

Line Failure Localization of Power Networks Part II: Cut Set Outages

Linqi Guo^{ID}, Chen Liang^D, *Student Member, IEEE*, Alessandro Zocca^{ID}, Steven H. Low^{ID}, *Fellow, IEEE*, and Adam Wierman

Abstract—Transmission line failure in power systems propagate non-locally, making the control of the resulting outages extremely difficult. In Part II of this paper, we continue the study of line failure localizability in transmission networks and characterize the impact of cut set outages. We establish a Simple Path Criterion, showing that the propagation pattern due to bridge outages, a special case of cut set failures, are fully determined by the positions in the network of the buses that participate in load balancing. We then extend our results to general cut set outages. In contrast to non-cut outages discussed in Part I whose subsequent line failures are contained within the original blocks, cut set outages typically impact the whole network, affecting the power flows on all remaining lines. We corroborate our analytical results in both parts using the IEEE 118-bus test system, in which the failure propagation patterns exhibit a clear block-diagonal structure predicted by our theory, even when using full AC power flow equations.

Index Terms—Cascading failure, Laplacian matrix, contingency analysis, spanning forests.

I. INTRODUCTION

IN PART I of this paper [1] we establish a spectral representation of power redistribution that precisely captures the Kirchhoff's Laws in terms of the distribution of different families of subtrees in the transmission network. This new representation enables us to precisely characterize how non-cut line outages propagate. In particular, a non-cut outage in a block will only impact the branch power flows on the transmission lines within that block, regardless of whether the outage involves a single line or multiple lines simultaneously. Moreover, a non-cut outage will almost surely impact flows on every remaining line within its block.

Contributions of Part II of this paper: We study cut set outages and analytically characterize how such failures impact the

Manuscript received May 27, 2020; revised November 1, 2020 and January 19, 2021; accepted March 1, 2021. Date of publication April 13, 2021; date of current version August 19, 2021. This work was supported by Resnick Fellowship, Linde Institute Research Award, NWO Rubicon under Grant 680.50.1529, and by NSF under Grants CCF 1637598, ECCS 1619352, ECCS 1931662, CNS 1545096, CNS 1518941, CPS ECCS 1739355, and CPS 154471. Paper no. TPWRS-00762-2020. (Corresponding author: Linqi Guo.)

Linqi Guo, Chen Liang, Steven H. Low, and Adam Wierman are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: lguo@caltech.edu; cliang2@caltech.edu; slow@caltech.edu; adamw@caltech.edu).

Alessandro Zocca is with the Department of Mathematics, The Vrije Universiteit Amsterdam, 1081HV Amsterdam, The Netherlands (e-mail: a.zocca@vu.nl).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TPWRS.2021.3068048>.

Digital Object Identifier 10.1109/TPWRS.2021.3068048

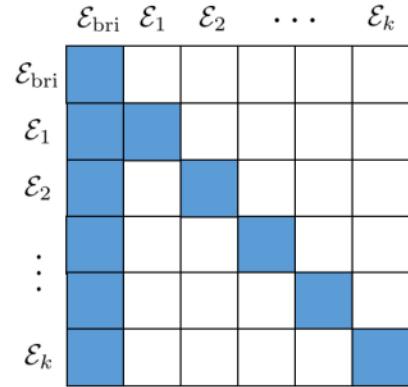


Fig. 1. Non-zero entries of the $K_{\hat{L}}$ matrix (depicted as the dark blocks) for a graph with non-bridge blocks $\{\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_b\}$ and bridge set \mathcal{E}_{bri} .

remaining lines. Our results demonstrate how the impact of cut set outages propagate globally in a way that depends on both the design of power balancing rules and the network topological structure. This characterization, together with our results from Part I, can be visualized in Fig. 1, where it becomes clear how the block decomposition of a network is linked to the sparsity of the LODF matrix $K_{\hat{L}}$. This new theory builds on recent work on line outage distribution factors, e.g., [2], [3], and shows that the block decomposition yields an extremely useful representation of these factors.

