

## A KINETIC MONTE CARLO APPROACH FOR SIMULATING CASCADING TRANSMISSION LINE FAILURE\*

JACOB ROTH<sup>†</sup>, DAVID A. BARAJAS-SOLANO<sup>‡</sup>, PANOS STINIS<sup>‡</sup>, JONATHAN WEARE<sup>§</sup>,  
AND MIHAI ANITESCU<sup>†</sup>

**Abstract.** In this work, cascading transmission line failures are studied through a dynamical model of the power system operating under fixed conditions. The power grid is modeled as a stochastic dynamical system where first-principles electromechanical dynamics are excited by small Gaussian disturbances in demand and generation around a specified operating point. In this context, a single line failure is interpreted in a large deviation context as a first escape event across a surface in phase space defined by line security constraints. The resulting system of stochastic differential equations admits a transverse decomposition of the drift, which leads to considerable simplification in evaluating the quasipotential (rate function) and, consequently, computation of exit rates. Tractable expressions for the rate of transmission line failure in a restricted network are derived from large deviation theory arguments and validated against numerical simulations. Extensions to realistic settings are considered, and individual line failure models are aggregated into a Markov model of cascading failure inspired by chemical kinetics. Cascades are generated by traversing a graph composed of weighted edges representing transitions to degraded network topologies. Numerical results indicate that the Markov model can produce cascades with qualitative power law properties similar to those observed in empirical cascades.

**Key words.** electrical grid, cascading failure, mean first passage time, irreversible diffusion

**AMS subject classifications.** Primary, 60H30, 68U20; Secondary, 37H10

**DOI.** 10.1137/19M1306865

**1. Introduction.** Because of recent blackout events, cascading failure in the power grid—whereby a local network disturbance may cause systemwide disruptions that result in loss of power to significant parts of the network [18, 42, 22, 51, 45]—has become an increasingly important field of study. The uncertain nature of power demand fluctuations in electrical transmission networks, together with the increasing penetration of uncertain power sources such as renewable energy, has made the threat of cascading failure more pronounced. The ability to quantify the probability of infrastructure component failure due to generation and demand variation is critical to the safe operation of the power grid. Furthermore, it is important to understand how these failure probabilities are related to the controllable properties of the network.

\*Received by the editors December 16, 2019; accepted for publication (in revised form) December 3, 2020; published electronically January 25, 2021. The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory (“Argonne”). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under contract DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan. <http://energy.gov/downloads/doe-public-access-plan>.

<https://doi.org/10.1137/19M1306865>

**Funding:** This work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR) under contract DE-AC02-06CH11347. The fifth author was partially supported by NSF awards FP061151-01-PR and CNS-1545046.

<sup>†</sup>Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439 USA (jmroth@mcs.anl.gov, anitescu@mcs.anl.gov).

<sup>‡</sup>Pacific Northwest National Laboratory, Richland, WA 99354 USA (David.Barajas-Solano@pnnl.gov, panos.stinis@pnnl.gov).

<sup>§</sup>Courant Institute of Mathematical Sciences, New York University, New York, NY 10012-1185 USA (weare@cims.nyu.edu).

Previous work in power system reliability has developed methods for studying both individual component and cascading grid failures. In [50, 35], Monte Carlo simulation and large deviation arguments were employed to estimate the probability of individual component failure. Among the various models of cascading failure that have been proposed, a central theme is that of a power law governing the relationship between cascade severity and frequency [18, 40]. In previous work, Dobson et al. [18] identified economic and engineering factors that counterbalance each other to drive the system's operating state toward a critical point. The “slow” dynamics help explain the prevalence of power laws describing cascade severity and have inspired recent work studying the cascading failure mechanism and its relation to power grid operating conditions. In particular, recent approaches have analyzed cascading failures with rare event simulation methodologies (such as splitting and importance sampling) or specialized models that address certain properties of the grid (such as loading levels, as in [43]).

In this work, we continue recent trends that analyze cascading failures with rare event methodologies, but we rely on the theory of large deviations to computationally streamline the calculation of failure probabilities. In particular, we model the dynamics of the power grid from a first-principles, physically inspired approach and reduce the problem of dynamical Monte Carlo simulation at low temperatures to the solution of a nonlinear program (NLP). Drawing inspiration from [13], we introduce stochastic fluctuations in demand and generation to a standard dynamical model of the classical power transmission network [6, 5, 15, 54, 55, 53] and interpret transmission line failures as first exit events for the resulting stochastic dynamics. Despite several assumptions of these classical system models, which include disregarding transmission losses, they have favorable properties, such as port-Hamiltonian structure [48], that render them amenable to analytical treatment.

We employ the theory of small noise transitions for irreversible diffusion processes [11, 20, 26] to derive asymptotically valid expressions that describe the rate of failure of transmission lines in a restricted network. Our analytic expression (i) characterizes the probability density of the time to failure for system dynamics initialized at a stable equilibrium, or “operating point,” (ii) is computable, and (iii) provides a clear and interpretable relation between line failure probability and the properties of the power grid. Based on the analytic approach to describing line failure, we adapt our model to more realistic scenarios (including finite-temperature regimes and more complex networks), and we develop a composite model of cascading failure that aggregates the models of individual component failure.

This manuscript is structured as follows. In section 2 we study the individual line failure problem. We present the stochastic dynamical model of the power grid, define line failure in a restricted model of the power grid, derive expressions for the asymptotic failure rate in this context, and present numerical results that justify our analysis. In section 3 we study potential pathologies of the line failure problem that complicate the straightforward application of results derived in section 2 to real systems. We find that the numerical evidence supports the claim that the simplified, tractable problem studied in section 2 may be a reasonable approximation of the line failure rates of practical interest. From our findings, in section 4 we employ our model of individual line failure to develop a chemical-kinetics-inspired Markov simulator of cascading line failures. In this section, we present validation results in the context of observed cascade data and discuss numerical and theoretical properties of the discrete state space Markov model. In section 5 we present our conclusions and briefly outline areas for future research.

**Contributions.** The key contributions of our work are (i) the development of a tractable expression for the failure rate of an individual transmission line due to small perturbations in power demand and generation, and (ii) the assembly of individual failures in a kinetic Monte Carlo (KMC) setting to describe macroscopic properties similar to real systems.

**2. Unconditional failure rate.** In this section we proceed to define the unconditional line failure problem (Definition 2.2), derive analytic expressions for the rate of failure ((2.26) and (2.27)), and present validation results (Figures 2.3 and 2.4) justifying their correctness.

**2.1. Transmission model.** We model the power transmission network as the undirected graph  $(\mathcal{B}, \mathcal{E})$ , where  $\mathcal{B}$  denotes the set of  $N$  network buses (graph nodes, and short for “busbar,” the physical connection nodes in a transmission network) and  $\mathcal{E}$  denotes the set of transmission lines (graph edges). The set of buses  $\mathcal{B}$  is split into three disjoint subsets corresponding to three types of bus: (i) electricity-producing *generator* buses (alternatively called “PV-buses” after the two exogenous variables associated with the node),  $\mathcal{G}$ , (ii) electricity-consuming *load* buses (alternatively called “PQ-buses” after the two exogenous variables associated with the node),  $\mathcal{L}$ , and (iii) network-balancing *slack* or reference buses (alternatively denoted “SL-buses”),  $\mathcal{S}$  (here we follow standard terminology from the power systems literature [7, 25]<sup>1</sup>). For simplicity in this work we assume that  $|\mathcal{S}| := 1$ , which corresponds to only one slack bus.

To model the short-term dynamics of the transmission network around operating conditions, which we define below, we employ a structure-preserving, energy-based model of port-Hamiltonian form [48, 44] based on the so-called swing equations (see [7] and the supplemental material for more detail). Port-Hamiltonian models differ from traditional Hamiltonian models by allowing nonconservative forces (in our case, damping) to enter the dynamical equations at predefined locations called “ports” while preserving important features of a Hamiltonian structure. Similar models have been employed for stability analysis [5, 15, 13] and cascading failure analysis [54, 53] of power grid systems. In our case, we assume that the short-term dynamics remain applicable over indefinitely long time horizons, since we are interested in studying operation around the synchronous point [44]. To present the dynamics model, we employ the following notation. For a vector  $\mathbb{R}^N$ ,  $[\cdot]_X$  denotes the subvector corresponding to the subset of indices  $X \subseteq [1, N]$ . Additionally,  $\mathbb{I}_X$  denotes the identity matrix of dimension  $|X|$ ,  $\mathbf{0}_{X \times Y}$  denotes the rectangular zero matrix of  $|X|$  rows and  $|Y|$  columns,  $\mathbb{1}_X$  denotes the column vector of ones of dimension  $|X|$ , and  $x^*$  denotes complex conjugation of  $x$ .

The network state at time  $t$  is described by the bus voltage magnitude,  $[V_t]_i$ , the bus phase angle with respect to the angular reference,  $[\theta_t]_i$ , and the bus angular frequency,  $[\omega_t]_i$ , for each bus  $i \in [1, N]$ . The electromechanical dynamics of the power transmission network are modeled by the set of (singularly perturbed) ordinary differential equations [5, 15]

$$(2.1) \quad \dot{x}_t = -K \nabla \mathcal{H}^y(x_t), \quad x_0 = \bar{x}, \quad K := S - J,$$

<sup>1</sup>In load flow analyses, the *slack* bus serves as a source/sink node capable of satisfying any network power imbalances; in lossless networks, this means ensuring that electrical production equals electrical demand. In a power network, the largest generator is traditionally selected to be the slack bus. In power system dynamics, the state space evolution is typically measured relative to a systemwide average phase angle (center of inertia) or a particular component; in our case we give the slack bus the dual role of both balancing the network and serving as the system’s angular reference.

where  $x_t := ([\omega_t]_{\mathcal{S} \cup \mathcal{G}}^\top, [\theta_t]_{\mathcal{G} \cup \mathcal{L}}^\top, [V_t]_{\mathcal{L}}^\top) \in \mathbb{R}^d$ ,  $d = 2N - 1$ , is the state vector of the network dynamics, composed of the vectors of generator angular frequencies,  $[\omega_t]_{\mathcal{S} \cup \mathcal{G}}$ , generator and load bus relative phase angles,  $[\theta_t]_{\mathcal{G} \cup \mathcal{L}}$ , and load bus voltage magnitudes,  $[V_t]_{\mathcal{L}}$ ;  $J$  is a skew-symmetric “connection” matrix, and  $S$  is a diagonal “damping” matrix with nonnegative entries such that (2.1) is of port-Hamiltonian form [48]. Historically, singular perturbation models have been introduced to address limitations in static load models, particularly at low voltages, and here we follow [5]. By  $\mathcal{H}^y(\cdot)$  we denote the system’s energy function parameterized by the network parameters  $y := \{[\delta]_{\mathcal{S}}, [V]_{\mathcal{S} \cup \mathcal{G}}, B^y, P_0^y, Q_0^y, M^g, D^g, D^d, D^\epsilon\}$  and given by [34, 36]

$$(2.2) \quad \begin{aligned} \mathcal{H}^y(x) := & \frac{1}{2} \langle [\omega_t]_{\mathcal{S} \cup \mathcal{G}}, M^g [\omega_t]_{\mathcal{S} \cup \mathcal{G}} \rangle + \frac{1}{2} \langle [v_t]_{\mathcal{L}}, B^y [v_t]_{\mathcal{L}} \rangle \\ & + \langle [P_0^y]_{\mathcal{G} \cup \mathcal{L}}, [\theta_t]_{\mathcal{G} \cup \mathcal{L}} \rangle + \langle [Q_0^y]_{\mathcal{L}}, \log [V_t]_{\mathcal{L}} \rangle, \end{aligned}$$

where  $M^g$  is the diagonal matrix of generation inertia constants,  $B^y$  is the network’s sparse susceptance matrix (the Laplacian weighted by the imaginary part of the reciprocal of the line admittance, describing the transmission system’s interconnections),  $[P_0^y]_{\mathcal{G} \cup \mathcal{L}}$  is the vector of net active power at the  $\mathcal{G} \cup \mathcal{L}$  buses at operating conditions,  $[Q_0^y]_{\mathcal{L}}$  is the vector of net reactive power at the load buses at operating conditions,  $[v_t]_{\mathcal{L}}$  is the vector of voltage phasors defined elementwise at the load buses as  $[v_t]_i := [V_t]_i \exp j[\theta_t]_i$  (where  $j$  denotes the imaginary unit),  $[\delta]_{\mathcal{S}}$  is the angular reference, and  $\langle \cdot, \cdot \rangle$  denotes the (Hermitian) inner product. Since we do not consider variations of the network parameters  $y$  in this work, for the remainder of the article we drop the dependence on  $y$ . This model, like other structure-preserving models, disregards transmission losses so that the line connectivity information is fully encoded in the susceptance matrix. The matrices  $J$  and  $S$  are given by [55] as

$$(2.3) \quad J = \begin{bmatrix} 0 & -(M^g)^{-1}T_1^\top & 0 \\ T_1(M^g)^{-1} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

$$(2.4) \quad S = \begin{bmatrix} (M^g)^{-1}D^g(M^g)^{-1} & 0 & 0 \\ 0 & T_2^\top(D^d)^{-1}T_2^\top & 0 \\ 0 & 0 & (D^\epsilon)^{-1}\mathbb{I}_{\mathcal{L}} \end{bmatrix},$$

where  $D^g$  is the diagonal matrix of generator damping constants,  $D^d$  is the diagonal matrix of real load damping constants, and  $D^\epsilon$  is a singular perturbation parameter governing the system’s proximity to instantaneously satisfying circuit law constraints (i.e., proximity to the power flow manifold [8]). Here, the matrices  $T_1$  and  $T_2$  are given by [55] as

$$(2.5) \quad T_1 = \begin{bmatrix} -\mathbb{1}_{\mathcal{G}} & \mathbb{1}_{\mathcal{G}} \\ -\mathbb{1}_{\mathcal{L}} & \mathbb{0}_{\mathcal{L} \times \mathcal{G}} \end{bmatrix}, \quad T_2 = \begin{bmatrix} \mathbb{0}_{\mathcal{G} \times \mathcal{L}} \\ \mathbb{1}_{\mathcal{L}} \end{bmatrix}.$$

We remark that  $K := S - J$  is invertible and positive definite (though not symmetric) [54], and hence the system qualifies as a “quasigradient” system [14] and is only a linear coordinate-transformation away from a gradient system. A more detailed derivation of our port-Hamiltonian model from the governing equations of power transmission network dynamics is summarized in the supplemental material.

It remains to specify the operating conditions, i.e., the initial conditions  $\bar{x}$  in (2.1). For given network parameters  $y$ , the equilibrium condition  $\nabla \mathcal{H}(x) = 0$ , or, equivalently,  $\arg \min_x \mathcal{H}(x)$ , defines a set of metastable critical points for the deterministic

dynamics (2.1) differing from one another by a factor of  $2n\pi$ ,  $n \in \mathbb{Z}$ , along each angular direction (see Figure 2.1(b)). We denote by  $\bar{x}$  the stable critical point closest to the origin in the angular directions, that is,

$$(2.6) \quad \bar{x} := \arg \min_{x \in \mathcal{X}} |\theta(x)|, \quad \mathcal{X} = \{x : \nabla \mathcal{H}(x) = 0, \nabla^2 \mathcal{H}(x) \succeq 0\}.$$

To account for fluctuations in power demand and generation, we introduce noise into (2.1) through scaled Gaussian perturbations at appropriate indices. Load perturbations are assumed to be the aggregate result of independent consumer device behavior, and generator perturbations are assumed to be small fluctuations in active generation capacity. The perturbations are assumed to affect net real ( $P_0$ ) and net reactive ( $Q_0$ ) power directly and are gathered into a simple additive term in the stochastic dynamics (see the supplemental material for more detail). The noise model proposed here differs slightly from previous work [13, 30] through the inclusion of generator perturbations (which are associated with generator damping, introduced for numerical purposes), but its specific form allows us to perform the analysis presented in subsection 2.3. Our stochastic port-Hamiltonian model for the grid dynamics then becomes [13]

$$(2.7) \quad dx_t^\tau = (J - S)\nabla \mathcal{H}(x_t^\tau) dt + \sqrt{2\tau}S^{1/2}dW_t, \quad x_0^\tau = \bar{x},$$

where  $W_t \in \mathbb{R}^d$  is a vector of  $d$  independent Wiener processes and the parameter  $\tau$  defines the strength of the scaled load and generation fluctuations.

