

Analysis of Cascading Failures Due to Dynamic Load-Altering Attacks

Maldon Patrice Goodridge*, Alessandro Zocca†, Subhash Lakshminarayana‡

*Global Development Initiatives

†Department of Mathematics, Vrije Universiteit Amsterdam, NL

‡School of Engineering, University of Warwick, UK

Emails: *gdi.uklimited@gmail.com, †a.zocca@vu.nl, ‡subhash.lakshminarayana@warwick.ac.uk

Abstract—Large-scale load-altering attacks (LAAs) are known to severely disrupt power grid operations by manipulating several internet-of-things (IoT)-enabled load devices. In this work, we analyze power grid cascading failures induced by such attacks. The inherent security features in power grids such as the $N - 1$ design philosophy dictate LAAs that can trigger cascading failures are *rare* events. We overcome the challenge of efficiently sampling critical LAAs scenarios for a wide range of attack parameters by using the so-called “skipping sampler” algorithm. We conduct extensive simulations using a three-area IEEE-39 bus system and provide several novel insights into the composition of cascades due to LAAs. Our results highlight the particular risks to modern power systems posed by strategically designed coordinated LAAs that exploit their structural and real-time operating characteristics.

Index Terms—Load-altering attacks, cascading failures, rare-event sampling, $N-1$ secure

I. INTRODUCTION

Load-altering attacks (LAAs) can cause sudden changes in power grid demand by manipulating large numbers of internet-of-things (IoT) enabled, high-wattage electrical appliances (smart heat pumps, electric vehicle charging stations), posing serious risks to system safety [1]. Unlike traditional attacks, such as false data injection attacks (FDI) against the power grid supervisory control and data acquisition (SCADA) systems, LAAs are easier to execute from an attacker’s perspective, targeting less secure end-user sites that typically lack sophisticated defence measures. Additionally, LAAs directly alter system load (which is one of the input signals of the power grid’s frequency control loop), setting them apart from load-redistribution attacks [2], which are a type of FDI that targets SCADA-reported load measurements. The increasing penetration of renewable energy resources results in low-inertia conditions that can exacerbate the impact of LAAs [3].

LAAs can be broadly divided into two categories: static (S-LAAs) and dynamic (D-LAAs) [4]. S-LAAs involve a sudden, one-time change in loads [1], whereas, D-LAAs refer to a series of load changes over a period of time [5], [6]. S-LAAs can result in network frequency and/or line flows exceeding safety limits, leading to component disconnections [1], [7]. In contrast, under D-LAAs, if an attacker changes system loads proportionally to the local frequency deviations [5] or proportional to the inter-area oscillations [6], they can potentially destabilize the frequency control loop, leading to cascading failures. An analytical approach to studying these

effects was proposed in [4] using second-order dynamical systems to identify high-impact attack nodes.

Despite the growing literature on LAAs, existing research lacks a framework to understand the extent of consequences, specifically, in terms of the cascading disconnections such large-scale attacks can cause. The first work in this direction is [1], but the analysis therein did not consider power grid protection features such as $N - 1$ security and load shedding and, as a result, the extent of cascades has been significantly overestimated. LAA-induced cascading failures considering the aforementioned security features were analyzed in [7]. However, the results presented in both [1] and [7] correspond to only a few specific LAA scenarios (i.e., specific load perturbations injected at a few of the victim nodes). They do not provide a thorough understanding of the distribution of cascades due to all possible spatial LAA scenarios over the victim nodes.

However, identifying LAAs that lead to cascades is challenging, since $N - 1$ design philosophies ensure that the power grid is resilient to the destabilizing effects of such load changes [7]. Thus, LAA instances that lead to cascading failures are in fact *rare events* and sampling them efficiently for a wide range of attack parameters becomes highly nontrivial. Despite their low likelihood, rare events are key to understanding power systems reliability, see e.g. [8]. In prior work, simulating cascading failures often necessitates computationally intensive models (e.g. [9]–[11]), making them ill-suited for rare event identification. We address this differently by coupling a computationally efficient dynamic model with a sampling technique to identify network and LAA parameters associated with cascading failures, extending the approach from [12]. Specifically, we employ the *skipping sampler*, a Markov Chain Monte Carlo (MCMC) algorithm specifically designed to efficiently sample low-likelihood events. The core idea behind this method is to traverse “uninteresting” regions of the parameter space (i.e. “skip”) in a systematic way to more efficiently sample rare but critical D-LAA characteristics. The skipping sampler has already been successfully used in the context of power systems security [13] and reliability [14].