The formal characterization of single bridge outage is given by a Simple Path Criterion as Theorem 1 in Section III, which shows that the relative positions of the buses participating in load balancing fully determines how such failures propagate. This result applies to the scenario in which the post-contingency network is disconnected into two or more connected components, known as *islands*, and the original power injections need to be balanced in each island. We then formulate the concept of *participating block* and show that bridge outages typically propagate globally across the network and impact the branch flows on all transmission lines. In Section IV, we extend Theorem 1 to the case of a general cut set outage, and show that the aggregate impact of such failures can be decomposed into two terms: (a) a first term that captures the effect of power redistribution, which can be decomposed in accordance to the blocks where the failures occur and is fully characterized by our study in Part I; and (b) a second term that describes the impact of the power balancing rule and generalizes the case of a single bridge failure, capturing how the system handles disconnected components.

Results in Parts I and II of this paper provide a complete characterization of line failure propagation and is illustrated in Section V using IEEE 118-bus test network. In particular, we show that the LODF matrix has a clear block-diagonal structure predicted by our theory, even when using full AC power flow equations.

II. ISLANDING MODEL

In contrast to Part I where we focus on simultaneous line outages that do not disconnect the network, we consider now the case in which the set of initial line outages disconnects the network into two or more connected sub-networks, called *islands*. We refer to such contingency as a *cut set outage* and, in the special case in which the cut set consists of a single line, as a *bridge outage*. We remark that, in practice, such islands can be created accidentally by line outages, but also deliberately as a defensive action to prevent a disturbance/attack from propagating across the entire network infrastructure [4].

A. Islands and Cut Set Outages

In this section we present a detailed model for a single island that is necessary for our analysis of cut set outages in the later sections. More specifically, we fully characterize the impact of a bridge outage in Section III and of a general cut set outage in Section IV.

We adopt the same notations as in Part I, which we now summarize. Let $\mathcal{G}' = (\mathcal{N}', \mathcal{E}')$ denote the pre-contingency network and consider a subset of lines $F' \subset \mathcal{E}'$ that is a cut set of \mathcal{G}' and denote by $\mathcal{G}_1, \dots, \mathcal{G}_k$ the multiple islands created by the removal of lines in F' . Let us focus on one of these islands, say $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N} \subseteq \mathcal{N}'$ is the set of buses that belong to the island and $\mathcal{E} \subseteq \mathcal{E}'$ is the set of lines that, *pre-contingency*, have both endpoints inside the island. From the viewpoint of the island \mathcal{G} , the lines in the cut set F' can then be classified into three types, $F' = F_{\text{external}} \cup F_{\text{tie}} \cup F$, as follows:

- $F_{\text{external}} := \{\hat{l} = (\hat{i}, \hat{j}) \in F' : \hat{i} \notin \mathcal{N}, \hat{j} \notin \mathcal{N}\}$ is the set of tripped *external lines* with neither endpoints in the island \mathcal{G} ;
- $F_{\text{tie}} := \{\hat{l} = (\hat{i}, \hat{j}) \in F' : \hat{i} \notin \mathcal{N}, \hat{j} \in \mathcal{N}\}$ is the set of tripped *tie lines* with exactly one endpoint (denoted by as $j(\hat{l})$) in the island \mathcal{G} ;
- $F := \{\hat{l} = (\hat{i}, \hat{j}) \in F' : \hat{i} \in \mathcal{N}, \hat{j} \in \mathcal{N}\}$ is the set of tripped *internal lines* with both endpoints inside the island \mathcal{G} .

Note that external line outages do not have a direct impact on the island \mathcal{G} since the post-contingency operations and power flow equations are decoupled in each island and thus can be ignored. Therefore, without loss of generality, we henceforth assume $F' = F_{\text{tie}} \cup F$. Since the graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ describes the pre-contingency topology of the island, its edge set \mathcal{E} includes the tripped internal lines F , but not the tie lines F_{tie} . We refer to the lines that are still active in the island post-contingency as *surviving lines* and denote their collection as $-F := \mathcal{E} \setminus F$. The post-contingency island is thus fully described by the graph $(\mathcal{N}, \mathcal{E} \setminus F)$, which is connected by construction. In particular, the tripped internal lines F is a non-cut set outage of island \mathcal{G} .