We mention three important properties of the stochastic differential equation (SDE) (2.7), which we summarize in the following remark.

*Remark 2.1* (SDE properties).

1. The SDE is *irreversible* because of the presence of the antisymmetric “connection” matrix  $J$  [47], and it represents a form of nonequilibrium Markovian dynamics. Typically, irreversibility makes a system difficult to analyze; however, this is mitigated in our case by the following property.
2. System (2.7) admits a transverse decomposition of the drift [20, 11, 10]; see Appendix A.
3. The stochastic port-Hamiltonian model (2.7) has the Gibbs measure

$$(2.8) \quad \mu^\tau = (Z^\tau)^{-1} \exp[-\mathcal{H}(x)/\tau], \quad Z^\tau := \int \exp[-\mathcal{H}(x)/\tau] dx$$

(for  $\tau > 0$ ), as its stationary distribution [24].

We choose a stochastic model for the grid dynamics with such properties (specifically the final two) because they will greatly simplify analyzing transmission line failures as first exit events [20, 11].

**2.2. Unconditional failure problem.** We now motivate our underlying problem, formally introduced in Definition 2.2. For the power transmission network described by the model (2.7), we are interested in quantifying the time it takes for the line energy of the  $l$ th line, denoted  $\Theta_l$ , to exceed a given safety threshold  $\Theta_l^{\max}$ . Disregarding transmission losses, the line energy  $\Theta_l$  is given in terms of the network’s state  $x$  by

$$(2.9) \quad i_l = (v_i - v_j) B_{ij},$$

$$(2.10) \quad \Theta_l(x) = |i_l|^2 = i_l i_l^* = B_{ij}^2 [V_i^2 - 2V_i V_j \cos(\theta_i - \theta_j) + V_j^2],$$

where  $(i, j)$  denote the start and end buses of the  $l$ th line, respectively;  $B_{ij}$  denotes the susceptance of the  $l$ th line; and  $i_l$  denotes the current phasor of the  $l$ th line. The safety criterion  $\Theta_l(x) < \Theta_l^{\max}$  defines a safety region in state space

$$(2.11) \quad D_l := \{x: \Theta_l(x) < \Theta_l^{\max}\},$$

which contains the equilibrium point  $\bar{x}$ . When the  $l$ th line energy exceeds the corresponding threshold, the line is said to fail irreparably, resulting in the disconnection of the line from the network. Therefore, we can interpret the  $l$ th line failure as the first exit of the stochastic dynamics (2.7) from the basin of attraction of  $\bar{x}$  through the boundary  $\partial D_l$  of the safety region (see Figure 2.1(a)). Note that each of the metastable equilibria has a safety region associated to it (see Figure 2.1(b)), but we are only concerned with the safety region that contains the operating conditions defined in (2.6).

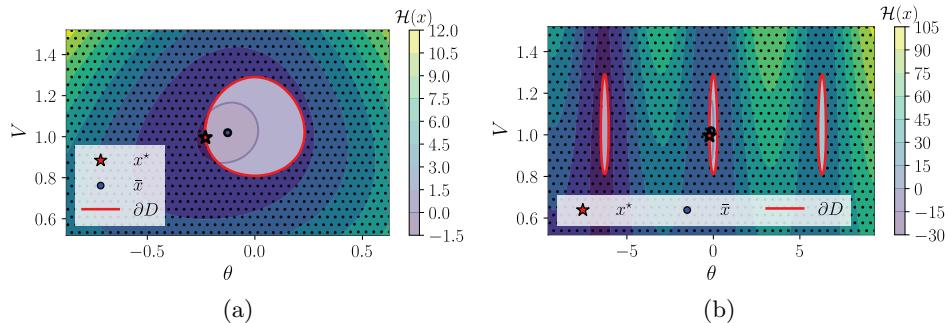


FIG. 2.1. Energy surface contour and failure region (dotted area) for a slack-load transmission line (line 2) in the 3-bus power transmission system [54]; see section SM2.3 of the supplement for more detail. (a) Local basin of attraction of  $\bar{x}$  and safety region. (b) Various metastable configurations with their respective basin of attraction and safety regions.

In a power system, the automatic line relay trigger event effectively removes that line from the network, which we model by setting its susceptance to zero, though continuous formulations have been proposed [54]. This discrete change to the system matrix  $B$  affects the energy function as well as the system's subsequent dynamics but not before the first failure, thus leaving the relevant properties of Remark 2.1 intact. When following the network's dynamics, the trajectory may wander into a region of state space where  $\Theta_k \geq \Theta_k^{\max}$  for  $k \neq l$ , thereby triggering the relay corresponding to line  $k$  before triggering the relay corresponding to line  $l$  of interest. In section 2, we assume that line  $l$  of interest is the only fallible line in the network. This consideration is summarized in the following definition.

**DEFINITION 2.2** (unconditional and conditional failure events). *An unconditional failure event for line  $l$  is defined as an event at time  $t$  such that  $\Theta_l(x_t) \geq \Theta_l^{\max}$ , regardless of any previous failure events  $\Theta_k(x_s) \geq \Theta_k^{\max}$  for  $k \neq l$  and  $s < t$ . This terminology arises in contrast to conditional failure events, in which the event of interest is the line  $l$  failure (transgression through  $\partial D_l$ ) conditioned on this transgression being the first transgression through any boundary for the full network.*

The conditional problem is the appropriate framework for computing the failure rate of practical interest, but drawing a mapping between the two types of failure problems is difficult. We return to the conditional-unconditional failure distinction in section 3 and study the unconditional failure problem until then.

**2.3. Quasistationary exit rate.** For the stochastic process  $x_t^\tau$  defined by (2.7), we are interested in the probability distribution of the first exit time, or “first passage time,”

$$(2.12) \quad T_{\partial D_l}^\tau := \inf \{t > 0 : x_t^\tau \in \partial D_l\},$$

where the boundary  $\partial D_l$  of the safety region  $D_l$  satisfies the following conditions:

- (A.i)  $\langle (J - S)\nabla \mathcal{H}(x), n(x) \rangle < 0$  for all  $x \in \partial D_l$ , where  $n(x)$  is the outward vector normal to  $\partial D_l$  at  $x$ . Equivalently, the surface  $\partial D_l$  is *noncharacteristic*.
- (A.ii)  $\langle \nabla \mathcal{H}(x), n(x) \rangle > 0$  for all  $x \in \partial D_l$ .

These assumptions ensure that the deterministic dynamics started at points along  $\partial D_l$  converge toward the stable point  $\bar{x}$ . Numerical analysis of various IEEE power transmission system test cases indicates that these assumptions are in general valid for line failure criteria of the form (2.10).

For the stochastic port-Hamiltonian system and line failure criteria satisfying assumptions (A.i) and (A.ii), we have that, in the limit  $\tau \rightarrow 0$ , the mean first passage time (MFPT)  $\mathbb{E}T_{\partial D_l}^\tau$  satisfies the limit [26]

$$(2.13) \quad \lim_{\tau \rightarrow 0} \tau \log \mathbb{E}T_{\partial D_l}^\tau = \min_{x \in \partial D_l} V(\bar{x}, x),$$

where  $V(\bar{x}, x)$  is the so-called Freidlin–Wentzell quasipotential [20]. Since our port-Hamiltonian system admits a transverse decomposition (Remark 2.1), calculation of the quasipotential is simplified considerably [26]. In this case, by definition, the quasipotential satisfies the properties  $V(\bar{x}, \bar{x}) = 0$  and  $V(\bar{x}, x) \geq 0$  for all  $x \in D_l$  and solves the Hamilton–Jacobi equations [26] for (2.7):

$$(2.14) \quad \langle \nabla V(x), SV(x) \rangle + \langle (J - S)\nabla \mathcal{H}(x), \nabla V(x) \rangle = 0, \quad x \in D_l.$$

One can verify by substitution that

$$(2.15) \quad V(\bar{x}, x) := \mathcal{H}(x) - \mathcal{H}(\bar{x})$$

satisfies (2.14) and the properties of the quasipotential listed above. Here we pause to make the following remark.

*Remark 2.3* (quasipotential evaluation). Since (2.15) satisfies the definition of a quasipotential for dynamics (2.7), quasipotential evaluation is equivalent to point evaluation of the energy function and *does not* require path optimization to compute the Freidlin–Wentzell action functional [20] (which forms the basis for most quasipotential computations). Fundamentally, this is due to the existence of a transverse decomposition for the drift vector.

The limit (2.13) indicates that the MFPT is of order  $\exp[V(\bar{x}, x^*)/\tau]$ , where we define  $x^*$  (the unconditional “exit point”) as the solution to the following nonlinear program:

$$(2.16) \quad x^* = \arg \min_{x \in \partial D_l} \mathcal{H}(x),$$

which corresponds to the most likely exit point of the stochastic dynamics along the surface  $\partial D_l$  in the limit  $\tau \rightarrow 0$ . For the remainder of this section, we assume that  $x^*$  is unique, although we return to this issue in section 3.

Over the period of time  $1 \ll t \ll \exp[V(\bar{x}, x^*)/\tau]$ , the diffusion process is characterized by a “quasistationary” (metastable) equilibrium in which the process forgets

its initial condition and the probability mass is concentrated around the equilibrium point  $\bar{x}$  according to the quasistationary density [11]. Probability mass is then lost through the boundary  $\partial D_l$  at an approximately constant rate  $\lambda^\tau$  known as the “quasistationary exit rate.” It is known that for Markov processes starting from quasistationary distributions, the distribution of first passage times is exponential with mean given by the reciprocal of the exit rate [12, 28]. Thus, in order to characterize the distribution of first passage times for the quasistationary range in the limit  $\tau \rightarrow 0$ , it suffices to approximate  $\lambda^\tau$ .

**2.4. Asymptotic expression for the exit rate.** We derive zeroth- and first-order approximations for the unconditional rate of transmission line failure.

**2.4.1. Zeroth-order approximation.** We derive an approximation to the quasistationary exit rate asymptotically valid in the small-noise limit. For this purpose we follow the procedure presented in [11] for irreversible diffusion processes. Here, we note that the theory of [11] was derived for the case of a full-rank diffusion matrix, which in our case would correspond to a full-rank matrix  $S$ . Unfortunately,  $S$  is diagonal with nonnegative, rather than strictly positive, entries, and thus the assumption of full rank does not hold.<sup>2</sup> Therefore, to employ the theory of [11], we must introduce the following additional assumption on the exit surface:

(A.iii)  $\langle n(x), Sn(x) \rangle > 0$  for all  $x \in \partial D_l$ ;

that is, we require there to be a diffusive probability flux in the direction orthogonal to the exit surface. This assumption limits the types of transmission lines that can be analyzed with the proposed approach. In particular, *generator-generator* (including *slack-generator*) lines do not satisfy this assumption for the proposed dynamics and line failure models; however, they may be captured by incorporating dynamical equations for the voltage magnitudes (and corresponding excitation noise) at the generator buses as in [13]. Fortunately, *generator-generator* lines are relatively insignificant in practice, since most *generator-generator* lines contain an intermediary substation. To illustrate this point, we evaluate the number of *generator-generator* lines in sixteen IEEE cases (that are identified by the number of buses they contain) and range in size from nine buses to 9,241 buses (with a mean of 2,302 buses). The IEEE cases are widely used synthetic models that resemble actual power systems.<sup>3</sup> In these IEEE cases, *generator-generator* lines are either uncommon or can be addressed by introducing a fictitious bus between the two nodes. Only one such line exists in the IEEE 30-bus system (see Figure 2.2(b)), and across the remaining systems, the proportion of such lines ranges between 0% and 30% (with each case falling under 7% except for the IEEE 118-bus system at 30%).

Having addressed this technical challenge, we can employ the results of [11]. Specifically, injecting the transverse decomposition of the drift of (2.7), derived in Appendix A, into (4.25) of [11], we obtain the following zeroth-order expression for the quasistationary exit rate in integral form:

$$(2.17) \quad \lambda_0^\tau = \int_{\partial D_l} \sqrt{\frac{\det \text{Hess } \mathcal{H}(\bar{x})}{(2\pi\tau)^d}} \exp\left[-\frac{\mathcal{H}(x) - \mathcal{H}(\bar{x})}{\tau}\right] \langle (J + S)\nabla \mathcal{H}(x), n(x) \rangle \, dA(x).$$

We can see that in the limit  $\tau \rightarrow 0$ , the value of this integral is dominated by the

<sup>2</sup>Zero elements correspond to the generator angle equations due to  $T_2^\top (D^d)^{-1} T_2$ .

<sup>3</sup>The IEEE cases were developed mainly for steady state analysis, but they are also frequently used for dynamic analysis due to their availability.

contribution of the exponential term  $\exp\{-[\mathcal{H}(x) - \mathcal{H}(\bar{x})]/\tau\}$  in the vicinity of the exit point. Therefore, we can employ the Laplace method for multiple integrals [52] to obtain a formula for the leading asymptotic behavior of the exit rate. This calculation is summarized in Appendix B. The resulting expression for the quasistationary exit rate is

$$(2.18) \quad \lambda_0^\tau \underset{\tau \rightarrow 0}{\sim} \nabla^\top \mathcal{H}(x^*) S \nabla \mathcal{H}(x^*) \sqrt{\frac{\det \text{Hess } \mathcal{H}(\bar{x})}{2\pi\tau|B^*|}} \exp\left[-\frac{\mathcal{H}(x^*) - \mathcal{H}(\bar{x})}{\tau}\right],$$

where  $B^*$  is a factor accounting for the curvature of  $\partial D_l$  in the vicinity of  $x^*$  (see (2.16)), given by

$$(2.19) \quad B^* \equiv \nabla_x \mathcal{H}(x^*)^\top L^{-1} \nabla_x \mathcal{H}(x^*) \det L, \quad L = \text{Hess } \mathcal{H}(x^*) - k \text{Hess } \Theta_l(x^*),$$

and  $k$  is the Lagrange multiplier on the  $\Theta_l$  constraint in the constrained optimization problem of the exit point (2.16).<sup>4</sup> For convenience, we refer to the subexponential term as the *prefactor* and the exponential term as the *energy factor* and write

$$(2.20) \quad \lambda_0^\tau \underset{\tau \rightarrow 0}{\sim} \frac{C^* C_0(\bar{x})}{\tau^{1/2}} \exp\left[-\frac{\Delta\mathcal{H}}{\tau}\right],$$

where we define

$$(2.21) \quad C^* := \frac{\nabla^\top \mathcal{H}(x^*) S \nabla \mathcal{H}(x^*)}{\sqrt{2\pi|B^*|}}, \quad C_0(\bar{x}) := \sqrt{\det \text{Hess } \mathcal{H}(\bar{x})},$$

and for convenience we also define (with implicit dependence on a line of interest)

$$(2.22) \quad \Delta\mathcal{H} := \mathcal{H}(x^*) - \mathcal{H}(\bar{x}).$$

It follows that the MFPT is given by  $\mathbb{E}T_{\partial D_l} \underset{\tau \rightarrow 0}{\sim} 1/\lambda_0^\tau$ .

