceptibility of the network to LAAs at both a local and global scope for three regimes of network vulnerability to LAAs- a 'secure network' ($\nu \leq 45\%$), a 'moderately vulnerable network' ($45\% < \nu \leq 65\%$) and a 'highly vulnerable network' ($\nu > 65\%$).

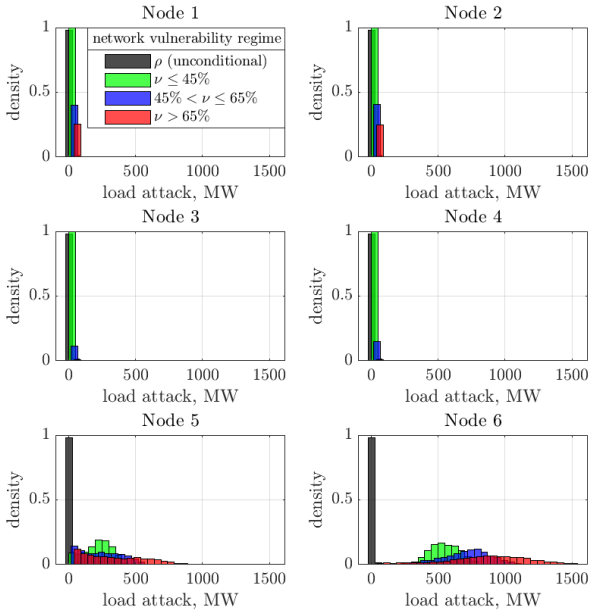

Fig. 2: Distribution of the absolute value LAAs at each node of the KTAS network conditioned on a network failure event, for different levels of network vulnerability ν and the unconditional distribution (ρ). Note - bins for successive values of ν are offset slightly to the right to improve readability.

Local Analysis – Identifying the vulnerable nodes: Figure 2 plots the distribution of LAA magnitudes at each node, conditioned on the occurrence of a network failure, for different degrees of network vulnerability ν. We observe the following : (i) only at nodes 5 and 6 does the conditional distribution of LAA magnitudes differ significantly from ρ, with network failures associated with larger magnitude LAAs, located in the low density region of ρ. This is driven by the design of the KTAS network, where most loads are concentrated at nodes 5 and 6, giving the attacker sufficient leverage over system frequency to trigger a network failure. Thus, network failures are primarily driven by LAAs at these nodes, mostly independent of the LAA magnitude at nodes $1-4$; (ii) as network vulnerability ν increases, the distribution of LAA magnitudes at node 6 shifts rightwards, implying larger magnitude LAAs at node 6 become more prevalent in the sample; however (iii) we note that large magnitude LAAs are rare under their assumed Log-normal distribution. Thus, Fig. 2 reveals network failures, driven by large LAA magnitudes at node 6, are indeed low probability events. Summarily, these results suggest nodes 5 and 6 are the critical nodes in the KTAS network from which the attacker can launch the most impactful attacks, regardless of the degree of network vulnerability.

Figure 3 illustrates the average number of activations of each ER per sample for different levels of network vulnerability, defined as

$$\epsilon_f = \frac{\text{Total ER}_f \text{ activations in the sample}}{\text{Sample size}}. \tag{9}$$

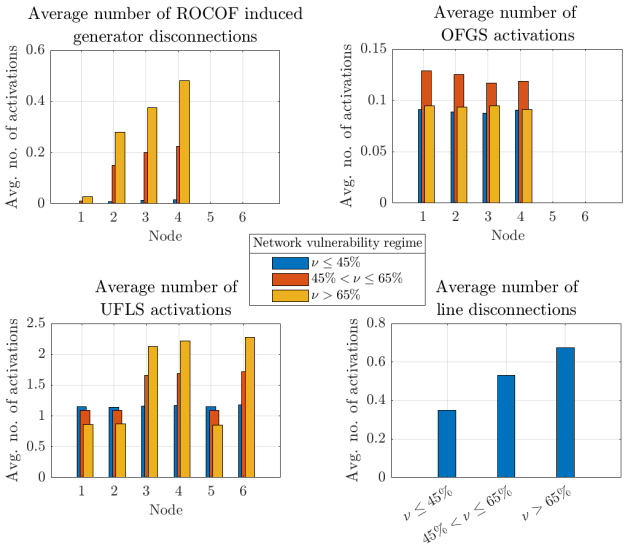

Fig. 3: Average number of each ER per sample $\epsilon_i, i = 1, \ldots, 4$ (defined in (9)) in the KTAS network for different degrees of network vulnerability to LAAs.