Designate any bus in \mathcal{N} to be the reference bus for \mathcal{G} and, without loss of generality, assume it is bus n . Let B denote the susceptance matrix associated with the island \mathcal{G} , C its incidence matrix, $L := CBC^T$ its Laplacian matrix, and define A in terms of the reduced Laplacian matrix \bar{L} as in Part I of this paper:

$$A = \begin{bmatrix} \bar{L}^{-1} & 0 \\ 0 & 0 \end{bmatrix}.$$

Let $f := (f_{\hat{l}}, \hat{l} \in \mathcal{E})$ and $f^{\text{tie}} := (f_{\hat{l}}, \hat{l} \in F_{\text{tie}})$ be the pre-contingency branch flows on the lines inside the island \mathcal{G} and on the tie lines respectively. We adopt the convention that $f_{\hat{l}} > 0$ for a line $\hat{l} = (\hat{i}, \hat{j})$ if power flows from bus \hat{i} to bus \hat{j} over line \hat{l} . In particular, for a tie line $\hat{l} \in F_{\text{tie}}$, $f_{\hat{l}} > 0$ implies that pre-contingency the island imports power over line \hat{l} and $f_{\hat{l}} < 0$ if it exports power over \hat{l} .

If the pre-contingency branch flow $f_{\hat{l}} = 0$ on any tie line or internal line $\hat{l} \in F'$, then its tripping has no impact on the post-contingency branch flows in this island, as modeled by the DC power flow equations. We therefore assume without loss of generality that the pre-contingency branch flows $f_{\hat{l}} \neq 0$ for all tripped lines $\hat{l} \in F'$ (otherwise, remove \hat{l} from F' and the surviving island \mathcal{G}).

B. Pre-Contingency Injections and Branch Flows

Let $p := (p_k, k \in \mathcal{N})$ denote the pre-contingency injections in the buses of island \mathcal{G} . The effect of pre-contingency tie line branch flows on the island \mathcal{G} can be modeled by additional injections $f_{\hat{l}}$ at buses $j(\hat{l})$:

$$\Delta p_{\text{tie}} := \sum_{\hat{l} \in F_{\text{tie}}} f_{\hat{l}} e_{j(\hat{l})}, \quad (1)$$

where $e_{j(\hat{l})}$ is the standard unit vector of size $n := |\mathcal{N}|$. Hence, for the purpose of computing pre-contingency branch flows in island \mathcal{G} , the injections can be taken to be $p + \Delta p_{\text{tie}}$. Let $f_F := (f_l, l \in F)$ denote the pre-contingency branch flows on lines in F , $f_{-F} := (f_l, l \in \mathcal{E} \setminus F)$ those on the surviving lines in \mathcal{G} , and $f := (f_{-F}, f_F)$. Partition the matrices (B, C) according to the two sets of lines, F and $-F = \mathcal{E} \setminus F$, as follows:

$$B =: \begin{bmatrix} B_{-F} & 0 \\ 0 & B_F \end{bmatrix}, \quad C =: \begin{bmatrix} C_{-F} & C_F \end{bmatrix}.$$

From (1) it follows that (f, θ) satisfies the DC power flow equations on the pre-contingency island $\mathcal{G} = (\mathcal{N}, \mathcal{E})$:

$$p + \Delta p_{\text{tie}} = Cf = C_{-F}f_{-F} + C_Ff_F, \quad (2a)$$

$$f = BC^T\theta, \quad (2b)$$

where θ are the pre-contingency voltage angles.

C. Post-Contingency Injections and Branch Flows

The effect of tie line outages F_{tie} on island \mathcal{G} can be modeled as the loss of the injections Δp_{tie} at the endpoints of the tie lines. The pre-contingency injections are then unbalanced over

the island \mathcal{G} with a total imbalance equal to

$$\sum_{k \in \mathcal{N}} p_k = - \sum_{\hat{l} \in F_{\text{tie}}} f_{\hat{l}}.$$

Post contingency, there is a surplus if the island net exports power pre contingency and a shortage otherwise, depending on the sign of this total imbalance. A balancing rule \mathbb{G} is invoked to rebalance power in the island by adjusting injections (generators and/or loads) in response to the contingency. A popular balancing rule, which we name *proportional control*, prescribes how to share the imbalance proportionally among a set of participating buses. More specifically, a proportional control \mathbb{G}_α is defined by a nonnegative vector $\alpha := (\alpha_k, k \in \mathcal{N})$ such that $\sum_{k \in \mathcal{N}} \alpha_k = 1$, with the interpretation that, post contingency, each bus $k \in \mathcal{N}$ adjusts its injection by the amount:

$$\mathbb{G}_\alpha : \Delta p_k := -\alpha_k \sum_{k \in \mathcal{N}} p_k = \alpha_k \sum_{\hat{l} \in F_{\text{tie}}} f_{\hat{l}}, \quad k \in \mathcal{N}. \quad (3)$$