Our zeroth-order expression for  $\lambda_0^\tau$  has the following properties. First, it is straightforward to evaluate, since it requires solving only two optimization problems (an unconstrained optimization problem for  $\bar{x}$  and a constrained optimization problem for the exit point  $x^*$ , (2.16)) and computing the curvature factor. Second, it shows how the exit rate (and thus the probability distribution of line failure times) depends on the properties of the network. Namely, the rate depends on the properties of the network directly through the energy function  $\mathcal{H}(x)$  and the matrix  $S$ , and indirectly through the values of  $\bar{x}$ ,  $x^*$ , and  $B^*$ .

**2.4.2. First-order approximation.** Next we extend the applicability of the zeroth-order failure rate to finite temperature regimes. Temperature enters the expression for line failure rate (2.20) in two competing ways. As temperature decreases, the subexponential prefactor ( $C^* C_0(\bar{x}) \tau^{-1/2}$ ) increases and the exponential energy factor ( $\exp[-\Delta\mathcal{H}\tau^{-1}]$ ) decreases. In the small-noise limit, the effect of the prefactor is dwarfed by the contribution of the energy factor. In the large-noise limit, the  $\tau^{-1/2}$  dependence in the prefactor leads to nonphysical behavior. Specifically, the failure rate decreases as temperature exceeds a “critical” temperature  $\tau$  of the same order of magnitude as  $\Delta\mathcal{H}$  where the energy factor contribution is approximately unity, and the system is expected to fail more easily in this regime.

<sup>4</sup>In numerical computation, we compute (2.18) and related expressions using the *magnitude* of  $\det \text{Hess } \mathcal{H}(\bar{x})$ .

We note that the subexponential prefactor stems from the normalization constant of the (quasi)stationary distribution  $p_{st}^\tau$  (see (2.8)). A higher-order analysis of  $p_{st}^\tau$  results in a modification to the failure rate, and we offer a brief physical justification of its appropriateness.

First, we examine the higher orders of the WKB approximation [4] for the stationary distribution  $p_{st}^\tau$  given by

$$(2.23) \quad p_{st}^\tau(x) = \frac{C_{st}^\tau(x)}{\tau^{1/2}} \exp \left[ \frac{\sum_{j=0}^{\infty} \mathcal{H}_j(x) \tau^j}{\tau} \right],$$

where  $\mathcal{H}_0(x) = -\mathcal{H}(x)$  and  $\mathcal{H}_j$  for  $j = 1, \dots$  are unknown functions of  $x$ . We show in Appendix C that for the power grid model we consider, the higher-order terms in the exponential are constant in  $x$  (depending only on  $\bar{x}$ ) and can be absorbed into the subexponential prefactor. The argument is built order by order in  $\tau$  using the Fokker–Planck equation for the stationary density  $p_{st}^\tau(x)$ . For each order, we use the result that the previous order approximation is constant in  $x$  and the two properties of the power grid model: (i) the matrix  $S$  is constant, and (ii)  $F(x) := \operatorname{div} J \nabla \mathcal{H}(x) + \langle A(x), S \nabla \mathcal{H}(x) \rangle = 0$ , where  $A_i(x) = \sum_{j=1}^d \partial_{ij} S_{ij}(x)$  (see Appendices A and C for more details).

Second, we examine the corrections that come from the expansion of the prefactor  $C_{st}^\tau(x)$ . We have chosen an expansion in integer powers of  $\tau$  given by

$$(2.24) \quad C_{st}^\tau(x) = \sum_{j=0}^n C_j(x) \tau^j.$$

We show in Appendix C that

$$(2.25) \quad C_{st}^\tau(x) = \sum_{j=0}^n C_j(\bar{x}) \tau^j,$$

where the  $C_j(\bar{x})$  in (2.25) are constants. As in the case of the higher-order terms in the exponent, the argument is again built order by order in  $\tau$  and utilizes the same properties of the power grid model (see Appendix C for more details).

Using a zeroth-order approximation of the prefactor (as in (2.18)), we obtain for the failure rate

$$(2.26) \quad \lambda_0^\tau \underset{\tau \rightarrow 0}{\sim} C^* C_0(\bar{x}) \tau^{-1/2} \exp \left[ -\frac{\Delta \mathcal{H}}{\tau} \right].$$

Similarly, using a first-order approximation of the prefactor, we obtain for the failure rate

$$(2.27) \quad \lambda_1^\tau \underset{\tau \rightarrow 0}{\sim} \left( \frac{C^* C_0(\bar{x})}{\tau^{1/2}} + \frac{C^{**} C_1(\bar{x}) \tau}{\tau^{1/2}} \right) \exp \left[ -\frac{\Delta \mathcal{H}}{\tau} \right]$$

for new constants  $C^{**}$  and  $C_1(\bar{x})$ .

Determining analytical approximations of  $C_1(\bar{x})$  and  $C^{**}$  requires a more detailed analysis of the quasistationary distribution. Instead, we present a heuristic supported by numerical evidence (Figure 2.4). In particular, at  $\tau \approx \Delta \mathcal{H}$ , the exponential factor is of order 1 and does not significantly influence the failure rate, which is dominated by  $C^* C_0(\bar{x}) \tau^{-1/2}$ . To interpolate between the low-temperature asymptotics dominated

by the energy factor and the high-temperature asymptotics dominated by  $\alpha\tau^{1/2}$  (for  $\alpha > 0$ ), we take  $\alpha = C^*C_0(\bar{x})\tau^{-1/2}$  to match the coefficient of the zeroth-order failure rate (2.26) at  $\tau = \Delta\mathcal{H}$ . This amounts to taking

$$(2.28) \quad C_1(\bar{x}) = C_0(\bar{x}) \quad \text{and} \quad C^{**} = \frac{C^*}{\Delta\mathcal{H}},$$

which has the added benefit of not requiring estimation of further parameters. From a physical perspective, we expect the failure rate to increase with temperature. The form we propose aligns with this intuition and interpolates between two regimes at the hypothesized critical point.

**2.5. Numerical validation of the asymptotic exit rate.** The goal of our numerical experimentation is to study the asymptotic properties of the failure rate theory expressed in (2.26) and (2.27). In particular, our experiments explore the theory's ability to measure system stability directly via initial operating conditions. The theory employed in subsection 2.4 provides guidance only for the small-noise limit, and we focus first on exploring asymptotic behavior before turning to the question of quantifying *when* such a regime might begin (if at all). In the remainder of this section, we present our simulation methodology and discuss our numerical findings.

**2.5.1. Methodology.** As introduced in section 1, the unconditional failure problem that we are concerned with measures the system's first exit through  $\partial D_l$  for a particular line  $l$  of interest, ignoring any previous transgressions through boundaries associated with lines other than the line of interest. The related *conditional* problem is the appropriate framework for computing the failure rate of practical interest, but here we focus on the unconditional failure problem and return to the unconditional-conditional failure issue in section 3.

To simulate an unconditional problem, we consider one boundary at a time and treat the line  $l$  of interest as an isolated component in a modified network where  $\Theta_k^{\max} \equiv \infty$  for lines  $k \neq l$ . Given specified operating conditions consisting of fixed demand, generation, and network topology (which can be summarized in parameters  $y^5$ ), the system is initialized at equilibrium point  $\bar{x}$  (2.6), the minimizer of the resulting energy  $\mathcal{H}^y(x)$ . At equilibrium, we assume that the grid state is synchronized, meaning that the momentum variables  $[\omega]_{\mathcal{G} \cup \mathcal{S}}$  are identically zero. Operating parameters are then augmented with dynamics parameters (including generator damping and masses) to specify the energy function, which remains continuous for the duration of the simulation. Next, we assume that emergency line limits ( $\Theta^{\max}$ ) are scaled as a constant factor of the base limits<sup>6</sup> (see Figure 2.2(c)). We select an individual line  $l$  for study and integrate the dynamics (2.7) with the Leimkuhler–Matthews (LM) method<sup>7</sup> until line  $l$  failure is observed. The procedure is summarized in Algorithm D.1.

We use the **Matpower** software package [37] implementation of the IEEE 30-bus system [1] (see Figure 2.2(b)) as a problem instance. It is a standard test system that contains six generators, 24 loads, and 41 transmission lines and was chosen in part because it also specifies line limits. This avoids the nontrivial issue of setting limits

<sup>5</sup>For fixed demand ( $P$  and  $Q$ ) data, only  $[V]_{\mathcal{G} \cup \mathcal{S}}$  and  $[P^g]_{\mathcal{G}}$  need to be specified; these can be obtained by solving an optimal power flow (OPF) [7, 19].

<sup>6</sup>Base limits are given by the first line MVA rating field in the IEEE case file format.

<sup>7</sup>The LM method is a modified Euler–Maruyama scheme in which each step's stochastic increment is an average of (a) the previous timestep's standard Wiener increment, and (b) the current timestep's standard Wiener increment [27]. In [30], the LM method was found to be more accurate than either vanilla Euler or Heun discretization schemes.

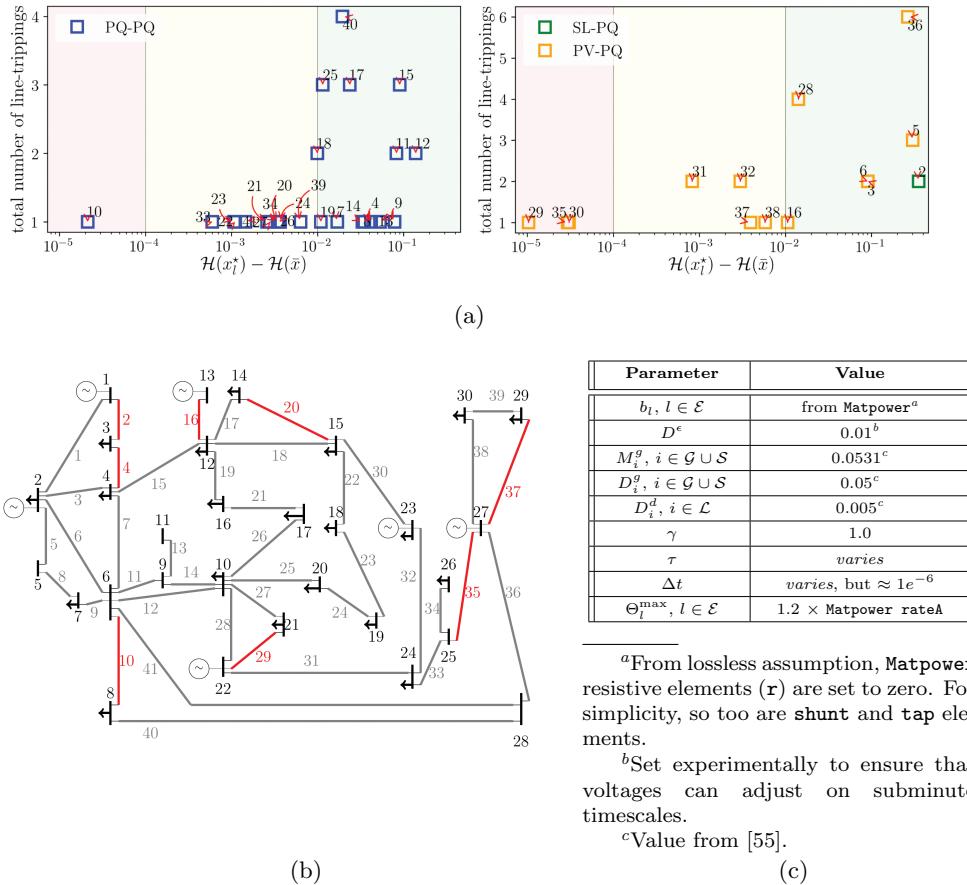


FIG. 2.2. (a) Energy difference between equilibrium and line  $l$  failure point ( $\mathcal{H}(x_l^*) - \mathcal{H}(\bar{x})$ ) for each acceptable load-load line (left) and generator-load line (right) plotted against the total number of failed lines at  $x_l^*$ . Particularly at dense connection points (such as buses 6 and 10), failure of a certain line may require simultaneous failures at other lines. (b) IEEE 30-bus network diagram [23]; black lines/labels correspond to buses (“~” denotes the presence of electrical generation and “←” denotes the presence of electrical demand; node 1 is assumed to be the slack bus), red lines/labels correspond to lines studied in subsection 2.5, and gray lines/labels correspond to lines not studied. (c) Parameter table for dynamics simulations. The singular perturbation parameter  $D^e$  governs the deviation from the power-flow manifold. Smaller values indicate closer proximity to the manifold but are associated with larger load voltage fluctuations (both deterministic and stochastic).

appropriately [29], although a reasonable proxy might be to ensure a standard level of system security, such as  $N-1$  stability [2, 31]. Of the 41 total system lines in the IEEE 30-bus system, 40 satisfy the required assumptions for our exit analysis (assumptions (A.i)–(A.iii); see also section 2.2.1 in [11]). Line 1 (Figure 2.2(b)) is a *generator-generator* line and violates assumption (A.iii) as discussed in subsection 2.4.1. For the set of 40 lines, we compute failure points, scan for pathologies (Definitions 3.2 to 3.4), and stratify lines based on type (*load-load*, *generator-load*, or *slack-load*) and risk class (*high*, *medium*, or *low*) using each line’s computed energy difference  $\Delta\mathcal{H}$  as a proxy for risk. We then choose eight representative lines to study: one (nonpathological) line from each risk class for both *load-load* and *generator-load* types, one *slack-load* line (the only such line in the network that is a pathological line), and one additional (nonpathological) high-risk line. For each simulation, we initialize the dynamics at

the same equilibrium point  $\bar{x}$  determined from the IEEE case file<sup>8</sup> and integrate until line  $l$  failure. We repeat failure experiments across temperatures ranging from  $\tau = 1.0$  to  $\tau = 9.8 \times 10^{-7}$  (though the range of temperatures is line dependent and limited by computational cost).

Briefly, we draw attention to a few details regarding the numerical implementation. First, we are limited in gathering low-temperature data. Using 40 parallel processors, Monte Carlo simulations for an individual line were performed with an approximate computation budget of 24 hours per line per temperature value. Given the exponential scaling of the failure rate (2.27) at low temperatures, the limiting experiment is the lowest temperature for each line. Second, since the additive noise term in the stochastic port-Hamiltonian model (2.7) is independent of the state variables  $x$ , the Euler–Maruyama integration scheme is identical to the Milstein method and has order-1 strong convergence<sup>9</sup> [32, 21]. We also note that because the system (2.7) originates from a singularly perturbed differential algebra equation model, the dynamics are stiff and require a “small” timescale discretization in order to handle the singular perturbation using simple explicit integrators. Determining an appropriate  $\Delta t$  is case dependent, and in what follows we take steps on the order of  $10^{-6}$  seconds.

**2.5.2. Experimental results.** The purpose of this section is threefold. First, we present and discuss empirical convergence and applicability results from our unconditional failure simulations. Second, we return to the question of adjusting the failure rate at high temperatures. Third, we discuss conclusions and limitations of our experimental framework and note that ancillary simulation results are included in the supplemental material.

**Failure simulations and convergence diagnostics.** Fundamentally, we are concerned with understanding the approximation quality of (2.26) and (2.27) to characterize unconditional failure rates for our model of grid dynamics (2.7). Relying on the asymptotic theory of subsection 2.4, we use absolute logarithmic error (henceforth  $\delta$ , defined as  $|\log(\lambda_{sim}^\tau / \lambda^\tau)|$ ) as a metric for convergence. Absolute logarithmic error symmetrically penalizes deviations from 1 and is measured in logarithmic units. We compute  $\delta$  for each line across the range of temperatures studied (findings are summarized in Table 2.1) and present detailed diagnostics for two particular simulation cases representing the best and worst lines (as measured by  $\delta$  averaged across all temperature experiments). Figure 2.3 displays the detailed convergence results for analytic zeroth-order and the heuristically determined first-order failure rate approximations.