We significantly generalize the framework proposed in [13] and extend it to analyse LAA-induced cascades. Specifically, as opposed to the exclusive focus on S-LAAs considered in [13], we instead present a framework that captures both static and dynamic LAAs in a unified manner. We model attacks that inject a sequence of periodic load alterations proportional to the frequency deviation (see [5]). By varying the interval

between attacks, we can move from a static attack model with long intervals between attacks to a dynamic one with frequent attacks. Lastly, we consider the sequence of power grid emergency responses, and how they cascade over the simulation period and investigate their distributions as a function of various attack parameters. We perform extensive simulations using a three-area IEEE-39 bus system and provide the following novel insights into the composition of cascades induced by LAAs. Specific contributions include:

- Analysis of power grid cascades as a function of the attack parameters, namely (i) the amount of vulnerable load accessible by the attacker, and (ii) the interval between successive load attacks (for DLAA).
- Analysis of cascades due to LAAs under different loading conditions and inter-area power balance.
- Identification of dominant failure modes under different attack and power grid operational regimes.

Our results show that there are two attack regimes in which the grid is particularly vulnerable: (i) low-magnitude attacks with a short interval between the load changes (D-LAAs), and (ii) very large-magnitude attacks with a long interval between the load changes (S-LAAs). Furthermore, the network is highly vulnerable to D-LAAs in peak demand periods and areas with large power imbalances are particularly vulnerable to D-LAAs.

The rest of the paper is organized as follows. Section II introduces the system model; Section III presents the rare-event sampling algorithm; Section IV describes the simulation results and Section V concludes.

II. SYSTEM MODEL

The power system is represented as a graph $\mathcal{S} = \{\mathcal{N}, \mathcal{W}\}$, with \mathcal{N} is the set of buses and \mathcal{W} is the set of transmission lines. We decompose $\mathcal{N} = \mathcal{G} \cup \mathcal{L}$, where \mathcal{G} is the set of generator buses and \mathcal{L} is the set of load buses. We employ a third-order model to capture the grid dynamics, accounting for frequency and voltage transients, generator governor action, and automatic voltage regulation. Under this model, for each generator $i \in \{1, \dots, N\}$, dynamics in voltage phase angle δ_i , voltage magnitude E_i and governor action ρ_i are given by:

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

Similarly, the dynamics for δ_i and E_i at each load bus $i \in \{N+1, \dots, N+L\}$ are given by:

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

In equations (1) and (2), ψ_i , Ω_{ij} and \mathcal{R}_i are indicator variables for frequency protection models and reflect the disconnection of network components, i.e., generators, lines and loads respectively (see Section II-B). Net power injection at nodes $i \in 1, \dots, N$ is given by $\chi_i^G = \min\{P_i^{\max}, P_i^G + \rho_i\}$, where P_i^{\max} is the stated maximum power output of generator i , P_i^G is the power of the generator in equilibrium and ρ_i is the power contributed by a governor unit (1c) [13]. The variable v_i represents the actions of automatic voltage regulation (see online Appendix). The net load at node i , χ_i^L , includes equilibrium loads, the dynamic LAA and a load disconnection scheme, and is discussed in Sections II-A and II-B. The remaining parameters are given in Table I.