We call a bus k *participating* if $\alpha_k > 0$. By design, all participating buses adjust their injections in the same direction. Examples of proportional control \mathbb{G}_α include participation factors used in automatic generation control or economic dispatch [5, Chapter 3.8], or equal sharing of total imbalance among participating buses [2], [6]; see also [7]–[10]. In this paper we focus solely on proportional control \mathbb{G}_α . For a different class of balancing rules in which the post-contingency injections are determined as a solution of an optimization problem to minimize the number of buses involved, see [11].

Under \mathbb{G}_α , the vector of injection adjustments are then

$$\Delta p_\alpha := (\Delta p_k, k \in \mathcal{N}) := \sum_{\hat{l} \in F_{\text{tie}}} f_{\hat{l}} \sum_{k \in \mathcal{N}} \alpha_k e_k, \quad (4)$$

where e_k is the standard unit vector in \mathbb{R}^n . The post-contingency injections are thus $p + \Delta p_\alpha$. Since $\sum_{k \in \mathcal{N}} \alpha_k = 1$, the identity $\sum_{k \in \mathcal{N}} (p_k + \Delta p_k) = 0$ holds for any initial injection vector p , which means that post-contingency power injections are always rebalanced under the proportional control \mathbb{G}_α .

Let $(\tilde{f}_{-F}, \tilde{\theta})$ denote the post-contingency branch flows and voltage angles, which satisfy the DC power flow equations on the post-contingency network $(\mathcal{N}, \mathcal{E} \setminus F)$:

$$p + \Delta p_\alpha = C_{-F} \tilde{f}_{-F}, \quad (5a)$$

$$\tilde{f}_{-F} = B_{-F} C_{-F}^T \tilde{\theta}. \quad (5b)$$

For post-contingency network, we denote the Laplacian matrix by $L_{-F} := C_{-F} B_{-F} C_{-F}^T$ and define the matrix A_{-F} correspondingly in terms of the reduced Laplacian matrix \bar{L}_{-F} .

In the next two sections we use this island model to analyze line outage localization within the island \mathcal{G} under the proportional control \mathbb{G}_α .

III. BRIDGE OUTAGE

In this section we focus on the case of a single bridge \hat{l} outage, i.e., $F_{\text{tie}} := \{\hat{l}\}$, and no internal line outages, i.e., $F := \emptyset$. In the next section we will then extend the results to a cut set outage where the post-contingency branch flows in island \mathcal{G} are impacted by both internal line outages in $F \subset \mathcal{E}$ and by tie line

outages in F_{tie} . Since F is not a cut set in \mathcal{G} we study this impact by combining the analysis of a non-cut outage in Part I of the paper and that of a bridge outage.

Consider a single bridge \hat{l} outage with pre-contingency branch flow $f_{\hat{l}}$ that disconnects the network into two islands. Focus on one of them, say \mathcal{G} . Let $\hat{j} := j(\hat{l})$ be the endpoint of \hat{l} in island \mathcal{G} . In this case,

$$\Delta p_{\text{tie}} = f_{\hat{l}} e_{\hat{j}}. \quad (6)$$

Post contingency, the injections are changed from $p + \Delta p_{\text{tie}}$ to $p + \Delta p_\alpha$ under the proportional control \mathbb{G}_α as defined in (4). Note that since $F = \emptyset$, we have $-F = \mathcal{E}$, $C_{-F} = C$, $B_{-F} = B$ and $A_{-F} = A$. Therefore, the post-contingency branch flows \tilde{f} in the island \mathcal{G} are given by

$$\begin{aligned} \tilde{f} &= BC^T A (p + \Delta p_\alpha) \\ &= BC^T A (p + \Delta p_{\text{tie}} + \Delta p_\alpha - \Delta p_{\text{tie}}) \\ &= f + BC^T A (\Delta p_\alpha - \Delta p_{\text{tie}}). \end{aligned} \quad (7)$$