In Figure 2.3, we observe good agreement between the analytic rates and the simulated rate as temperature decreases (Figures 2.3(I.a), (I.b) and 2.3(II.a), (II.b)), as predicted by theory. In both cases, the ratio between simulated and approximate failure rates shows encouraging progress toward unity as temperature decreases. The relationship between approximation quality and temperature is moderated by numerical parameters such as integration step size, but we find that numerical error can be controlled by setting integration parameters appropriately. Accordingly, we expect agreement to improve at smaller step sizes and lower temperatures. Further, we re-

<sup>8</sup>In the 30-bus system, we use demand and topology data from the case file and solve a lossless OPF problem to determine dispatch settings. Under the prespecified line limits, a few lines are heavily stressed at the resulting  $\bar{x}$  (lines 10, 29, and 35 have initial line energy loadings that are close to 100% of their capacity, as specified by the `RateA` limits in the case files). For comparison, we also solve the OPF without transmission losses and observe similar congestion, indicating that the lossless assumption is not the cause.

<sup>9</sup>We expect this to hold for the LM scheme as well.

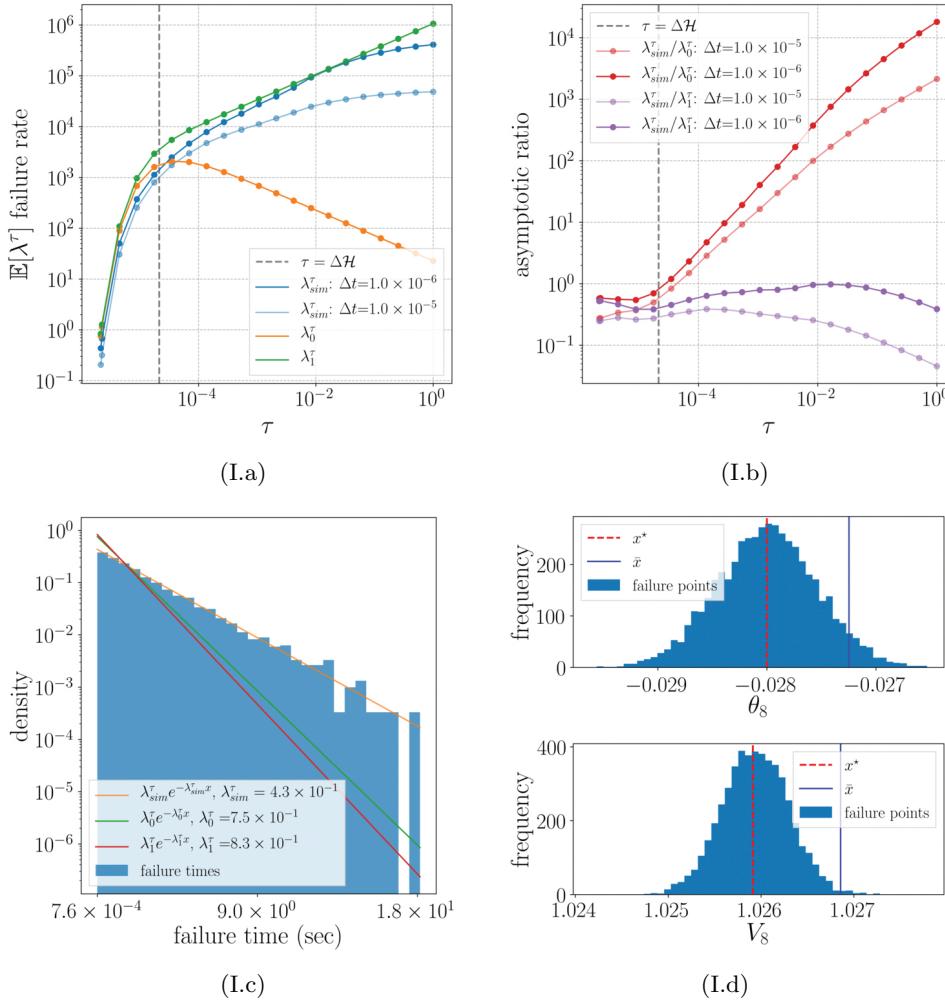


FIG. 2.3. Unconditional line failure diagnostics for the best (5,000 simulations of line 10, part (I)) and worst (500 simulations of line 16, part (II); see continuation of figure on following page) experiments, based on  $\delta$  (see Table 2.1). (a) Failure rate comparison across temperature. (b) Asymptotic ratio across temperature. (c) (Scaled) histogram of unconditional failure times at lowest temperature tested. (d) Histogram of voltage, angle, or frequency state variable at observed failures relative to  $\bar{x}_i$  and  $x_i^*$ , where  $i$  corresponds to the bus index of (one of) the line's connecting buses; shown at lowest temperature tested.

mark that this behavior is typical across all lines studied, as evidenced by the average (first-order) error of less than one logarithmic unit (see Table 2.1).

Panels (c)–(d) in parts I and II of Figure 2.3 present additional diagnostics. In panel (c) in parts I and II, the distribution of exit times approximately follows an exponential distribution parametrized by the reciprocal of the mean of the empirical data, the simulated mean failure rate. This is encouraging because it indicates that expressions (2.26) and (2.27) describe the distribution of failure times in addition to the mean failure time. Further, at low temperatures, the zeroth- and first-order failure rate expressions appear to provide a conservative approximation of the simulated failure rate. In panel (d) in parts I and II, we display the value of voltage, angle, and

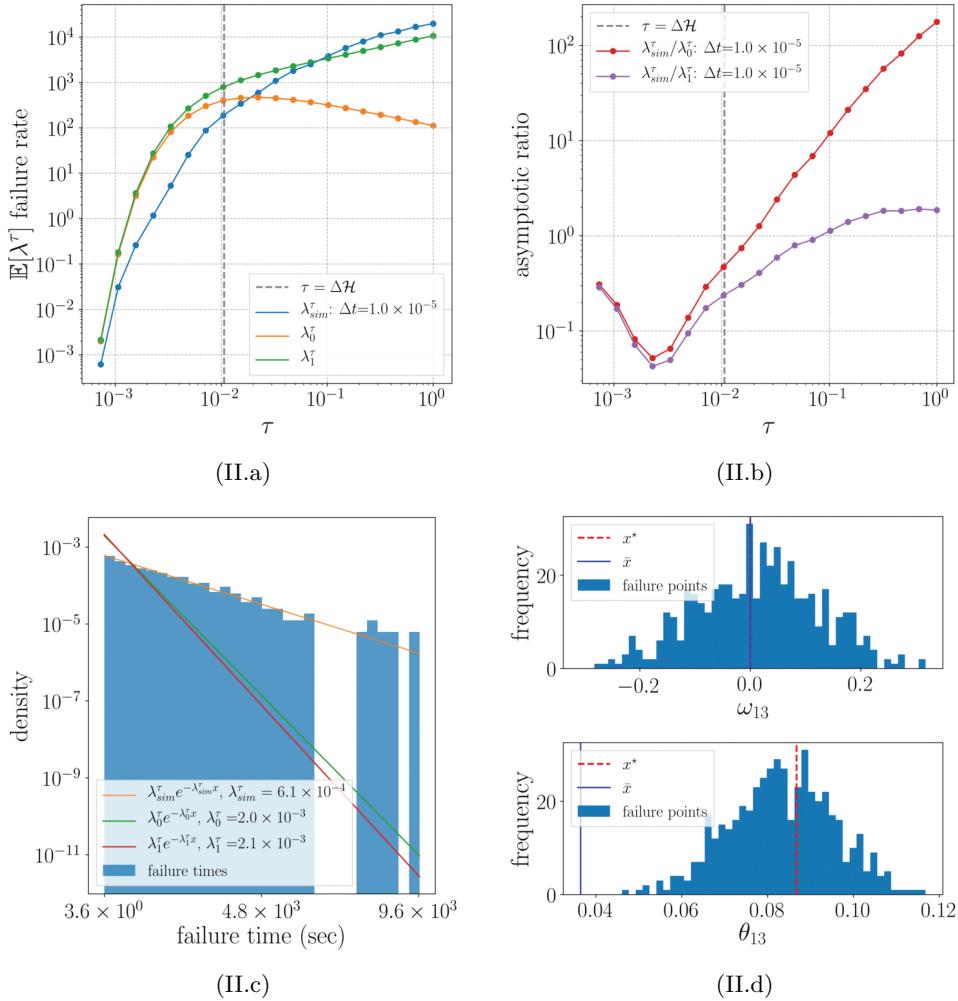


FIG. 2.3. (cont.)

frequency state variables at simulated failure points. Since observed failures tend to cluster around the lowest energy failure point  $x^*$ , in our opinion this provides evidence of the reasonableness of approximating the flux through the failure boundary by the first-order Laplace approximation at  $x^*$ .

Next we turn to the question of determining when the approximation belongs to a “small-noise” or “large-noise” asymptotic regime. In Figure 2.3, panels (a)–(b) in parts I and II show that  $\tau \approx \Delta\mathcal{H}$  is an approximate critical point for each line. We remark that since  $\Delta\mathcal{H}$  is a line-dependent quantity, regime classification by this measure is also line dependent. For example, as depicted in panel (b) in parts I and II, the exponential factor begins to contribute for  $\tau < \Delta\mathcal{H}$  in both simulations, and the failure ratio begins to improve soon thereafter. Conversely, the prefactor contribution begins to dominate for  $\tau > \Delta\mathcal{H}$ . In the zeroth-order approximation, this regime results in divergence of theory and simulation since the prefactor is dominated by  $\tau^{-1/2}$  but failure rate increases with increasing temperature. The first-order approximation captures this intuition, however, and we observe that failure rate scaling of  $\lambda^\tau \propto \tau^{1/2}$

generally well approximates the simulated behavior (except for line 37, which exhibits scaling  $\lambda^\tau \propto \tau^r$  for  $r > 1/2$ ).

**First-order approximation justification.** In this section, we study the first-order approximation of the failure rate from (2.27) (which we call the *first-order heuristic method*) and estimate the prefactor constants  $C_0, C^*, C_1$ , and  $C^{**}$  based on least-squares fits from simulation data (which we call the *first-order empirical method*). To estimate the prefactor constants, we first transform the simulated failure rate data to the “prefactor” space by removing the contribution of the analytic exponential factor. Since the analysis that results in the zeroth-order rate (2.18) is valid in the small-noise limit, we fit  $\hat{C}_0$  from numerical simulation data only for temperatures such that  $\tau \ll \Delta\mathcal{H}$ . We define this regime as  $5\tau < \Delta\mathcal{H}$  and estimate

$$(2.29) \quad \hat{C}_0 = \arg \min_c E_{C_0}, \quad E_{C_0} := \left\| \lambda_{sim}^\tau \times \exp \left[ \frac{\Delta\mathcal{H}}{\tau} \right] - C^* c \tau^{-1/2} \right\|_2^2.$$

Across the seven lines studied in our experiments, we find that our estimate of  $\hat{C}_0$  is within an order of magnitude of  $C_0(\bar{x})$  (see Figure 2.4).

Similarly, we estimate  $\hat{C}_1$  from simulation data in the large-noise limit in the regime where  $\tau \gg \Delta\mathcal{H}$ . We define this regime as  $\tau > 5\Delta\mathcal{H}$  and estimate

$$(2.30) \quad \hat{C}_1 = \arg \min_c E_{C_1}, \quad E_{C_1} := \left\| \lambda_{sim}^\tau \times \exp \left[ \frac{\Delta\mathcal{H}}{\tau} \right] - (C^* \hat{C}_0 + C^* \Delta\mathcal{H}^{-1} c \tau) \tau^{-1/2} \right\|_2^2$$

after estimating  $\hat{C}_0$  from (2.29). Across all lines studied in our experiments (excluding the slack-load line for which we did not have high-temperature observations), we find agreement to within an order of magnitude indicating that  $\hat{C}_1(\bar{x}) \approx C_0(\bar{x})$  and that our  $C^{**}$  approximation is reasonable. Across the range of lines and temperatures studied, we note that the benefit from the empirical method is marginal—less than 0.2 logarithmic units (see Table 2.1) when compared with the heuristic method. We also studied the inclusion of a fractional power  $\tau^{1/2}$  in the prefactor estimate of (2.25), but empirical estimation did not meaningfully improve the fit over the first-order method, so we did not pursue higher fractional orders (3/2, 5/2, ...). Because of this and the convenience of not requiring any additional parameter estimation, we use (2.27) with prefactor estimates from (2.28) in the remainder of this document.

**Summary, impact, and limitations.** Our numerical experiments provide evidence in support of the applicability of large deviation theory to the problem of unconditional line failure. Specifically, our simulations validate the low-noise behavior predicted by both zeroth- and first-order theories for nonpathological lines. Outside the low-noise regime, we find that the first-order approximation introduces additional robustness across a range of temperatures, and we observe an aggregate mean error of around 0.7 logarithmic units across all lines and temperatures studied. Since regime classification is line dependent, this robustness is particularly important in “relatively” high-temperature cases where a line’s energy difference may be smaller than the system’s ambient temperature. In these cases, the  $\lambda \propto \tau^{1/2}$  behavior of the first-order theory generally coincides with our simulations, although in a single case (line 37) the failure rate grew faster. For practical uses, the large-noise limit is less of a concern, however; if a line is expected to fail in the near future, whether it fails in the next second or the next fraction of a second is of less importance.

We acknowledge a number of limitations in our framework. First, the deterministic power transmission model that we employ in (2.1) was developed for local stability

TABLE 2.1

Mean absolute logarithmic error ( $\bar{\delta} = N^{-1} \sum_{i=1}^N |\log(\lambda_i / \log \hat{\lambda}_i)|$ ) across  $N$  temperatures for simulated rate  $\lambda$  and analytic rate  $\hat{\lambda}$ ) [46] and standard deviation of absolute logarithmic error ( $sd$ ) across all studied lines over all temperatures. The “Aggregate” row is computed from the combined dataset.

Line			0th order		1st-order empirical		1st-order heuristic	
Number	Type	Risk	$\bar{\delta}$	( $sd$ )	$\bar{\delta}$	( $sd$ )	$\bar{\delta}$	( $sd$ )
2 <sup>a</sup>	SL-PQ	low	0.82	(0.49)	0.29	(0.19)	0.64	(0.45)
4	PQ-PQ	low	1.20	(1.00)	0.31	(0.20)	0.49	(0.32)
10	PQ-PQ	high	4.18	(3.37)	0.33	(0.20)	0.47	(0.31)
16	PV-PQ	low	2.37	(1.44)	0.91	(0.82)	1.19	(0.94)
20	PQ-PQ	med	2.59	(2.25)	0.51	(0.40)	0.70	(0.31)
29	PV-PQ	high	4.32	(3.43)	0.46	(0.31)	0.64	(0.30)
35	PV-PQ	high	4.58	(3.64)	0.50	(0.35)	0.55	(0.35)
37	PV-PQ	med	2.76	(2.52)	0.91	(0.58)	0.86	(0.56)
Aggregate	—	—	2.88	(2.87)	0.53	(0.49)	0.69	(0.53)

<sup>a</sup>No “high-temperature” observations for line 2;  $C_1$  taken as (2.28). Note also that line 2 is pathological (Definition 3.2).