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

A. Dynamic load-altering attack model

We denote the maximum load at each node $i \in \mathcal{L}$ by $P_{i,\max}^L$. At the start of the simulation (effectively modelling different times of the day, see Section IV), let us denote the load at node $i \in \mathcal{L}$ by $P_i^L \in [0, P_{i,\max}^L]$. A D-LAA can be modelled as a time-dependent sequence of load changes at the vulnerable load nodes. More specifically, we denote by $\lambda_i(t_k) \in \mathbb{R}$ the magnitude in MW of the LAA at node i occurring at time epoch t_k for $k \in \mathbb{N}$. We assume that the time epochs are separated by a pre-determined, constant time interval \mathcal{I} . Each load change is applied as an instantaneous impulse and kept constant until the next update in the sequence. The initial load change, $\lambda_i(t_0)$, is selected randomly by the sampling procedure (see Section III). Subsequent load changes $\lambda_i(t_k)$ for $k \geq 1$ are deterministic, calculated using a frequency-dependent 'reverse governor' model. This describes an attacker with access to network data, intent on using the sequence of attacks to exacerbate frequency deviations (the premise of D-LAAs [5]). In mathematical terms, for every vulnerable load node i and for every index $k \geq 1$

$$\lambda_i(t_k) = \begin{cases} C\dot{\delta}_i(t_k) & \chi_i^{L-} - C\dot{\delta}_i(t_k) \in [0, P_{i,\max}^L], \\ P_{i,\max}^L - \chi_i^{L-} & \chi_i^{L-} - C\dot{\delta}_i(t_k) > P_{i,\max}^L, \\ -\chi_i^{L-} & \chi_i^{L-} - C\dot{\delta}_i(t_k) < 0, \end{cases} \quad (3)$$

where $\lambda_i(t_k)$ is the k^{th} D-LAA component at time $t_k = k\mathcal{I}$ for $k \in \mathbb{N}$, $\delta_i(t_k)$ is the frequency at node i at time $t = t_k$, $C \in \mathbb{R}^+$ is a network variable which relates the change in load to the frequency deviation at $t = t_k$. The expression $\chi_i^{L-} = P_i^L + \sum_{j=0}^{k-1} \lambda_i(t_j)$ refers to the total net loads at node i just before the application of the k^{th} load change, inclusive of all previous load changes and load shedding.

In essence, (3) ensures that the load at node $i \in \mathcal{L}$ following the D-LAA remains within the set $[0, P_{i,\text{max}}^L]$. The framework presented above models S- and D-LAAs in a unified manner, specifically by varying t_k . Note that $t_k \rightarrow 0$ represents a continuous load attack (i.e., the D-LAA framework presented in [5]), while $t_k \rightarrow \infty$ models an S-LAA (i.e., only a single load attack over the entire simulation interval). The D-LAA at each node can be aggregated in a straightforward manner, giving rise to two key metrics:

- 1) The *cumulative D-LAA* (in MW), $\Sigma_i(\lambda) := \sum_k |\lambda_i(t_k)| \in \mathbb{R}^L$, which measures the size of the attack at node i as the sum of the magnitudes of D-LAAs at that node over the duration of the simulation.
- 2) The *average network load change* (in MW); $\mu_i(\lambda, h) := \sum_{t=1}^h (\sum_i^{N+L} |\lambda_i(t_k)|) / h \in \mathbb{R}$, which is a measure of the average size of each load change across the network. It provides information about the average magnitude of network loads the attacker must manipulate to trigger the cascade observed.

B. Network emergency responses and cascading failures

Emergency responses (ER) are the systems safety mechanisms that safeguard sensitive network equipment from dangerous deviations in frequency-related profiles. We provide a list of ER employed in the following.

- 1) *Generation shedding*: we model two independent schemes which disconnect generators from the network: (i) RoCoF-induced generation shedding (RIGS) – generation is disconnected when nodal RoCoF $|\dot{\delta}_i|$ 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.
- 2) *Under-frequency load shedding (UFLS)*: this scheme disconnects of 10% of equilibrium nodal loads when the frequency $\dot{\delta}_i$ falls below a strictly decreasing sequence of four frequency thresholds.
- 3) *Line disconnection*: If the power flowing through an interconnector line linking different areas in the power grid exceeds a pre-set upper threshold, then the line gets disconnected.