Using (4), (6), and the fact that $\sum_{k \in \mathcal{N}} \alpha_k = 1$, we get

$$\begin{aligned} \Delta p_\alpha - \Delta p_{\text{tie}} &= f_{\hat{l}} \sum_{k \in \mathcal{N}} \alpha_k e_k - f_{\hat{l}} e_{\hat{j}} \sum_{k \in \mathcal{N}} \alpha_k \\ &= f_{\hat{l}} \sum_{k \in \mathcal{N}} \alpha_k (e_k - e_{\hat{j}}). \end{aligned} \quad (8)$$

Substituting this expression in (7) gives

$$\tilde{f} - f = f_{\hat{l}} \sum_{k \in \mathcal{N}} \alpha_k BC^T A (e_k - e_{\hat{j}}).$$

Recalling that $f_{\hat{l}} \neq 0$ by assumption, for any line $l = (i, j) \in \mathcal{E}$ we obtain

$$\begin{aligned} \frac{\tilde{f}_l - f_l}{f_{\hat{l}}} &= \sum_{k \in \mathcal{N}} \alpha_k B_l (A_{ik} + A_{j\hat{j}} - A_{i\hat{j}} - A_{jk}) \\ &= \sum_{k: \alpha_k > 0} \alpha_k D_{l,k\hat{j}} \end{aligned}$$

where $D_{l,k\hat{j}}$ is the power transfer distribution factor (PTDF) for island \mathcal{G} discussed in Part I of the paper. Therefore, the branch flow change on line l is the superposition of impacts due to injecting $\alpha_k f_{\hat{l}}$ at participating buses $k \in \mathcal{N}$ and withdrawing them at bus \hat{j} . We can thus extend the definition of line outage distribution factor (LODF) $K_{l\hat{l}}$ to allow bridge outages as follows: given a bridge outage \hat{l} , under the proportional control \mathbb{G}_α for all $l \in \mathcal{E}$ we have

$$K_{l\hat{l}} := \frac{\tilde{f}_l - f_l}{f_{\hat{l}}} = \sum_{k: \alpha_k > 0} \alpha_k D_{l,k\hat{j}}. \quad (9)$$

Remember that $\hat{l} \notin \mathcal{E}$ in the island model that we consider here.

The next result is analogous to the Simple Cycle Criterion (see Theorem 8 in Part I of this paper) for non-bridge outages. It states that all (μ -almost surely), and only, lines on a simple path¹ between bus \hat{j} and a participating bus $k \in \mathcal{N}$ with $\alpha_k > 0$ will be impacted by the bridge \hat{l} outage.

¹A simple path is a path that visits each node at most once.

Theorem 1 (Simple Path Criterion: bridge outage): For a single bridge \hat{l} outage, under the proportional control \mathbb{G}_α , for every line l in the island \mathcal{G} , $K_{\hat{l}l} \neq 0$ “if” and only if there exists a simple path in \mathcal{G} that contains l from \hat{j} to a participating bus $k \in \mathcal{N}$.

See Appendix A for a proof. Note that if \hat{j} (the endpoint of \hat{l} in \mathcal{G}) is the only participating bus, then $K_{\hat{l}l} = 0$ for all $l \in \mathcal{E}$.

Denote by $\mathcal{E}_1, \dots, \mathcal{E}_b$ the unique block decomposition of \mathcal{G} .² We say a block \mathcal{E}_k is on a simple path between bus \hat{j} and a participating bus $i \in \mathcal{N}$ if there is a simple path between bus \hat{j} and bus i with $\alpha_i > 0$ that contains a line $l \in \mathcal{E}_k$. From Theorem 1 we can deduce the following localization property after a single bridge \hat{l} outage in terms of the block structure of the island \mathcal{G} .

Corollary 2 (Bridge outage): Under the proportional control \mathbb{G}_α (4), for any block \mathcal{E}_k of the island \mathcal{G} the following statement hold:

- 1) $\tilde{f}_l = f_l$ for all lines l in block \mathcal{E}_k if \mathcal{E}_k is not on a simple path between bus \hat{j} and a participating bus.
- 2) Conversely, $\tilde{f}_l \neq f_l$ for all lines l in \mathcal{E}_k “if” \mathcal{E}_k is on a simple path between bus \hat{j} and a participating bus.