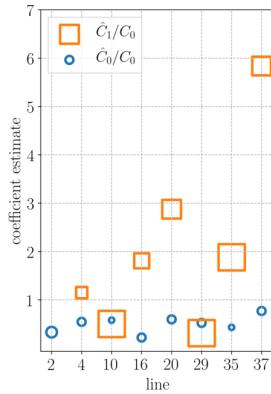


FIG. 2.4.  $\hat{C}_0/C_0$  and  $\hat{C}_1/C_0$  relative estimates from data. Marker size corresponds to the number of data points used in the estimation.

analysis and hence is typically only valid near an equilibrium point; we assume that the stochastic forcing is small enough that the dynamics are applicable.<sup>10</sup> Next, we were limited in exploring low-temperature (and, to a lesser extent, high-temperature) regimes because of the computational cost. In addition, aside from potential modeling concerns such as appropriate system settings and network structure,<sup>11</sup> the failure rate is intimately connected to the system temperature, which may be difficult to estimate.<sup>12</sup>

<sup>10</sup>A rough estimate is that the model may be valid for  $\pm 5\%$  deviations of the nominal operating frequency, but a more detailed analysis, which depends both on the applicability of the power flow equations and the dimension of the generator dynamics, is beyond the scope of this paper. We note, however, that we do not model automatic controls such as generator droop control, which might be necessary to accurately model larger voltage magnitude deviations ( $> 10\%$ ).

<sup>11</sup>The IEEE 30-bus system is an approximation of a portion of the American Electrical Power System from the 1960s.

<sup>12</sup>One approach, proposed in [30], is to utilize the approximately ergodic properties of the system (2.7) to compute the mean magnitude of power grid frequency fluctuations over a long period of time.

Most notably, however, the simulation methodology discussed above handles the tractable *unconditional* failure problem, but the *conditional* failure problem is of interest in practice. Since conditional line  $l$  failures must avoid certain regions of the state space (by definition, regions where other lines fail), computing the conditional failure rate requires solving a different set of Fokker–Planck equations [38] to evaluate the corresponding conditional exit rate. Solving the appropriate Fokker–Planck equation may be nontrivial, and for the remainder of this document we treat the unconditional failure rate as an approximation of the conditional failure rate. In the next section, we present numerical evidence in support of the claim that the approximation is reasonable in the low-noise limit.

**3. Approximating the conditional failure rate.** At this time, we do not have a general approach for modeling and validating conditional failure rates. In fact, the numerical simulation of conditional failure experiments for a single line could be computationally prohibitive if the line of interest has a failure rate significantly lower than the top few failure rates, in which case failures of the desired line become nearly impossible to observe. For the 30-bus model, we have observed that conditional simulation results align with the unconditional failure rate for the line with the largest failure rate at temperatures studied in Figure 2.3, but at finite temperatures, the approximation may become problematic for lines without significant separation between failure rates. Motivated by the engineered structure of the power grid, and to avoid this computational issue while extending our study to the majority of lines, we instead study the validity of the unconditional failure rate as an approximation for the conditional failure rate in the low-noise limit from the computationally tractable optimization perspective.

Since power systems are sparsely connected and designed to prevent simultaneous line failures, state variables that contribute to line  $l$ 's energy often do not significantly influence the line energies of other lines. Accordingly, we anticipate that the regions of state space that correspond to line  $l$  failure are in some sense isolated from regions of state space that may trigger other failures, especially at low loadings (see [18]). Based on these observations, we make the following hypothesis.

*Hypothesis 3.1* (conditional failure approximation). The conditional failure rate can be well approximated by the unconditional failure rate in the low-noise limit.

In the remainder of this section, we study the implications of this hypothesis and its validity by analyzing both failure points and failure paths via optimization.

**3.1. Nested failure point analysis.** A direct indicator that the unconditional failure problem will be a poor approximation of the conditional failure problem is the presence of multiple failures associated with an unconditional failure point (2.16). For a line  $l$  of interest, such a scenario indicates that the lowest energy failure point  $x^*$  is contained within the failure region of another line. The lowest energy failure point such that *only* the line of interest is tripped can be computed by the “conditional” NLP formulated as

$$(3.1a) \quad \min_x \quad \mathcal{H}(x)$$

$$(3.1b) \quad \text{s.t.} \quad \Theta_l(x) = \Theta_l^{\max},$$

$$(3.1c) \quad \Theta_k(x) \leq \Theta_k^{\max} - \epsilon \quad \forall k \in \mathcal{A},$$

where  $\mathcal{A}$  is the set of acceptable lines (those which are still active and exclude line  $l$ ) and  $\epsilon$  is a small positive parameter that enforces strict/loose conformity. When

(3.1) is infeasible, the line of interest will never fail by itself, meaning that the line's failure region is nested entirely within the region of another line and that the MFPT theory is not applicable. When (3.1) is feasible, however, the unconditional problem provides a low-noise asymptotic upper bound on the failure rate. We summarize the scenario as follows.

**DEFINITION 3.2** (nested pathology). *A transmission line has a conditionally nested failure region if its conditional failure point problem (3.1) is infeasible. Similarly, a transmission line is unconditionally nested if the solution to its unconditional failure point problem (2.16) results in line failures besides the line of interest.*

To quantify the prevalence and impact of nestedness, we proceed in two steps. First, we determine prevalence by computing the proportion of infeasible conditional NLPs across all system lines (for conditional nestedness) and the proportion of NLP solutions (2.16) that trigger multiple lines (for unconditional nestedness), as summarized in Figure 3.1(a). We scan for nestedness in both the IEEE 30-bus and 118-bus systems [37], which we augment with physical parameters from Figure 2.2(c).<sup>13</sup> Second, we use (2.27) to compare failure rates computed at both the conditional and unconditional failure point solutions, as summarized in Figure 3.1(b)–(c).

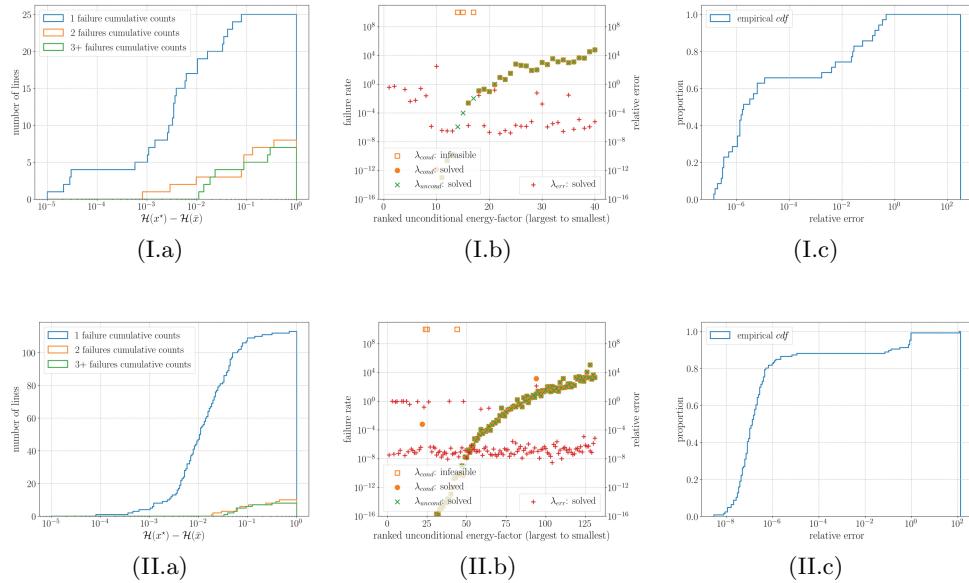


FIG. 3.1. Nested failure diagnostics for 30-bus (I) and 118-bus (II) IEEE systems. (a) Empirical cumulative counts of total number of line failures associated with line  $l$  failure as a function of (unconditional) energy difference  $\Delta\mathcal{H}$ . One line failure means that line  $l$  is the only line to fail. (b) Unconditional failure rate computed at unconditional failure rate and relative error (versus unconditional failure rate computed at conditional failure point) ordered from largest  $\Delta\mathcal{H}$  to smallest  $\Delta\mathcal{H}$ . (c) Empirical cumulative distribution function of relative error.

In each problem instance, we observe that the majority of the system's lines are neither conditionally nor unconditionally nested. Furthermore, we find (i) that nestedness becomes rarer in the larger system, and (ii) that more vulnerable (lower  $\Delta\mathcal{H}$ ) lines have a lower rate of nestedness than the overall system rate. Three of forty

<sup>13</sup>The 118-bus case file does not contain line limits, so we set limits to ensure a standard measure of system security ( $N-1$  security, defined and studied in [31, 2]).

lines in the 30-bus system exhibit conditional NLP infeasibility, but only three of 131 lines studied in the 118-bus system do (see Figures 3.1(I.b) and 3.1(II.b)). Similarly, fifteen lines are unconditionally nested in the 30-bus system, but only eighteen lines are in the 118-bus system (see Figures 3.1(I.a) and 3.1(II.a)). Additionally, we observe that the set of conditionally nested lines is a subset of the set of unconditionally nested lines, indicating that unconditional nestedness is a useful proxy for conditional nestedness. As expected for nonnested lines, panels (b)–(c) of Figure 3.1 show that the relative error between failure rates (2.27) computed from the conditional failure point and the unconditional failure point is small (below  $10^{-5}$ ). For unconditionally nested lines, larger errors are observed, indicating that such lines cannot be studied by the theory in subsection 2.4. Across all lines, about 75% of the lines in the 30-bus case have a relative error less than 1%, and about 85% of the lines in the 118-bus case have relative error less than  $10^{-5}$ . We find these results encouraging, not only because nestedness is uncommon, but also because the low energy difference lines that are the limiting factors for system reliability exhibit lower rates of nestedness. For larger systems, we expect that nestedness will become even less prevalent, and we use the unconditional failure point (2.16) when computing individual failure rates.

**3.2. Nonisolated failure point analysis.** When considering failure boundaries in isolation of each other, the constrained energy minimization problems (2.16) and (3.1) may have multiple local solutions, since they are nonconvex, nonlinear programs. Solutions that are isolated from each other are straightforward to diagnose; in the limit of low noise, the minimal energy solution will dominate. Conversely, lines whose constraint intersects a level set of the energy function cannot be handled by the asymptotic theory. We collect such cases in the following definition.

**DEFINITION 3.3** (nonisolated pathology). *A line  $l$  is pathological in the nonisolated sense if solutions to the unconditional failure problem form a continuum in state space.*

We study isolatedness by solving the unconditional NLP from 100 randomly initialized seeds corresponding to perturbed solutions of the initial dispatch point, in addition to one seed at  $\bar{x}$ . For those problems that converge, we round the failure point  $x^*$  to four digits of accuracy and obtain a lower bound for the number of isolated solutions. We compute the failure rate at each isolated solution and compare the rates with unconditional failure rate (that is, the solution to the unconditional NLP initialized at the operating point  $\bar{x}$ ).

In the vast majority of lines studied numerically, the presence of multiple solutions does not pose a significant problem to failure rate estimation (see Figure 3.2), and hence lines with nonisolated solutions are in some sense “pathological.” In Figure 3.2 we find good agreement between the maximum failure rate (across isolated solutions) and the  $\bar{x}$ -initialized failure rate across both test systems, especially for lines with higher unconditional failure rates. In the larger system, we observe that increased NLP complexity contributed to an increased prevalence of local minima. In the 30-bus system, across all non-*generator-generator* lines, 78% of NLP solutions find a single failure point, 20% find two, and 2% (line 13, a unique case of an isolated load connected to the network through another load) find a continuum of failure points. In the 118-bus system, across all non-*generator-generator* lines, 44% of NLP solutions find a single failure point, 40% find two, 8% find three, and 7% find four. Ultimately, solutions initialized at  $\bar{x}$  well approximate the lowest energy minimizer obtained from multiple restarts in both test cases.

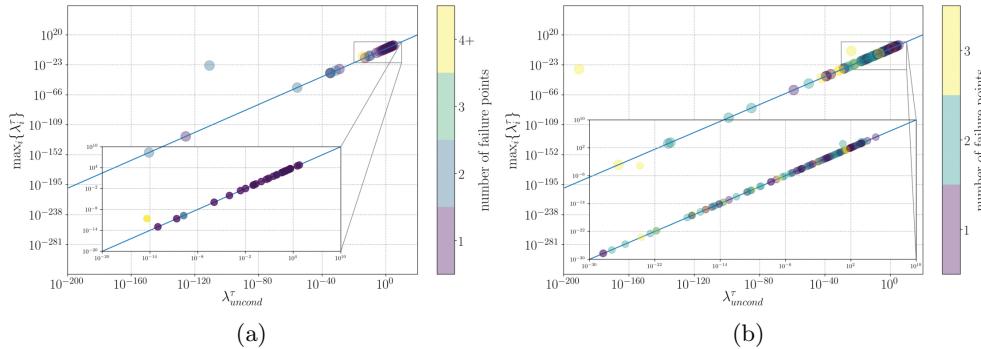


FIG. 3.2. Study of multiplicity of solutions to (2.16). First-order failure rates (2.27) computed at  $\tau = 10^{-3}$ ; the  $y$ -axis shows the maximum first-order failure rate computed across solutions to the randomly initialized unconditional NLP (2.16), and the  $x$ -axis shows the first-order failure rate computed from the solution to (2.16) initialized at  $\bar{x}$ . The blue line indicates perfect agreement. (a) 30-bus system. (b) 118-bus system.

We remark that violation of the applicability conditions (A.i)–(A.iii) can be detected first by inspection (that is, by line type) and second by solving the unconditional NLP. Optimality conditions require  $\nabla \mathcal{H}(x^*) = k \nabla \Theta_l(x)$ , and if  $k < 0$ , we can conclude that conditions (A.i) and (A.ii) are not satisfied at  $x^*$ . In practice, we have observed numerical satisfaction of the applicability criteria near  $x^*$ , although this is not consistent along the entire boundary  $\partial D_l$  of the safety region (2.11).

**3.3. Failure path analysis.** Another indicator that the unconditional problem will poorly approximate the conditional problem is if the most likely exit path from the equilibrium to the constrained energy minimizer passes through a region of state space which triggers another line failure. We collect such instances in the following definition.

**DEFINITION 3.4** (dynamically inaccessible pathology). *A line  $l$  is pathological in the (dynamically) inaccessible sense if the most likely exit path triggers line relays for lines  $k \neq l$ .*

We study inaccessibility by computing the most likely exit path of the dynamics through the boundary  $\partial D_l$ . For the power system dynamics (2.7), the most likely exit path solves the backwards problem (see section 2.1.1 in [11] and Chapter 4, Theorem 3.1 in [20])

$$(3.2) \quad \dot{\phi}_t^x = S \nabla \mathcal{H}(x) + J \nabla \mathcal{H}(x), \quad \phi_0^x = x^*, \quad t \leq 0.$$

By numerically solving (3.2) backwards in time, we compute failure paths for each line in both the IEEE 30-bus and 118-bus networks and find that the majority do not trigger other line failures. In particular, we remark that uniformly increasing the line limits generally separates failure pathways in state space, as shown in Figure 3.3. Similarly, although not directly comparable because of network differences, a comparison of failure pathways under  $N-1$  secure line limits in both 30-bus and 118-bus systems indicates that well-designed limits also increasingly tend to separate in higher dimensions. In summary, the most likely exit path analysis provides evidence in favor of using the unconditional failure rate as an approximation of the conditional failure rate in realistic models at low temperatures.

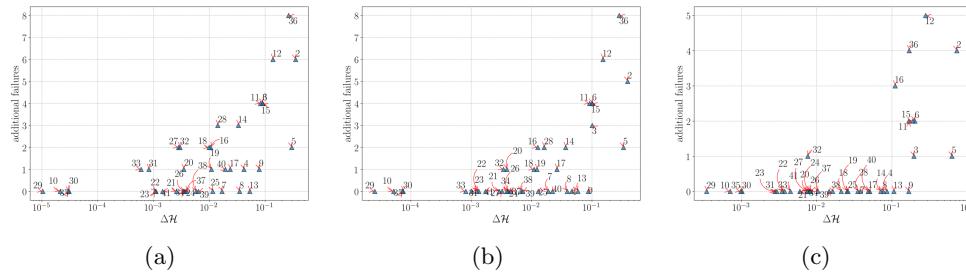


FIG. 3.3. Study of fluctuation dynamics paths displaying number of unintended failures along the pathway between  $x^*$  and  $\bar{x}$  for each line in the 30-bus system (see the supplemental materials for more detail). (a) Standard (Matpower) line limits. (b) Line limits uniformly increased by 10%. (c) Line limits uniformly increased by 100%.