While ERs are intended to arrest large deviations in frequency-related dynamics when coupled with the effects of D-LAAs, multiple disconnection events may transpire on the network even when such power grids are designed to be $N - 1$ secure [15]. Such *cascading failures* threaten network integrity and result in significant costs to various network stakeholders. We measure the *cascade size* \mathcal{X} resulting from a D-LAA as the cumulative power (in MW) of network components (loads

and generators) disconnected during the power grid operation following the LAA.

III. SAMPLING METHODOLOGY FOR D-LAAs

In this study, we use a sampling technique to create diverse spatiotemporal LAAs, which serve as inputs to the power system model outlined in Section II and assess the resulting cascading failures. The detailed models are presented next.

- 1) We model the **D-LAA magnitudes** at time t_0 at all L vulnerable nodes as independent, uniformly distributed random variables, namely $\boldsymbol{\lambda} = (\lambda_1(t_0), \dots, \lambda_L(t_0)) \sim \mathcal{U}[0, \lambda_{\text{max}}^0]^L$, where λ_{max}^0 denotes the attack limits. The D-LAA magnitudes at subsequent epochs are then uniquely determined by the dynamics described in Section II-A. The wide support for the uniform distribution allows us to capture also extreme scenarios in which an attacker manipulates a large proportion of the load.
- 2) We model the **interval between load changes** using a uniform discrete distribution, i.e. $\mathcal{I} \sim \mathcal{U}[1, T_{\text{max}}]$, where T_{max} (in seconds) denotes the duration of the dynamic simulations. Note that we do not allow subsequent attacks to be less than 1 second apart to account for the time an attacker may need to estimate the system frequency to calibrate their next move.
- 3) We sample the **scenario** τ uniformly among a set of four, each representative of the intra-day power equilibrium at a different moment of the day (see Section IV-A).
- 4) We model the **D-LAA-frequency response** as a uniform random variable $C \sim \mathcal{U}[C_{\text{min}}, C_{\text{max}}]$, where C_{min} and C_{max} denote practical limits for the D-LAA.

We note our sampling procedure is not constrained to the uniform distribution and, in fact, it can be easily extended to accommodate any underlying distribution.

The sampling assumptions for the various D-LAA features outlined in the previous subsection result in an unconditional product density ρ over \mathbb{R}^{L+3} , where each vector of attack parameters $\boldsymbol{x} = (\boldsymbol{\lambda}, \mathcal{I}, \tau, C)$ fully characterizes a D-LAA attack in view of the dynamics we described in Section II-A. Since we are interested in studying the cascading failures triggered by D-LAAs, we need to sample attack parameters $\boldsymbol{x} = (\boldsymbol{\lambda}, \mathcal{I}, \tau, C)$ that results in the activation of at least one network ER. Let $A \subset \mathbb{R}^{L+3}$ be the subset of such “successful” attack parameters. We are interested in sampling D-LAA attacks according to the density ρ but *conditionally* on the event A . Let π_A be such conditional distribution over \mathbb{R}^{L+3} , i.e., $\pi_A(\boldsymbol{x}) := \rho(\boldsymbol{x}) \mathbb{1}_A(\boldsymbol{x}) / \rho(A)$, where $\rho(A)$ is the probability of the event A occurring, and $\mathbb{1}_A(\boldsymbol{x}) = 1$ if $\boldsymbol{x} \in A$ and 0 otherwise. Due to the pre-existing network security mechanisms, however, it is very unlikely that a D-LAA would trigger a network ER, which means that A is a rare event with $\rho(A) \ll 1$. Any standard MCMC sampling method would then struggle enormously at sampling from the conditional density π_A .