Herein, the indices $f = 1, \ldots, 4$ correspond to RoCoF-induced generation shedding (RIGS), over-frequency generator shedding (OFGS), UFLS activations, and inter-connector line trips respectively. Note that $\epsilon_f \in [0, 1]$ for $f = 1, 2, 4$. However, $\epsilon_3 \in [0, 4]$, as each node can experience a maximum of 4 UFLS events ($f = 3$) during the simulation (see Section II-C).

For secure networks, i.e, when $\nu \leq 45\%$, the average number of RIGS is comparatively small, as low-magnitude LAAs are unable to induce sufficiently large RoCoF deviations to trigger generator disconnection. Additionally, we observe all generator nodes experience OFGS and UFLS events at similar rates of $\epsilon_2 \approx 0.1$ and $\epsilon_3 \approx 1$ respectively. Together, these imply for a secure KTAS network, each class of ER is similarly likely at any node in the network. Thus, to counter the potential threat of LAAs in a secure KTAS network will require a system-wide solution, e.g. - a coordinated automatic generation control (AGC) system.

As ν increases, it grants the attacker greater authority over network loads and the ability to significantly disrupt the active power balance of the network. This is associated with increased rates of RIGS, UFLS and line disconnections, as these are triggered by large power deviations. However, when $\nu > 65\%$, the average number of OFGS responses declines, as large changes in loads result in RIGS dominating generation shedding events. In contrast to secure networks, where nodes are similarly vulnerable to UFLS events under LAAs, Figure 3 reveals a differential in the nodal risk of an UFLS event in moderately and highly vulnerable networks. For example, when $\nu > 65\%$, nodes $3, 4$ and 6 experience approximately twice as many UFLS events as nodes $1, 2$ and 5. This can be best understood through a global analysis of the KTAS network, which is discussed next.

Global Analysis - Spatial Distribution of LAAs: Recall that the KTAS network is comprised of two areas- Area 1 (nodes 1, 2 and 5) with excess generation, and Area 2 (nodes

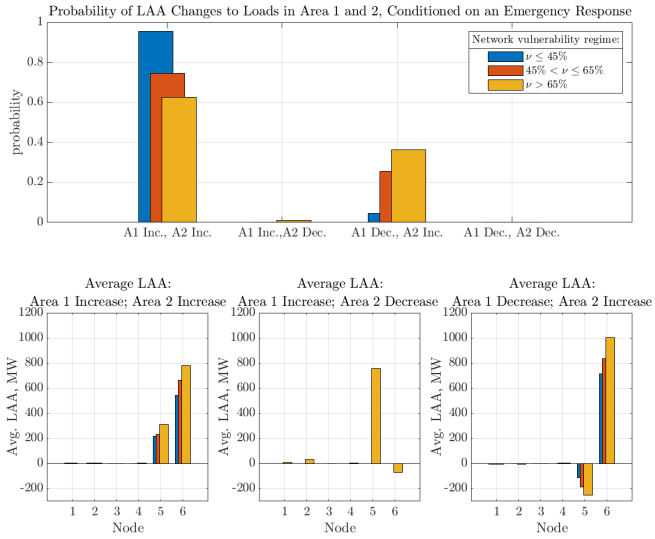

Fig. 4: Probability of network failure for increase/decrease of system load in each area for the KTAS network.

3, 4 and 6) with excess demand. With sufficient authority over network loads, an attacker can exploit the global pre-LAA power imbalances of each area to trigger an ER by decreasing loads in Area 1 to exacerbate the excess generation, and increasing loads in Area 2 to exacerbate the generation deficit. Thus, as ν increases, we observe fewer UFLS responses in Area 1, more UFLS responses in Area 2, and a substantial increase in power transferred across the inter-connector, increasing the rates of line disconnections (Figure 3).

The global susceptibility of the KTAS network to LAAs is also illustrated by Figure 4, which plots the probability of the relative changes in loads in each Area of the KTAS network, conditioned on the occurrence of a network failure. First, we note the KTAS network is secure against global reduction of loads in both Areas, as it is likely governors are able to adjust generator output to handle the loss of loads in all vulnerability regimes. Instead, we observe that failure events occur primarily for large increases in loads, dominated by two scenarios: (i) when the LAA increases the demand in both Areas- this scenario is the primary driver of failure events when KTAS is secure, as the attacker increases loads at both nodes 5 and 6 beyond the maximum power capabilities of generators, triggering a failure anywhere in the network; (ii) the attacker decreases loads in Area 1 and increases them in Area 2- this scenario is rare when the KTAS is secure, as the attacker generally lacks sufficient authority to exploit the imbalances of the KTAS network. However, as ν (and the attacker's authority) increases, so too does the conditional probability of scenario (ii) (see Figure 3).