**4. Chemical kinetics Markov model.** In this section, we introduce a chemical kinetics-inspired Markov model of cascading line failure under specified operating conditions. Specifically, we interpret the unconditional failure rate (2.27) as a chemical “reaction rate” in the sense of transition state theory (see [3] for a recent overview) and aggregate individual rates into a composite model derived from kinetic Monte Carlo (KMC, also known as dynamic Monte Carlo) methods. KMC methods were developed to discretize a system’s continuous dynamical behavior into a sequence of distinct events representing transitions between discrete states that occur at known rates [49]. An attractive feature of these models is that if the rates corresponding to mutually exclusive events are both correct and independent of the preceding states, then the event-based system evolves at the same timescale as the original system [49, 41].

Our model discretizes the dynamics of (2.7) by employing the first-order failure rate expression (2.27) to catalog transition rates such that each event corresponds to (i) an individual line failure, and (ii) an irreversible transition to an increasingly degraded operating point. By storing an on-the-fly event list and rate catalog in a weighted directed acyclic graph (DAG), we can generate realizations of possible failure pathways by traversing the DAG rather than by integrating the dynamics of (2.7). The event-based approach can generate sequences of transmission line failures with computational effort that is independent of the temperature  $\tau$ . As a result, the method allows for the study of cascades in settings that may otherwise be intractable because of the prohibitively high computational cost of direct numerical integration at low temperatures. Below, we detail the model structure and simulation methodology (subsection 4.1), its assumptions (subsection 4.2), and our numerical experiments (subsection 4.3).

**4.1. DAG failure pathways and Markov simulation.** The organizing principle of the DAG is to enumerate all possible failure paths between the fully operational topological network state (equivalently, *energy surface*)  $S_0$  (where no lines have failed) and the fully degraded topological network state  $S_\infty$  (where all lines have failed). In this context, each directed edge hierarchically links two energy surfaces,  $S$  and  $S'$ , such that  $S' \setminus S$  contains the set of failed lines linking  $S$  to  $S'$  and its weight is specified by the transition rate from  $S$  to  $S'$ . For simplicity, we disregard simultaneous line failures so that edges connect only topological states that are separated by a single line failure.

At state  $S$ , under certain assumptions described in subsection 4.2, the failure rate to any child state  $\lambda_{S \rightarrow S'}$  can be computed directly from the information contained in

the parent state  $S$  and hence defines a Markov process. Specifically, the set of line failures in  $S$  (in conjunction with the original dispatch data) defines an energy surface with equilibrium minimizer  $\bar{x}_S$ , and the computation of failure point  $x_{S \rightarrow S'}^*$  proceeds as before (2.16). Since the failure rate model depends only on the combination of failed lines in a topological state  $S$ , the full DAG  $\mathcal{G} = (\mathcal{V}, \mathcal{D})$  composed of vertices  $\mathcal{V}$  and directed edges  $\mathcal{D}$  (depicted in Appendix E) can be expressed with a number of vertices that scale exponentially rather than combinatorially. Once all such topological states have been enumerated, failure sequences can be generated by traversing the graph according to the relative weights of each edge.

In order to enumerate the DAG, three inputs are required: (i) an initial network topology, (ii) dispatch data (which defines an initial operating point), and (iii) a set of exogenous dynamics parameters (including temperature  $\tau$  and physical parameters like damping and mass). The first input determines each vertex in  $\mathcal{V}$  as well as the hierarchical interconnections of each edge in  $\mathcal{D}$ , while the second two inputs determine the weights of each edge. In this way, the entire DAG stems from the original dispatch point. For nontrivial networks, enumerating every possible vertex in the graph a priori is infeasible, so we employ a depth-first search of the failure pathway space to enumerate “likely” sequences of failures, as summarized in Algorithm E.1.

**4.2. Markov model hypotheses and evidence.** The efficient exploration of the combinatorial state space of likely failure pathways depends on identifying susceptible lines at each vertex. We use the unconditional failure rate model (2.27) as a proxy for susceptibility along each directed edge.

To compute directed edge weights for the Markov model using transition state theory, we rely on the notion of “equilibration” for the stochastic dynamics (2.7).

**DEFINITION 4.1** (equilibration). *A stochastic process has equilibrated if its law has converged to an appropriate quasistationary distribution.*

Following a line failure in the continuous power grid dynamics (2.7), the system has equilibrated according to Definition 4.1 if (after updating the susceptance matrix  $B$ ), an ensemble of trajectories from the parent state has relaxed to a quasistationary distribution over the child state. In certain cases, however, the process may fail to equilibrate or may fail to do so in an amount of time that is negligible relative to the time of the next failure. The theory in section 2 does not handle the computation of failure rates for processes that have not equilibrated, so we propose the following hypothesis.

**Hypothesis 4.2** (rarity of nonequilibration). Instances of nonequilibration are uncommon when modeling cascading line failure.

Researchers have shown that there exist deterministic failure pathways between cascade topologies [53], but we assume that such cases are rare. Hypothesis 4.2 allows for the formulation of conditional and unconditional failure rate problems for individual lines. Though the conditional failure problem gives the correct line failure rates at each level, we rely on Hypothesis 3.1 and its encouraging validation results to justify computing edge weights via the unconditional failure rate (2.27).

Next we study the reasonableness of Hypothesis 4.2. Symptoms of equilibration can be observed numerically since it is known that the distribution of first exit times for a Markov process that has equilibrated is approximately exponential [12].<sup>14</sup> While

---

<sup>14</sup>In fact, a positive exponential moment of the failure time distribution is a necessary condition for the existence of a quasistationary distribution [12].

we cannot definitively verify equilibration via simulation, we can identify lack of equilibration. To explore equilibration, we return to studying the dynamics (2.7) in the 30-bus system. Specifically, we initialize an ensemble of trajectories at points of failure in degraded networks (according to failure sequences generated by the Markov model), randomize each trajectory's momentum variables (angular frequency), and simulate dynamics until the first overall line failure. Equilibration experiments in Figure 4.1 are performed on the longest sequence of failures generated by the Markov model and carried out at decreasing temperatures until computation time exceeds 2 hours. Convergence in failure time distribution toward an exponential distribution is presented in Figure 4.1.

In Figure 4.1, we summarize convergence toward an exponential distribution (or lack thereof) in two ways, first visually and then quantitatively. Panels (a)–(c) visually demonstrate an example of lack of equilibration across decreasing temperatures (each at the same topology), indicating the presence of deterministic failure pathways as studied in [53]. Similarly, panels (d)–(f) visually demonstrate an example of successful equilibration across decreasing temperatures (each at the same topology). Together, panels (a)–(f) offer a visual key for interpreting the quantitative summary of failure time distribution convergence summarized in panel (g). In particular, panel (g) displays mean absolute deviation  $\bar{\delta} := \frac{1}{N} \sum_{i=1}^N |\log(\hat{q}_i/q_i)|$  across (i) temperature, and (ii) number of initially failed lines (i.e., failure-“index”), where  $\hat{q}_i$  denotes the  $i$ th quantile of the empirically observed failure time distribution and  $q_i$  denotes the  $i$ th quantile of the exponential distribution fitted to the empirical data.

Along this particular line failure sequence, we find that in topologies with fewer line failures (i.e., a lower failure-index), the distribution of first failure times tends to approach a Dirac delta function in the zero noise limit, thus indicating lack of equilibration. In fact, line failures in Figure 4.1(c) correspond to a deterministic failure pathway associated with a single dynamically inaccessible line (line 32). Furthermore, the failure to equilibrate in the child state is a likely indication that the child state's equilibrium may not be genuinely metastable (that is, the relaxation time exceeds the first passage time). Accordingly, when failure times do not follow an approximately exponential distribution, we believe that the reason is the presence of dynamically inaccessible lines that dominate the most likely equilibration pathways. At higher-indexed topologies (those with a greater number of line failures), however, we observe that the system generally exhibits equilibration in the low-noise limit. This indicates that the higher-indexed topologies are relatively stable and that the dynamically inaccessible lines have previously failed. At moderately higher temperatures, we observe approximate equilibration in the majority of cases, even at topologies with multiple dynamically inaccessible lines. In such cases, noise appears to stabilize the equilibration pathway (see Figure 4.1(a)), even at unstable operating points that contain dynamically inaccessible lines.

In summary, results in Figure 4.1(g) indicate the existence of both “stable” and “unstable” equilibration pathways at various levels of degradation (i.e., failure-index), but this relationship is moderated by temperature such that there appears to be an intermediate regime ( $10^{-2}$  to  $10^{-4}$ ) where approximate equilibration is the predominant response. These results provide encouraging evidence in support of the claim that at moderate temperatures, equilibration is not a strong assumption, even for unstable operating points.

**4.3. Experimental results.** In this section, we explore the practical implications of the Markov model. Validating the model against “real” data is difficult since

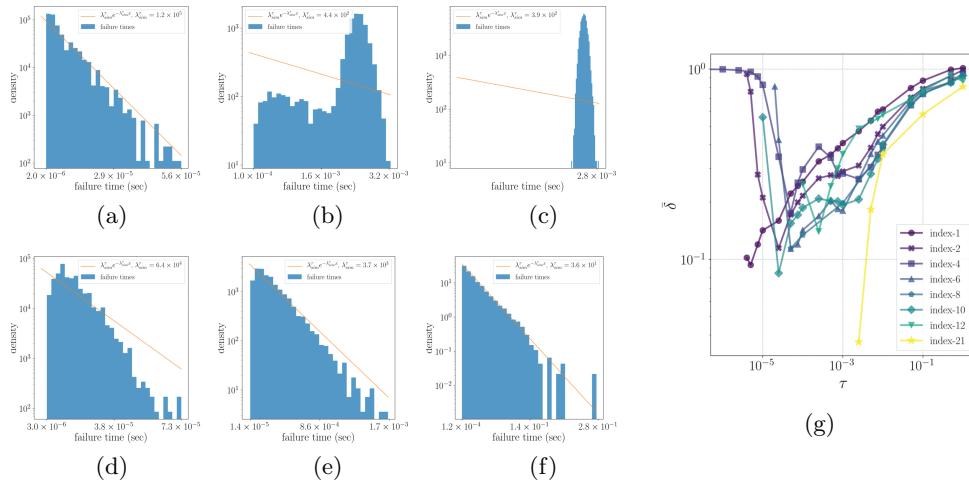


FIG. 4.1. *Equilibration diagnostics.* Panels (a)–(f) show observed and (estimated) exponential failure time distributions. (a)–(c) 4 failures at  $\tau = 10^{-2}$ ,  $10^{-5}$ ,  $10^{-6}$ . (d)–(f) 21 failures at  $\tau = 10^{-1}$ ,  $10^{-2}$ ,  $2.5 \times 10^{-3}$ . (g)  $\delta$  measures mean error between observed and (estimated) exponential failure time distribution across different temperatures and levels of network degradation (i.e., “index”).

cascades are rare events and line failure data is often private. One point of reference, however, is the timestamped line status data provided by the Bonneville Power Administration (BPA [9], part of the Western Interconnect with thousands of buses), which has been analyzed by Dobson and collaborators [17, 16].

The dataset contains times of line outage and reinstatement events, and following [16], cascades can be constructed by appropriately grouping outages that occur in close succession to one another. Specifically, Dobson separates successive outage events by event time according to two bandwidths: one that determines failures that are part of the same *cascade* (failures occurring within 60 minutes of each other), and one that determines failures that are part of the same *generation* (failures occurring within 1 minute of each other within the same cascade) [39].

His results indicate that the distribution of the number of cascades containing  $g$  generations is approximately governed by a power law with exponent  $-s$  [17] and that the number of generations is a useful measure of cascade severity [56]. Furthermore,  $s$  (the negative of the Zipf distribution slope) can be interpreted as a measurement of the system’s risk of cascading failure, called SEPSI (system event propagation slope index).

From observed BPA outage frequency data, we estimate SEPSI by maximum likelihood based on the first nine generations. In Figure 4.2(a), Dobson computes SEPSI of 3.02 for similar data, and the choice of nine generations best aligns with the SEPSI calculated in [17]. Dobson finds that SEPSI ranges from 2.2 to 3.2 under various grid operating conditions [17].

SEPSI provides a tractable metric for comparing observed failure data with simulated failure data, and we study the Markov model’s capability of producing “realistic” cascades through estimates of SEPSI. To generate raw failure data for cascade-processing, we repeat 100 independent experiments of Algorithm E.1 in the IEEE 118-bus system. Each repetition generates a sequence of timestamped line failures which continues until either the entire network collapses or the cumulative sequence

time exceeds  $10^{10}$  seconds.<sup>15</sup> We then aggregate failures to count the frequencies of cascades and generations under 1-hour and 1-minute bandwidth assumptions [17, 16], assuming that the timescale of the dynamics is in seconds. In this way, each repetition can generate multiple sets of cascades from topologies that are related to the original topology through the appropriate number of line failures. Hence it also allows for exploring network stability after likely configurations of initial lines have failed.

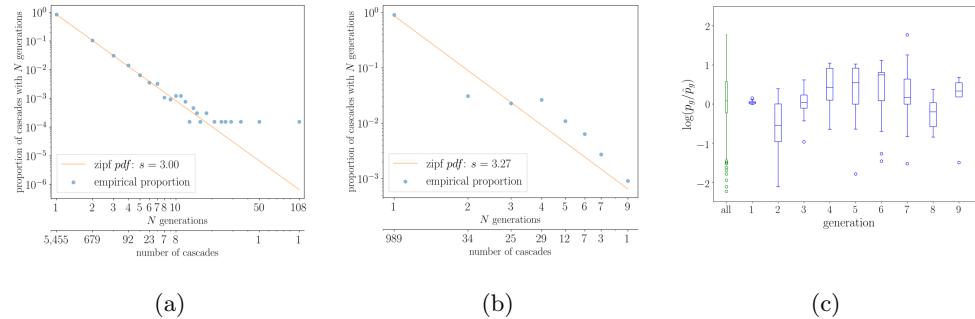


FIG. 4.2. *Empirical failure proportion versus Zipf fit (estimated from the first nine generations).* (a) SEPSI computed from 6,563 cascades from BPA outage data. (b) SEPSI computed from 1,100 cascades from the Markov model at  $\tau = 10^{-3}$ . (c) Aggregated Zipf residuals from Markov model at twelve temperatures between  $10^{-4}$ – $10^{-2}$ .

Findings in Figure 4.2 indicate that the Markov model is capable of reproducing sequences of cascading line failures with approximate power law behavior. The quality of the power law relationship is moderated by network configuration, system temperature, and the cascade-splitting mechanism, but across a range of parameter configurations, we have observed variations in SEPSI that fall within the realistic range.

Across a range of system temperatures, we find that observations exhibit oscillatory behavior around the power law (Figure 4.2(c)). An analysis of logarithmic residuals (defined as  $\log(p_g/\hat{p}_g)$ , where  $p_g$  is the empirical proportion of cascades with  $g$  generations and  $\hat{p}_g$  is the Zipf prediction) indicates that the Zipf fit overestimates the number of cascades with two generations (though we do not expect this amount to be of practical significance in realistic situations). The moderate over- and underestimation varies with generation and may be attributable to the network's size. In larger networks, we expect our preliminary results to stabilize.

We next discuss some numerical properties of the Markov model. First, the computational effort of generating cascading failure sequences is significantly more stable than that of direct numerical integration. At higher temperatures, direct integration may outperform the Markov model, but at temperatures where numerical integration may be prohibitively expensive, the computational cost of the Markov model is fixed and involves only solving a sequence of NLPs. Such examples may occur in intermediate stages of the cascade where the system settles into a relatively stable configuration. Second, solution of individual NLP problems to determine the set of outgoing failure rates at each failure level can be parallelized and therefore should have positive scaling properties for larger networks. Third, once the majority of dominant Markov pathways have been computed, cascade generation amounts to scanning a lookup table and generating exponential random variables.