We thus employ the *skipping sampler* MCMC algorithm proposed in [16] to efficiently draw samples from π_A . The steps followed are presented in Algorithm 1. As a Metropolis-class algorithm, the skipping sampler can be understood as a

Algorithm 1: Skipping sampler algorithm

```

1 Input: initial state  $U_1$ ;
2 for  $i = 1, \dots, n$  do
3   Generate an initial proposal  $Z_1$  distributed
   according to the density  $q(y - U_i)dy$ ;
4   Calculate the direction  $\Phi = (Z_1 - U_i) / \|Z_1 - U_i\|$ ;
5   Generate a halting index  $K \sim K_\varphi$ ;
6   Set  $k = 1$ ;
7   while  $Z_k \notin C$  and  $k < K$  do
8     Generate a distance increment  $R$  distributed
     according to  $q_{r|\Phi}(r|\Phi)$ ;
9     Set  $Z_{k+1} = Z_k + \Phi R$ ;
10     $k = k + 1$ ;
11  end
12  Set  $Z := Z_k$ ;
13  Evaluate the acceptance probability:
      
$$\alpha(U_i, Z) = \begin{cases} \min\left(1, \frac{\pi(Z)}{\pi(U_i)}\right) & \text{if } \pi(U_i) \neq 0, \\ 1, & \text{otherwise.} \end{cases} \quad (4)$$

14  Generate a uniform random variable  $V$  on  $(0, 1)$ ;
15  if  $V \leq \alpha(U_i, Z)$  then
16     $U_{i+1} = Z$ ;
17  else
18     $U_{i+1} = U_i$ ;
19  end
20  return  $U_{i+1}$ .
21 end
22 Output: final sample  $[U_1, U_2, U_3, \dots]$ 

```

two-step procedure consisting of a *proposal step* (Steps 3-12 of Algorithm 1) and an *acceptance/rejection step* (Steps 13-21 of Algorithm 1). We omit a detailed description of Algorithm 1 and provide an intuitive explanation of the working of the skipping sampler next.




Fig. 1: Illustration of a skipping proposal in \mathbb{R}^2 attempting to sample from the subset A (coloured regions). Starting at X , the initial random-walk proposal Z_1 would be rejected since $Z_1 \notin A$ by a classical MCMC method. Instead, the skipping proposal generates random distant increments and finds new Z_k 's in a linear fashion ('skipping') along the original direction until $Z_k \in A$ or updates are halted.

In essence, the skipping sampler improves the conditional sampling from the subset A by using a specialized proposal step that 'skips' over proposals in A^c (the set of D-LAAs

parameters that do not lead to the activation of an ER) until a point $x \in A$ is sampled, or the skipping process is halted in a randomized fashion (see Fig. 1 for an illustration). A key advantage of the skipping sampler is its ability to efficiently transition between potentially disconnected components of the subset A and different modes of the conditional density π_A . This is particularly important given the nonlinear dynamics of the network and the presence of heterogeneous ER mechanisms, which can result in a disconnected subset A of successful attacks. With a standard MCMC sampling strategy, exploring multiple disconnected components of A would be challenging, leading to incomplete exploration of the range of attack parameters and results heavily dependent on the starting configuration. In contrast, the skipping sampler is designed to effortlessly navigate across different connected components of A . In the context of D-LAA attacks, this means the skipping sampler can generate samples with diverse attack parameters and features, allowing for the discovery of all major network vulnerabilities.

IV. SIMULATIONS

A. Simulation Settings

We implement the power system dynamics and D-LAAs on a Kron-reduced version of IEEE 39-bus test network, consisting of $N = 10$ generation buses (2 of which also have loads present) and $L = 17$ pure load buses [17].

At $t = 0^-$, the system equilibrium power balance is modelled to be in one of four load-scenario states $\tau \in \{1, 2, 3, 4\}$, corresponding to four peaks of the diurnal load cycle of a typical European network, colloquially labelled *night*, *morning*, *afternoon* and *evening* respectively. Equilibrium active loads (and power) are calculated as proportional changes to the published equilibrium power balance of the IEEE 39-bus network [17], whose relative proportions of $\{0.4, 1, 0.85, 1.3\}$ for $\tau = 1, \dots, 4$ respectively are based on the UK power grid's cycle [18]. Broadly speaking, evening and morning periods have the highest load, whereas night-time has the lowest load conditions. The choices of protection system parameters are such that the network is $N - 1$ secure, such that the loss of a single component (generator, load, or line) in the absence of any other disturbance does not trigger a subsequent ER.