V. Conclusions and Future Research

In this work, we present a framework to evaluate the impact of LAAs on power grid operations using a rare-event sampling approach. The proposed approach provides a comprehensive framework to examine the impact of LAAs under different potential spatial distributions of LAAs across the power network. Our results identify the nodes from which the attacker can launch the most impactful LAA and further illustrate how the attacker can exploit the inter-area power imbalances in the network to trigger ER events. Future work includes (i) an extension of the analysis to dynamic LAAs against power grids [7], (ii) considering correlations in LAA injections, (iii) showing scalability of the proposed method to large-dimensional systems.

References

[1] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," in *Proc. ACM Conference on Computer and Communications Security (CCS)*, 2009, pp. 21–32.

[2] S. Lakshminarayana, E. V. Belmega, and H. V. Poor, "Moving-target defense against cyber-physical attacks in power grids via game theory," *IEEE Transactions on Smart Grid*, vol. 12, no. 6, pp. 5244–5257, 2021.

[3] S. Lakshminarayana, A. Kammoun, M. Debbah, and H. V. Poor, "Data-driven false data injection attacks against power grids: A random matrix approach," *IEEE Transactions on Smart Grid*, vol. 12, no. 1, pp. 635–646, 2021.

[4] A. Dabrowski, J. Ullrich, and E. R. Weippl, "Grid shock: Coordinated load-changing attacks on power grids: The non-smart power grid is vulnerable to cyber attacks as well," in *Proc. ACSAC*, 2017, pp. 303–314.

[5] S. Soltan, P. Mittal, and H. V. Poor, "BlackIoT: IoT botnet of high wattage devices can disrupt the power grid," in *Proc. USENIX Security Symposium*, Baltimore, MD, Aug. 2018, pp. 15–32.

[6] B. Huang, A. A. Cardenas, and R. Baldick, "Not everything is dark and gloomy: Power grid protections against iot demand attacks," in *Proc. USENIX Security Symposium*, Aug. 2019, pp. 1115–1132.

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

[8] S. Lakshminarayana, S. Adhikari, and C. Maple, "Analysis of IoT-based load altering attacks against power grids using the theory of second-order dynamical systems," *IEEE Transactions on Smart Grid*, vol. 12, no. 5, pp. 4415–4425, 2021.

[9] S. Amini, F. Pasqualetti, M. Abbaszadeh, and H. Mohsenian-Rad, "Hierarchical location identification of destabilizing faults and attacks in power systems: A frequency-domain approach," *IEEE Transactions on Smart Grid*, vol. 10, no. 2, pp. 2036–2045, 2019.

[10] S. Lakshminarayana, S. Sthapit, H. Jahangir, C. Maple, and H. V. Poor, "Data-driven detection and identification of iot-enabled load-altering attacks in power grids," *IET Smart Grid*, pp. 1–16, 2022.

[11] P. Kundur and N. Balu, *Power System Stability and Control*. McGraw-Hill, 1994.

[12] K. Schmietendorf, J. Peinke, R. Friedrich, and O. Kamps, "Self-organized synchronization and voltage stability in networks of synchronous machines," *The European Physical Journal Special Topics*, vol. 223, no. 12, pp. 2577–2592, 2014.

[13] B. Edwards, S. Hofmeyr, and S. Forrest, "Hype and heavy tails: A closer look at data breaches," *Journal of Cybersecurity*, vol. 2, no. 1, pp. 3–14, 12 2016.

[14] "Online appendix," https://tinyurl.com/ycknrnzv.

[15] M. Patrice Goodridge, J. Moriarty, and A. Pizzoferrato, "Distributions of cascade sizes in power system emergency response," pp. 1–6, 2020.

[16] ——, "A rare-event study of frequency regulation and contingency services from grid-scale batteries," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 379, no. 2202, p. 20190433, 2021.

[17] J. Moriarty, J. Vogrinc, and A. Zocca, "The Skipping Sampler: A new approach to sample from complex conditional densities," pp. 1–20, 2019. [Online]. Available: http://arxiv.org/abs/1905.09964

[18] D. Gamerman and H. Lopes, *Markov Chain monte Carlo, Stochastic Simulatoin for Bayesian Inference*, 2nd ed. Boca Raton: Chapman & Hall/CRC, 2006.