<sup>15</sup>We choose a long stopping time ( $10^{10}$  seconds) to introduce variability in the original network topology to better reflect the heterogeneous conditions across which the BPA data were collected.

**Limitations.** The Markov model can reproduce approximate power law behavior for cascades with  $N$  generations under certain parameter settings, but it is limited by its complex dependence on tunable parameters, including system temperature and cascade-generation timescale. The choices of system temperature and cascade-generation bandwidths affect the intracascade failure time properties of the Markov model, and it is not clear how best to determine these without numerical experimentation. The faithfulness with which the Markov simulator can reproduce failure sequences corresponding to numerical integration of the network depends on the validity of the reequilibration assumption (Hypothesis 4.2). Our discussion from subsection 4.2 leads us to believe that this is not a significant effect. We also note that the BPA data is not directly comparable to the failure data generated by the Markov model, since the BPA data allows for line reconnection, whereas the Markov model has no such mechanism. A continuous mechanism for line reinstatement and analytic adjustments to the system's energy function is introduced in [54], but we do not address this extension.

**5. Conclusions.** In this article, we described a transition rate theory approach to modeling both individual and cascading transmission line failure in power systems. We introduced the interpretation of individual line failures as large-deviation events on an energy surface defined by a first-principles model of the power system, and we derived expressions for the asymptotic failure rates. Although the power grid dynamics (2.7) were modeled as an irreversible SDE, their close relation to a gradient system (by virtue of the transverse decomposition of the drift vector field) provided a framework for analytical treatment of the problem of first escape through a dividing surface. In the limit of small noise, validation results indicated that our analytic approximations were well suited for determining the unconditional failure rate of any applicable network line. Further, a first-order expansion of the prefactor suggested a modification that increased the approximation's applicability to higher temperature regimes as well. In the first-order model of line failure (2.27), we observed an average (absolute) logarithmic error of 0.7 across all experiments and inferred that the model is capable of describing the failure mechanism of isolated transmission lines in a restricted network.

Next we studied the applicability of the unconditional failure model in the context of realistic networks, relaxing the assumption of studying line failure as an isolated phenomenon. We described various network pathologies (Definitions 3.2 to 3.4) and their potential effects on the validity of the unconditional failure model as an approximation to the conditional failure model. From numerical experiments, we inferred that the prevalence and impact of these pathologies in realistic networks is relatively limited. In addition, such pathologies are related to network parameters (e.g., line limits) that can be adjusted.

We incorporated individual models of line failure into a finite-temperature, KMC-based Markov model to simulate cascades. The Markov model generates cascades by traversing a directed graph containing edges with weights specified by the individual line failure model. Computation within a cascade sequence can be parallelized, and once dominant pathways have been computed, simulation amounts to finding values in a lookup table (albeit an exponentially large one) and generating exponential random variables. Our numerical experiments show that under certain parameter settings, the Markov model can approximate the power law behavior of empirically observed cascades. Tuning the Markov model parameters is currently based on experimentation (in particular  $\tau$ , making the interpretation of cascade timescales difficult for use in

engineering applications); however, we anticipate that future research will address this issue, as well as the following potential directions:

1. Further explore the relationship between temperature, cascade splitting parameters, and cascade severity for describing cascade behavior.
2. Refine the cascade sampling mechanism by sampling cascades from an appropriate *joint* distribution composed of initiating failure events (prior to KMC) and subsequent failure events (from KMC) to aid in estimating system temperature.<sup>16</sup>
3. Analyze the qualitative features of failure time sequences of cascades.
4. Explore alternative characterizations of cascade severity beyond SEPSI (including the number of outages in a cascade or generation and the amount of load lost in a cascade).
5. Introduce mechanisms for line reconnection.
6. Consider more complex stochastic drivers, including flux-decay dynamics and voltage excitation perturbations per [13] to extend the applicability of the failure rate analysis to *generator-generator* lines.
7. Relax the lossless power flow assumption.

**Appendix A. Transverse decomposition of the port-Hamiltonian dynamics.** Bouchet and Reygner [11] note that a stochastic differential equation of the form

$$(A.1) \quad dx_t^\tau = -K \nabla U(x_t^\tau) dt + \sqrt{2\tau} \sigma dW_t$$

for potential  $U$  and any appropriately sized (constant) matrix  $K$  with symmetric part  $\frac{1}{2}(K + K^\top) = \sigma\sigma^\top$  has two salient properties:

1. Its generator preserves stationarity of the Gibbs measure.
2. It possesses a natural transverse decomposition of the drift.

Denoting symmetric and antisymmetric components by  $S = \frac{1}{2}(K + K^\top)$  and  $J = \frac{1}{2}(K - K^\top)$ , respectively, we can express our model (2.7) in the form of (A.1). Specifically, we can perform a transverse decomposition of the port-Hamiltonian drift term into

$$(A.2) \quad b(x) := (J - S)\nabla\mathcal{H}(x) = -S\nabla\mathcal{H}(x) + \ell(x),$$

with  $\ell(x) := J\nabla\mathcal{H}(x)$ . Since  $J$  is skew-symmetric, we have that  $\langle \nabla\mathcal{H}(x), \ell(x) \rangle = 0$ . By the same property, we also have  $\text{div}(\ell(x)) = \text{div}(J\nabla\mathcal{H}(x)) = \text{tr}(J\nabla^2\mathcal{H})(x) = 0$  (using the cycle property of the trace operator and symmetry of the Hessian). Hence, despite being irreversible, the dynamics (2.7) preserve the Gibbs measure and do not introduce any “non-Gibbsianness” into estimates involving the stationary density.

**Appendix B. Derivation of the quasistationary exit rate.** The surface integral in (2.20) can be evaluated by using the Laplace method. Specifically, we are interested in evaluating the asymptotic leading term of integrals of the form

$$\int_{x \in \partial D_l} v^\top(x) n(x) e^{-f(x)/\tau} dS(x), \quad \tau \rightarrow 0,$$

where the minima of  $f(x)$  along the surface  $\partial D_l$  is unique and is not a critical point of  $f(x)$ . In the limit  $\tau \rightarrow 0$ , the integral is dominated by the contribution due to  $e^{-f(x)/\tau}$  around  $x^* = \arg \min_{x \in \partial D_l} f(x)$ .

---

<sup>16</sup>Based on discussion with Ian Dobson.

To evaluate the leading term, we parametrically represent  $\partial D_l \equiv \{x: \Theta_l(x) = \Theta_l^{\max}\}$  in the vicinity of  $x^*$  by

$$x = x(u), \quad u = (u_1, \dots, u_{d-1})^\top \in \partial D_l,$$

where  $u$  denotes orthogonal curvilinear coordinates so that  $x^* = x(u^*)$ . Formula (5.20) of [52, Chapter 9] provides the leading term, which reads

$$(B.1) \quad \int v^\top(x) n(x) e^{-f(x)/\tau} dS(x) \sim (v^\top n)(x^*) D(u^*) (\det \mathcal{B})^{-1/2} (2\pi\tau)^{(d-1)/2} e^{-f(x^*)/\tau},$$

where

$$\mathcal{B} = \text{Hess}_u f(x(u^*)), \quad D(u) = \sqrt{\det \mathcal{J}^\top(u) \mathcal{J}(u)}, \quad \mathcal{J}_{ij}(u) = \frac{\partial x_i}{\partial u_j}, \quad \mathcal{J} \in \mathbb{R}^{d \times (d-1)}.$$

To write  $\det \mathcal{B}$  in terms of  $f(x)$  and  $\Theta_l(x)$ , we employ formula (3.24) of [52], which reads

$$(B.2) \quad \det \mathcal{B} = \frac{D^2(u^*)}{|\nabla f(x^*)|^2} B^*,$$

where

$$(B.3) \quad B^* = \sum_{p=1}^d \sum_{q=1}^d \frac{\partial f(x^*)}{\partial x_p} \frac{\partial f(x^*)}{\partial x_q} \text{cof}_{pq} [\text{Hess } \mathcal{H}(x^*) - k \text{Hess } \Theta_l(x^*)],$$

and  $k$  is a constant such that  $\nabla f(x^*) = k \nabla \Theta_l(x^*)$ . In the NLPs (2.16) and (3.1),  $k$  corresponds to the Lagrange multiplier on the line energy constraint.

We now evaluate these expressions for the stochastic port-Hamiltonian model. First, we have the equivalencies  $f(x) = V(\bar{x}, x) = \mathcal{H}(x) - \mathcal{H}(\bar{x})$  and  $v(x) = (J + S)\nabla \mathcal{H}(x)$ . Second, we have

$$\begin{aligned} (v^\top n)(x^*) D(u^*) (\det \mathcal{B})^{-1/2} &= (v^\top n) D(u^*) \frac{|\nabla f(x^*)|}{D(u^*) \sqrt{B^*}} = \frac{(v^\top \nabla f)(x^*)}{\sqrt{B^*}} \\ &= \frac{1}{\sqrt{B^*}} \nabla^\top \mathcal{H}(x^*) (J + S)^\top \nabla \mathcal{H}(x^*) \\ &= \frac{1}{\sqrt{B^*}} \nabla^\top \mathcal{H}(x^*) S \nabla \mathcal{H}(x^*), \end{aligned}$$

because  $n = \nabla f(x^*)/|\nabla f(x^*)|$ , given that  $f(x^*)$  over  $\partial D_l$  has its minimum on  $x^*$ , and  $\nabla^\top \mathcal{H}(x^*) J \nabla^\top \mathcal{H}(x^*) = 0$ . Substituting this expression above into (B.1) leads to (2.18).

We also note that  $\nabla \Theta_l(x)$  points outwards of  $D_l$ . Furthermore, given that  $f(x)$  over  $\partial D$  is minimal at  $x^*$ ,  $\nabla f(x)$  is parallel and in the direction of  $n(x)$ . Therefore,  $\nabla f(x^*)$  and  $\nabla \Theta_l(x^*)$  point in the same direction, so that  $k = |\nabla f(x^*)|/|\nabla \Theta_l(x^*)| > 0$ . Substituting into (B.2) results in (2.19).

**Appendix C. Higher-order expansion of the exponent and the prefactor of the stationary density.** We use the same notation as the zero-order analysis for the stationary density that can be found in [11]. For the first order of the expansion in the exponent (using up to  $j = 1$  in (2.23)), we assume

$$(C.1) \quad p_{st}^\tau(x) = \frac{C_{st}^\tau(x)}{\tau^{1/2}} \exp \left[ -\frac{\mathcal{H}(x)}{\tau} + \mathcal{H}_1(x) \right],$$

where  $\mathcal{H}_1(x)$  is an unknown function of  $x$ . We plug (C.1) for the density into the stationary Fokker–Planck equation

$$(C.2) \quad \tau \sum_{i,j=1}^d \partial_{ij} [S_{ij}(x)p_{st}^\tau(x)] = \sum_{i=1}^d \partial_i [b(x)p_{st}^\tau(x)],$$

where  $b(x)$  is given by (A.2). We now collect terms in orders of  $\tau$ . We follow section 3 in [11] and use the properties of the power grid model that (i) the matrix  $S$  is constant, (ii)  $F(x) := \operatorname{div} J \nabla \mathcal{H}(x) + \langle A(x), S \nabla \mathcal{H}(x) \rangle = 0$ , and (iii)  $A_i(x) = \sum_{j=1}^d \partial_{ij} S_{ij}(x)$ . The equation of order  $\tau^{-1}$  (*eikonal*) is the Hamilton–Jacobi equation for the quasipotential. Using the eikonal equation and the aforementioned properties, the equation of order  $\tau^0$  gives

$$(C.3) \quad \langle \nabla(\log C_{st}^\tau(x) + \mathcal{H}_1(x)), S \nabla \mathcal{H}(x) + J \nabla \mathcal{H}(x) \rangle = 0.$$

Through the method of characteristics we obtain

$$(C.4) \quad C_{st}^\tau(x) = C_{st}^\tau(\bar{x}) \exp[-\mathcal{H}_1(x) + \mathcal{H}_1(\bar{x})]$$

and for the stationary density

$$(C.5) \quad p_{st}^\tau(x) = \frac{C_{st}^\tau(\bar{x})}{\tau^{1/2}} \exp\left[-\frac{\mathcal{H}(x)}{\tau} + \mathcal{H}_1(\bar{x})\right].$$

Thus, the contribution from the first-order term  $\mathcal{H}_1(x)$  in the exponent is a constant that can be absorbed in the prefactor. We note that we would arrive at the same result even if  $F(x) \neq 0$ . However, starting from the second order and onwards, the condition  $F(x) = 0$  is required in order to show the contribution of only constants from the higher-order terms in the exponent.

Using a similar analysis (albeit with more involved algebra), one can show that the second-order approximation (using up to  $j = 2$  in (2.23))

$$(C.6) \quad p_{st}^\tau(x) = \frac{C_{st}^\tau(x)}{\tau^{1/2}} \exp\left[-\frac{\mathcal{H}(x)}{\tau} + \mathcal{H}_1(x) + \tau \mathcal{H}_2(x)\right]$$

becomes

$$(C.7) \quad p_{st}^\tau(x) = \frac{C_{st}^\tau(\bar{x})}{\tau^{1/2}} \exp\left[-\frac{\mathcal{H}(x)}{\tau} + \mathcal{H}_1(\bar{x}) + \tau \mathcal{H}_2(\bar{x})\right],$$

and thus the contributions of the first- and second-order terms can be absorbed in the prefactor. Similarly, one can show that all the higher-order terms in the exponent end up contributing constants. We omit the details.

For the expansion of the prefactor we use the expression (2.24). The analysis to show that all terms contribute a constant (in  $x$ ) is performed in a similar way as for the expansion of the exponent. For example, for the first order ( $n = 1$  in (2.24)) we have

$$(C.8) \quad p_{st}^\tau(x) = \frac{C_0(x) + \tau C_1(x)}{\tau^{1/2}} \exp\left[-\frac{\mathcal{H}(x)}{\tau}\right].$$

We plug (C.8) into (C.2) and collect terms in orders of  $\tau$ . From the equations at orders  $\tau^{-1}$ ,  $\tau^0$ , and  $\tau^1$  and using properties (i) and (ii) above for the power grid model, we find  $C_0(x) = C_0(\bar{x})$  and

$$(C.9) \quad \langle \nabla C_1(x), b(x) + 2S \nabla \mathcal{H}(x) \rangle = \langle \nabla C_1(x), S \nabla \mathcal{H}(x) + J \nabla \mathcal{H}(x) \rangle = 0.$$

Using the method of characteristics, we get  $C_1(x) = C_1(\bar{x})$ , and so

$$(C.10) \quad p_{st}^\tau(x) = \frac{C_0(\bar{x}) + \tau C_1(\bar{x})}{\tau^{1/2}} \exp\left[-\frac{\mathcal{H}(x)}{\tau}\right].$$

The analysis for the higher-order terms in the prefactor expansion proceeds in a similar way. We omit the details.

**Appendix D. Failure rate dynamics algorithm.** The procedure for unconditional failure rate simulation is summarized below.