To generate dynamic LAAs associated with the disconnection of network components, we utilise the *skipping sampler algorithm* as follows. Starting from an initial nodal LAA load change $\lambda \in \mathbb{R}^L$, we sample During each proposal step, we sample a tuple $x = (\lambda, \mathcal{I}, \tau, C)$ consisting of: the initial nodal LAA vector λ , an interval between attacks \mathcal{I} , a scenario τ and a frequency response C . The sampling parameters are set to $\lambda_{\max}^0 = 1000$ MW, $T_{\max} = 60$ s, $C_{\min} = 0.5$ and $C_{\max} = 5$. We then apply the components of x as an input to the power system model [15] at time $t = 0$, with frequency dynamics simulated for 60s using MATLAB. We conduct $n = 100,000$ proposals, which resulted in a final sample of $S \approx 16,500$ dynamic LAAs that activated at least one ER. This corresponds to an acceptance rate of 16.5%, which is within the optimal 15–48% acceptance rate range [19]. From each sample, we

estimate the metrics of the cascade size \mathcal{X} in MW and the average magnitude of loads changed $\{\Sigma_i(\lambda)\}_{i=1}^L$.

B. Simulation Results

1) **Overview of Cascade Sizes:** The results shown in Fig. 2 indicate that cascades in the system are on average primarily caused by load shedding and RoCoF-induced generation disconnections. These findings have three important implications: (1) globally, the IEEE-39 bus system is most vulnerable to D-LAAs which increase loads, resulting in the preponderance of UFLS events. (2) The network exhibits greater resilience to LAA-induced net reductions in loads, as evidenced by the relatively small proportion of cascades associated with OFGS-related generation disconnections. (3) Among all types of disconnections, individual UFLS events are the most prevalent in the sample. Although each UFLS event only sheds a small magnitude of load in MW, they can be activated multiple times at each node, thus the 700MW of load disconnections in fact consists of a large number of individual UFLS events. In the following sections, we explore how both network and LAA parameters contribute to these results.




Fig. 2: Average cascade size by area for the IEEE 39-bus network.

Impact of the vulnerability ratio (ν) on cascade size:

The vulnerability ratio reflects the proportion of the total load an attacker can potentially control at each node. Hypothesising that $\nu < 30\%$ may be possible for current and near-future power systems, we observe that disconnections in this regime consist of UFLS and RIGS. As per Fig. 3, the IEEE-39 network demonstrated resilience against D-LAAs for $\nu \leq 10\%$, with no components disconnected in this regime as the attacker lacks sufficient leverage to disrupt network operations. For $\nu \in [10, 20)\%$, disconnections are exclusively UFLS, reinforcing the network's susceptibility to this form of disconnection following a D-LAA. Further, the network exhibits greater resilience against generation shedding, with RIGS requiring $\nu \geq 20\%$, and OFGS disconnections $\nu \geq 30\%$.

When $\nu \in [10\%, 60\%]$, our analysis reveals a positive relationship between cascade size \mathcal{X} and ν – as the attacker gains more authority over loads nodes of the network, each effective load change in the dynamic LAA sequence can be larger in magnitude, inducing larger frequency deviations and thus larger cascades. In this vulnerability range, the increase in cascade size is driven primarily by the disconnection of large generating units, mainly due to RoCoF violations, with minor contributions from over-frequency generation shedding.

For $\nu \geq 60\%$, total cascade size growth plateaus while average load shed rises. This marks a phase transition in




Fig. 3: Average cascade size (in MW) by ν .

the susceptibility of the network: for $\nu < 60\%$, generation disconnections drive cascades, while for $\nu \geq 60\%$, they decrease and load disconnections become dominant as attacker leverage over network loads increases.

Interval between dynamic load changes \mathcal{I} : Fig. 4 reveals two regimes particularly susceptible to large cascades: (a) rapidly changing ($\mathcal{I} < 10s$), smaller-magnitude ($\mu(\lambda) < 4GW$) LAAs and (b) static ($\mathcal{I} > 50s$), large magnitude ($\mu(\lambda) > 7GW$) LAAs. D-LAAs in region a result in larger cascades with an average size of 9,000MW and are more threatening to network integrity. Conversely, the region b is associated with large, static attacks, resulting in an average cascade size of 4,000MW. Thus, by deploying multiple attacks, informed either by quasi-real-time or simulated frequency data, an attacker can exacerbate frequency and RoCoF deviations in an intentional and strategic manner, inducing cascades using smaller magnitude D-LAAs when compared to S-LAAs. The efficacy of this strategy is observed in Fig. 4, where, for example, when the interval between attacks is 50s, the attacker must be able to manipulate at least 1,550MW of network loads to trigger a disconnection event. In comparison, when $\mathcal{I} = 10s$, the attacker needs only manipulate $\mu(\lambda) \geq 300MW$ of loads across the network in a coordinated fashion to trigger a disconnection event. For reference, 300MW of loads may represent 170,000 typical space heaters or 45,000 charging EVs, a magnitude of loads possible in modern networks. In general, the lower boundary of Fig. 4 represents the *critical D-LAA characteristic threshold* – the minimum average load change which must be manipulated for each attack interval \mathcal{I} to induce a disconnection event. The positive gradient of this boundary establishes that smaller magnitude D-LAAs require shorter attack intervals to trigger similar-sized cascades.




Fig. 4: Color map showing the cascade size \mathcal{X} (in MW), with respect to the interval between dynamic LAAs \mathcal{I} (s) and the average size of each LAA's load change μ_λ , (GW).

2) *Impact of Power Grid Operating Conditions:* Last, we consider the impact of power grid operating conditions.

Inter-Area Power Balance: Decomposing results by the areas of the IEEE-39 network, cascade sizes \mathcal{X} are related to the initial ($t < 0^-$) net generation imbalance of each area. Area 1, with excess demand, is particularly susceptible to large cascades, half of which are attributable to a large number of loss of load events. This is related to the network’s susceptibility to D-LAAs which increase loads – the excess demand profile of Area 1 is exacerbated by such D-LAAs, triggering UFLS events. Area 3, conversely, with near parity between generation and demand, experiences the smallest cascade sizes on average. Being less dependent on other areas to maintain its power balance, it is more resilient to D-LAAs.




Fig. 5: Average cascade size (in MW) by scenario τ .

Network Load Conditions: On average, \mathcal{X} is minimised during nadir demand (denoted in Fig. 5 as *night*). In this scenario, RIGS dominate cascades, as adversaries induce large changes in frequency through a combination of frequent, alternating changes in loads, or a few, large magnitude increases in load changes, referencing the network’s susceptibility to increases in loads. Conversely, \mathcal{X} is maximised during periods of peak demand. In this scenario, with most of the network’s generating capabilities deployed to serve demand, the network becomes more susceptible to UFLS events from relatively small, positive D-LAA shocks. This is observed in Fig. 5, where UFLS dominates cascades during the evening, peak demand period. Note that, while load disconnections in MWs will naturally be higher during peak demand, the number of UFLS events, which is invariant to the load scenario, is also maximised during the evening (not shown in the figure).

V. CONCLUSIONS AND FUTURE RESEARCH

In this work, we apply a rare-event sampling approach to assess how network variables coupled with LAA parameters influence the size of cascading failures (in MW) in the IEEE39 network. With respect to network variables, our results indicate that the average cascade sizes are larger during peak demand periods and the areas with significant excess demand are particularly vulnerable to cascades. With regard to D-LAA parameters, our results show that the network is resilient against any disconnection event when the proportion of load vulnerable to LAAs is less than 10% of the base load, while a greater amount of vulnerable load leads to increasing cascade sizes. Crucial however is the impact of varying the interval between

attacks, with results clearly highlighting that shorter intervals between attacks enable an attacker to trigger larger cascades while manipulating a smaller quantity of vulnerable loads, within the range of IoT penetration in the present and near future networks. This exposes a key attack strategy that can be exploited by an adversary with access to network data. Our future work includes exploring optimum strategies and methodologies to mitigate the impact of network D-LAAs, including the usage of battery energy storage systems, the tuning of protection systems, and line parameters optimization.