---

**Algorithm D.1** Line  $l$  unconditional failure.

---

**Require:** Equilibrium point  $\bar{x}$ , system parameters  $y$ , temperature  $\tau$ , step size  $\Delta t$ , and line limits  $\Theta_l^{\max}$   
 Construct  $J$  and  $S$  per (2.3)  
 $x \leftarrow \bar{x}$   
 $\Theta_l \leftarrow \Theta_l(x)$   
 $t \leftarrow 0$   
**while**  $\Theta_l < \Theta_l^{\max}$  **do**  
 $x \leftarrow x + (J - S) \cdot \nabla_x \mathcal{H}^y(x) \Delta t + \sqrt{2\tau} \sqrt{S} \Delta W$   
 $t \leftarrow t + \Delta t$   
 $\Theta_l \leftarrow \Theta_l(x)$   
**end while**

---

**Appendix E. Markov model.** The rejectionless, on-the-fly Markov simulation procedure is described below.

---

**Algorithm E.1** Rejectionless (Markov) KMC Model.

---

**Require:** Network line data  $\mathcal{N}_L$ , network state  $\mathcal{N}_S$ , equilibrium point  $\bar{x}$ , system parameters  $y$ , temperature  $\tau$ , and stopping time  $t_{\max}$   
 Initialize topological network state:  $S \leftarrow \{\emptyset\}$   
 Initialize DAG:  $\mathcal{G} \leftarrow (\emptyset, \emptyset)$   
 Initialize time, state, and unfailed lines:  $S \leftarrow \mathcal{N}_S$ ,  $t \leftarrow 0$ ,  $U \leftarrow \{\mathcal{N}_L\}$   
**while**  $t \leq t_{\max}$  or  $S \neq S_\infty$  **do** ▷ For full blackout state  $S_\infty$   
 Compute  $\bar{x}_S$  (more detail in supplemental material) and update  $U$  for any failed lines  
**for**  $l \in U$  **do**  
 Look up or compute  $x_{S \rightarrow \{S \cup l\}}^*$  (more detail in supplemental material)  
 Compute  $\lambda_{S \rightarrow \{S \cup l\}}^\tau$  from  $\bar{x}_S$  and  $x_{S \rightarrow \{S \cup l\}}^*$   
 $\mathcal{V} \leftarrow \mathcal{V} \cup \{S \cup l\}$   
 $\mathcal{D} \leftarrow \mathcal{D} \cup \lambda_{S \rightarrow \{S \cup l\}}^\tau$   
**end for**  
 Compute aggregate rate:  $\lambda_{S \rightarrow S'} = \sum_l \lambda_{S \rightarrow \{S \cup l\}}$   
 Sample failure time: draw  $\Delta t \sim \text{Exp}(\lambda_{S \rightarrow S'})$   
 Sample failure line: draw  $l'$  proportionally from aggregate rate  $\lambda_{S \rightarrow S'}$   
 $t \leftarrow t + \Delta t$ ;  $S \leftarrow S' \equiv \{S \cup l'\}$ ;  $U \leftarrow U \setminus l'$   
**end while**

---

The entire discrete event space for a four-line model is enumerated in Figure E.1.

Each edge of the Markov model relates the rate of failure from an equilibrium point  $\bar{x}$  to a failure of a given line  $l$  through a failure point  $x^*$ . In degraded surfaces, the computation of  $\bar{x}$  may result in (i) a nonphysical solution (for example, one with low voltage at a given bus, bus- $i$ ), or (ii) direct infeasibility. We encounter both cases in the Markov model and adjust the solution procedure by shedding load [33]; see the supplemental material for further details.

**Acknowledgments.** We thank Christopher DeMarco, Ian Dobson, Charles Matthews, Adrian Maldonado, and Anirudh Subramanyam for insightful discussions.

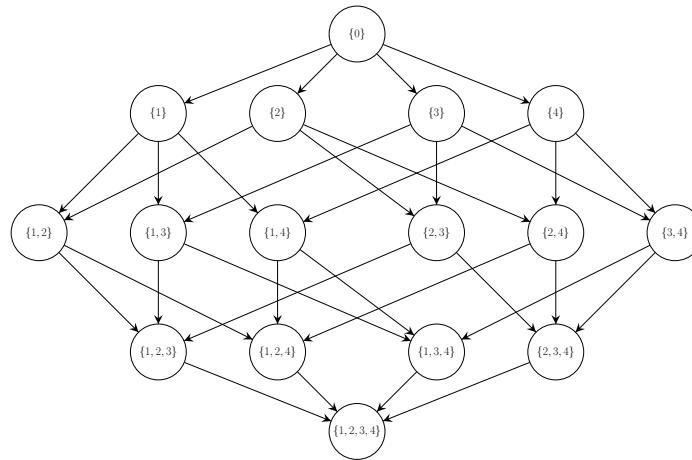


FIG. E.1. (Unweighted) Hasse diagram of all possible failure states for a four-line (labeled 1, 2, 3, and 4) model. Each node contains a topological state  $S$ , and the (directed) edges represent transitions to further degraded surfaces  $S'$  (from left to right). Each directed edge represents an additional line failure  $l$  (where  $l = S' \setminus S$ ) and is weighted by  $\lambda_S^{S'}$ , the analytic transition rate between  $S$  and  $S'$ .

## REFERENCES

- [1] 30bus Matpower Case File, <https://matpower.org/docs/ref/matpower5.0/case30.html> (accessed: 2020-09-28).
- [2] O. ALSAC AND B. STOTT, *Optimal load flow with steady-state security*, IEEE Trans. Power Apparatus Syst., 93 (1974), pp. 745–751.
- [3] J. L. BAO AND D. G. TRUHLAR, *Variational transition state theory: Theoretical framework and recent developments*, Chem. Soc. Rev., 46 (2017), pp. 7548–7596.
- [4] C. M. BENDER AND S. A. ORSZAG, *Advanced Mathematical Methods for Scientists and Engineers I: Asymptotic Methods and Perturbation Theory*, Springer Science & Business Media, 2013.
- [5] A. R. BERGEN AND C. L. DEMARCO, *Application of Singular Perturbation Techniques to Power System Transient Stability Analysis*, Tech. Rep. UCB/ERL-M84/7, Electronics Research Laboratory, UC Berkeley College of Engineering, 1984.
- [6] A. R. BERGEN AND D. J. HILL, *A structure preserving model for power system stability analysis*, IEEE Trans. Power Apparatus Syst., 100 (1981), pp. 25–35.
- [7] A. R. BERGEN AND V. VITTAL, *Power Systems Analysis*, 2nd ed., Pearson, 2000.
- [8] S. BOLOGNANI AND F. DÖRFLER, *Fast power system analysis via implicit linearization of the power flow manifold*, in 2015 53rd Annual Allerton Conference on Communication, Control, and Computing, IEEE, 2015, pp. 402–409.
- [9] BONNEVILLE POWER ADMINISTRATION, *BPA Outage and Reliability Reports*, <https://transmission.bpa.gov/Business/Operations/Outages/> (accessed: 2019-09-04).
- [10] F. BOUCHET, K. GAWEDZKI, AND C. NARDINI, *Perturbative calculation of quasi-potential in non-equilibrium diffusions: A mean-field example*, J. Statist. Phys., 163 (2016), pp. 1157–1210.
- [11] F. BOUCHET AND J. REYGNER, *Generalisation of the Eyring-Kramers transition rate formula to irreversible diffusion processes*, Ann. Henri Poincaré, 17 (2016), pp. 3499–3532.
- [12] P. COLLET, S. MARTÍNEZ, AND J. SAN MARTÍN, *Quasi-Stationary Distributions: General Results*, Springer, 2013, pp. 17–29.
- [13] C. DEMARCO AND A. BERGEN, *A security measure for random load disturbances in nonlinear power system models*, IEEE Trans. Circuits Syst., 34 (1987), pp. 1546–1557.
- [14] C. L. DEMARCO, *A new method of constructing Lyapunov functions for power systems*, in 1988 IEEE International Symposium on Circuits and Systems (ISCAS), Vol. 1, 1988, pp. 905–908.
- [15] C. L. DEMARCO AND A. R. BERGEN, *Small noise in power system dynamics and its relation to security assessment*, in 1985 American Control Conference, IEEE, 1985, pp. 559–560.

- [16] I. DOBSON, *Estimating the propagation and extent of cascading line outages from utility data with a branching process*, IEEE Trans. Power Syst., 27 (2012), pp. 2146–2155.
- [17] I. DOBSON, *Finding a Zipf Distribution and Cascading Propagation Metric in Utility Line Outage Data*, preprint, <https://arxiv.org/abs/1808.08434>, 2018.
- [18] I. DOBSON, B. A. CARRERAS, V. E. LYNCH, AND D. E. NEWMAN, *Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization*, Chaos, 17 (2007), 026103.
- [19] S. FRANK AND S. REBENNACK, *A Primer on Optimal Power Flow: Theory, Formulation, and Practical Examples*, Working Papers 2012-14, Division of Economics and Business, Colorado School of Mines, 2012.
- [20] M. I. FREIDLIN AND A. D. WENTZELL, *Random Perturbations of Dynamical Systems*, 3rd ed., Grundlehren Math. Wiss. 260, Springer-Verlag, 2012.
- [21] D. J. HIGHAM, *An algorithmic introduction to numerical simulation of stochastic differential equations*, SIAM Rev., 43 (2001), pp. 525–546, <https://doi.org/10.1137/S0036144500378302>.
- [22] P. HINES, E. COTILLA-SANCHEZ, AND S. BLUMSACK, *Do topological models provide good information about electricity infrastructure vulnerability?*, Chaos, 20 (2010), 033122.
- [23] P. K. HOTA AND A. P. NAIK, *Analytical review of power flow tracing in deregulated power system*, American J. Electrical Electronic Engrg., 4 (2016), pp. 92–101.
- [24] C.-R. HWANG, S.-Y. HWANG-MA, AND S.-J. SHEU, *Accelerating Gaussian diffusions*, Ann. Appl. Probab., 3 (1993), pp. 897–913.
- [25] P. KUNDUR, *Power System Stability and Control*, McGraw-Hill, 1994.
- [26] H. J. KUSHNER, *Approximation and Weak Convergence Methods for Random Processes, with Applications to Stochastic Systems Theory*, MIT Press, 1984.
- [27] B. LEIMKUHLER AND C. MATTHEWS, *Numerical Methods for Stochastic Molecular Dynamics*, Springer, 2015, pp. 261–328.
- [28] T. LELIÉVRE AND G. STOLTZ, *Partial differential equations and stochastic methods in molecular dynamics*, Acta Numer., 25 (2016), pp. 681–880.
- [29] P. A. LIPKA, C. CAMPAIGNE, M. PIRNIA, R. P. O'NEILL, AND S. S. OREN, *Constructing transmission line current constraints for the IEEE and Polish systems*, Energy Syst., 8 (2017), pp. 199–216.
- [30] C. MATTHEWS, B. STADIE, J. WEARE, M. ANITESCU, AND C. DEMARCO, *Simulating the Stochastic Dynamics and Cascade Failure of Power Networks*, preprint, <https://arxiv.org/abs/1806.02420>, 2018.
- [31] A. V. MEIER, *Electric Power Systems: A Conceptual Introduction*, John Wiley & Sons, 2006.
- [32] G. MIL'SHTEJN, *Approximate integration of stochastic differential equations*, Theory Probab. Appl., 19 (1975), pp. 557–562.
- [33] K. MOLLAH, M. BAHADORNEJAD, N. K. C. NAIR, AND G. ANCELL, *Automatic under-voltage load shedding: A systematic review*, in 2012 IEEE Power and Energy Society General Meeting, 2012, pp. 1–7.
- [34] N. NARASIMHAMURTHI AND M. MUSAVI, *A generalized energy function for transient stability analysis of power systems*, IEEE Trans. Circuits Syst., 31 (1984), pp. 637–645.
- [35] T. NESTI, J. NAIR, AND B. ZWART, *Temperature overloads in power grids under uncertainty: A large deviations approach*, IEEE Trans. Control Network Syst., 6 (2019), pp. 1161–1173.
- [36] M. PAI, *Energy Function Analysis for Power System Stability*, Springer, 1989.
- [37] R. D. ZIMMERMAN, C.E. MURILLO-SANCHEZ, AND R. J. THOMAS, *MATPOWER: Steady-state operations, planning and analysis tools for power systems research and education*, IEEE Trans. Power Syst., 26 (2011), pp. 12–19.
- [38] S. REDNER, *A Guide to First-Passage Processes*, Cambridge University Press, 2001.
- [39] H. REN AND I. DOBSON, *Using transmission line outage data to estimate cascading failure propagation in an electric power system*, IEEE Trans. Circuits Syst. II Express Briefs, 55 (2008), pp. 927–931.
- [40] M. ROHDEN, D. JUNG, S. TAMRAKAR, AND S. KETTEMANN, *Cascading failures in AC electricity grids*, Phys. Rev. E, 94 (2016), 032209.
- [41] S. A. SEREBRINSKY, *Physical time scale in kinetic Monte Carlo simulations of continuous-time Markov chains*, Phys. Rev. E, 83 (2011), 037701.
- [42] I. SIMONSEN, L. BUZNA, K. PETERS, S. BORNHOLDT, AND D. HELBING, *Transient dynamics increasing network vulnerability to cascading failures*, Phys. Rev. Lett., 100 (2008), 218701.
- [43] F. SLOOTHAAK, S. C. BORST, AND B. ZWART, *Robustness of power-law behavior in cascading line failure models*, Stoch. Models, 34 (2018), pp. 45–72.
- [44] T. W. STEGINK, C. D. PERSIS, AND A. J. V. D. SCHAFT, *An energy-based analysis of reduced-order models of (networked) synchronous machines*, Math. Comput. Model. Dyn. Syst., 25 (2019), pp. 1–39.

- [45] S. TAMRAKAR, M. CONRATH, AND S. KETTEMANN, *Propagation of disturbances in AC electricity grids*, Sci. Rep., 8 (2018), 6459.
- [46] L. TÖRNQVIST, P. VARTIA, AND Y. O. VARTIA, *How should relative changes be measured?*, Amer. Statist., 39 (1985), pp. 43–46.
- [47] H. TOUCHETTE, *Introduction to dynamical large deviations of Markov processes*, Phys. A, 504 (2018), pp. 5–19.
- [48] A. VAN DER SCHAFT AND D. JELTSEMA, *Port-Hamiltonian systems theory: An introductory overview*, Found. Trends Syst. Control, 1 (2014), pp. 173–378.
- [49] A. F. VOTER, *Introduction to the kinetic Monte Carlo method*, in *Radiation Effects in Solids*, K. E. Sickafus, E. A. Kotomin, and B. P. Uberuaga, eds., Springer, 2001, pp. 1–23.
- [50] W. S. WADMAN, D. T. CROMMELIN, AND B. P. ZWART, *A large-deviation-based splitting estimation of power flow reliability*, ACM Trans. Model. Comput. Simul., 26 (2016), pp. 23:1–23:26.
- [51] D. WITTHAUT AND M. TIMME, *Nonlocal effects and countermeasures in cascading failures*, Phys. Rev. E, 92 (2015), 032809.
- [52] R. WONG, *Higher dimensional integrals*, in *Asymptotic Approximations of Integrals*, SIAM, 2001, pp. 477–515, <https://doi.org/10.1137/1.9780898719260.ch9>.
- [53] Y. YANG AND A. E. MOTTER, *Cascading failures as continuous phase-space transitions*, Phys. Rev. Lett., 119 (2017), 248302.
- [54] H. ZHENG AND C. L. DEMARCO, *A bi-stable branch model for energy-based cascading failure analysis in power systems*, in *North American Power Symposium 2010*, IEEE, 2010, pp. 1–7.
- [55] H. ZHENG AND C. L. DEMARCO, *A new dynamic performance model of motor stalling and FIDVR for smart grid monitoring/planning*, IEEE Trans. Smart Grid, 7 (2016), pp. 1989–1996.
- [56] K. ZHOU, I. DOBSON, Z. WANG, A. ROITERSHEIN, AND A. P. GHOSH, *A Markovian Influence Graph Formed from Utility Line Outage Data to Mitigate Cascading*, preprint, <https://arxiv.org/abs/1902.00686>, 2019.