REFERENCES

- [1] 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*, Aug. 2018, pp. 15–32.
- [2] A. Pinceti, L. Sankar, and O. Kosut, “Detection and localization of load redistribution attacks on large-scale systems,” *Journal of Modern Power Systems and Clean Energy*, vol. 10, no. 2, pp. 361–370, 2022.
- [3] S. Lakshminarayana, J. Ospina, and C. Konstantinou, “Load-altering attacks against power grids under covid-19 low-inertia conditions,” *IEEE Open Access Journal of Power and Energy*, vol. 9, pp. 226–240, 2022.
- [4] 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 Trans Smart Grid*, vol. 12, no. 5, pp. 4415–4425, 2021.
- [5] S. Amini, F. Pasqualetti, and H. Mohsenian-Rad, “Dynamic load altering attacks against power system stability: Attack models and protection schemes,” *IEEE Trans Smart Grid*, vol. 9, no. 4, pp. 2862–2872, 2018.
- [6] F. Alanazi, J. Kim, and E. Cotilla-Sanchez, “Load oscillating attacks of smart grids: Vulnerability analysis,” *IEEE Access*, vol. 11, pp. 36 538–36 549, 2023.
- [7] 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.
- [8] T. Nesti, A. Zocca, and B. Zwart, “Emergent failures and cascades in power grids: A statistical physics perspective,” *Physical Review Letters*, vol. 120, no. 25, 2018.
- [9] J. Song, E. Cotilla-Sanchez, G. Ghanavati, and P. D. H. Hines, “Dynamic modeling of cascading failure in power systems,” *IEEE Transactions on Power Systems*, vol. 31, no. 3, pp. 2085–2095, 2016.
- [10] J. M. Stürmer, A. Plietzsch, and M. Anvari, “The risk of cascading failures in electrical grids triggered by extreme weather events,” 2021. [Online]. Available: <https://arxiv.org/abs/2107.00829>
- [11] Y. Dai, M. Noebels, R. Preece, M. Panteli, and I. Dobson, “Risk assessment and mitigation of cascading failures using critical line sensitivities,” *IEEE Transactions on Power Systems*, pp. 1–12, 2023.
- [12] M. P. Goodridge, J. Moriarty, and A. Pizzoferrato, “Distributions of cascade sizes in power system emergency response,” in *Proc. PMAPS*, IEEE, 2020.
- [13] M. P. Goodridge, S. Lakshminarayana, and C. Few, “Analysis of load-altering attacks against power grids: A rare-event sampling approach,” in *Proc. PMAPS*. IEEE, 2022, pp. 1–6.
- [14] J. Moriarty, J. Vogrin, and A. Zocca, “Frequency violations from random disturbances: an MCMC approach,” in *Proc. IEEE CDC*, 2018.
- [15] M. P. Goodridge, J. Moriarty, and A. Pizzoferrato, “A rare-event study of frequency regulation and contingency services from grid-scale batteries,” *Philos. Trans. Royal Soc. A*, vol. 379, no. 2202, p. 20190433, 2021.
- [16] J. Moriarty, J. Vogrin, and A. Zocca, “A Metropolis-class sampler for targets with non-convex support,” *Statist. Comput.*, vol. 31, no. 6, 2021.
- [17] T. M. Athay, R. Podmore, and S. Virmani, “A practical method for the direct analysis of transient stability,” *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-98, pp. 573–584, 1979.
- [18] A. J. Pimm, T. T. Cockerill, and P. G. Taylor, “The potential for peak shaving on low voltage distribution networks using electricity storage,” *Journal of Energy Storage*, vol. 16, pp. 231–242, 2018.
- [19] D. Gamerman and H. Lopes, *Markov Chain Monte Carlo, Stochastic Simulation for Bayesian Inference*, 2nd ed. Chapman & Hall, 2006.