Corollary 2 shows that the positions of participating buses play an important role in distributing the power imbalance across the network. In particular, the proportional control \mathbb{G}_α almost surely changes the branch flow on every line that lies in a path from the failure endpoint \hat{j} to the set of participating buses. As a result, if l is a bridge connecting two sub-networks \mathcal{G}_1 and \mathcal{G}_2 post contingency (each of which contains one or more blocks), assuming $\hat{j} \in \mathcal{G}_1$, then $\Delta f_l \neq 0$ “if” and only if there is a participating bus in \mathcal{G}_2 since a path from \hat{j} to any node in \mathcal{G}_2 must pass through the bridge l . If l is not a bridge, i.e. l belong to a non-bridge block, then we can devise a simple sufficient condition for $\Delta f_l \neq 0$ using *participating blocks*, defined as follows:

Definition 3: Consider an island \mathcal{G} with block decomposition $\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_b$ operating under proportional control \mathbb{G}_α with a set $\mathcal{N}_\alpha := \{i \in \mathcal{N} : \alpha_i > 0\}$ of participating buses. A non-bridge block is said to be a **participating block** if there is a non-cut vertex in this block that is a participating bus.

If all generators participate in AGC and load-side participation is implemented at all load buses, then every node in the network is a participating bus and, hence, every block is participating.

The following result, whose proof is presented in Appendix B, shows that if a non-bridge block is participating, then all lines inside it are impacted when the original bridge \hat{l} is disconnected.

Corollary 4: Consider a bridge outage \hat{l} with non-zero branch flow $f_{\hat{l}} \neq 0$ and let the block decomposition of island \mathcal{G} be $\mathcal{E}_1, \dots, \mathcal{E}_b$. If \mathcal{E}_k is a participating block, then $\Delta f_l \neq 0$ μ -almost surely for any $l \in \mathcal{E}_k$, i.e., $\mu(\Delta f_l \neq 0) = 1$.

IV. CUT SET OUTAGE

We now extend our results to a cut set outage F' . Consider an island $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ and, as before, partition the tripped lines into tie lines and internal lines, i.e. $F' = F \cup F_{\text{tie}}$.

The impact on post-contingency branch flows is a superposition of the impact of internal line outages in F , weighted

by generalized line outage distribution factor (GLODF) with multiple tripped lines K^F , and the impact of tie line outages in F_{tie} , weighted by the proportional control α_k as well as PTDF of the post-contingency network, as stated in the following theorem.

Theorem 5: Given an island $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with a cut set outage F' , under the proportional control \mathbb{G}_α (4) the branch flow changes on the surviving lines in $-F = \mathcal{E} \setminus F$ are given by

$$\Delta f_{-F} = K^F f_F + \sum_{\hat{l} \in F_{\text{tie}}} f_{\hat{l}} \sum_{k \in \mathcal{N}} \alpha_k \hat{D}^F (e_k - e_{j(\hat{l})}), \quad (10)$$

where $K^F := B_{-F} C_{-F}^T A C_F (I - B_F C_F^T A C_F)^{-1}$ is the GLODF³ of island \mathcal{G} with a non-cut outage F , and \hat{D}^F is the PTDF for the post-contingency network $(\mathcal{N}, \mathcal{E} \setminus F)$ defined as

$$\hat{D}^F := B_{-F} C_{-F} A_{-F}, \quad (11a)$$

which can equivalently be expressed in terms of the pre-contingency island \mathcal{G} as

$$\hat{D}^F = (B_{-F} C_{-F}^T + K^F B_F C_F^T) A. \quad (11b)$$

The theorem reduces to (9) for a single bridge \hat{l} outage, that is when $F = \emptyset$ and $F_{\text{tie}} = \{\hat{l}\}$. When the cut set outage contains both internal line outages in F and tie line outages in F_{tie} , the post-contingency branch flows depend in an intricate way on both types of outages. Our theorem makes this relationship explicit and precise:

- 1) The first term on the right-hand side of (10) represents the impact of the outage of a non-cut set F of internal lines in \mathcal{G} through the GLODF K^F of island \mathcal{G} . If there are no tie line outages $F_{\text{tie}} = \emptyset$, then the formula reduces to the GLODF for a non-cutset outage as discussed in Part I.
- 2) The second term on the right-hand side of (10) represents the impact of the proportional control \mathbb{G}_α in response to tie line outages in F_{tie} . If there are no internal line outages $F = \emptyset$ then the formula reduces to

$$\Delta f_{-F} = \sum_{\hat{l} \in F_{\text{tie}}} f_{\hat{l}} \sum_{k \in \mathcal{N}} \alpha_k \hat{D}^F (e_k - e_{j(\hat{l})}),$$

which generalizes (9) from a single bridge \hat{l} outage to a cut set F_{tie} outage with multiple tripped tie lines. In this case the impact of simultaneous outage of a set F_{tie} of tie lines is simply the sum of the impacts of single-bridge outages as if \hat{l} is a bridge incident on \mathcal{G} . The expressions for \hat{D}^F (11a) and (11b) in terms of pre- and post-contingency network trivially coincide when $F = \emptyset$.

- 3) When both $F \neq \emptyset$ and $F_{\text{tie}} \neq \emptyset$, the effect of their outages on post-contingency branch flows can be interpreted in terms of either the pre-contingency network A or the post-contingency network A_{-F} , through the expressions for both K^F and \hat{D}^F . In particular the expression for \hat{D}^F in terms of the pre-contingency network A has two components. The first component $B_{-F} C_{-F}^T A$ says that the injection adjustments due to the proportional control α_k will change the branch flows on the surviving lines in $-F$ according to PTDF. The second component $K^F B_F C_F^T A$

²See Part I of the paper for more details on block decomposition.

³See Part I of the paper for more details on GLODFs.

says that the injection adjustments are first mapped to flow changes on lines in F through PTDF $B_F C_F^T A(e_k - e_{j(\hat{l})})$ in the *pre-contingency* network and they are then mapped to flow changes on surviving lines in $-F$ through K^F when lines in the non-cut set F are disconnected.

4) According to the expression (11a) for \hat{D}^F , the post-contingency network integrates both effects: changes to post-contingency branch flows are the sum of the impact of injection adjustments under proportional control α_k through PTDF $B_{-F} C_{-F}^T A_{-F}(e_k - e_{j(\hat{l})})$ on the *post-contingency* network.

Considering a surviving line $l \in -F$, the flow change is given by, in terms of pre-contingency island \mathcal{G} :

$$\begin{aligned} \Delta f_l &= \sum_{\hat{l} \in F} K_{l\hat{l}}^F f_{\hat{l}} \\ &+ \sum_{\hat{l} \in F_{\text{tie}}} f_{\hat{l}} \sum_{k \in \mathcal{N}} \alpha_k \left(D_{l,kj(\hat{l})} + \sum_{\tilde{l} \in F} K_{l\tilde{l}}^F D_{\tilde{l},kj(\hat{l})} \right) \\ &= \sum_{\hat{l} \in F} K_{l\hat{l}}^F \left(f_{\hat{l}} + \sum_{\hat{l} \in F_{\text{tie}}} \sum_{k \in \mathcal{N}} \alpha_k D_{\hat{l},kj(\hat{l})} f_{\hat{l}} \right) \\ &+ \sum_{\hat{l} \in F_{\text{tie}}} f_{\hat{l}} \sum_{k \in \mathcal{N}} \alpha_k D_{l,kj(\hat{l})}. \end{aligned} \quad (12)$$

From Theorem 12 in Part I (which covers the case of a non-cut outage) it follows that the first term is zero if l is not in the same block as any disconnected internal line $\hat{l} \in F$ (since $K_{l\hat{l}}^F = 0$). Applying the Simple Path Criterion in Theorem 1 implies that the second term is zero if, for every tie line $\hat{l} \in F_{\text{tie}}$ and every participating bus $k \in \mathcal{N}$ with $\alpha_k > 0$, l does not lie on any simple path in \mathcal{G} connecting $j(\hat{l})$ and k . We can thus derive the following localization property in terms of the block decomposition of \mathcal{G} .

Corollary 6 (Cut set outage): For a cut set F' outage, under the proportional control \mathbb{G}_α (4), for any surviving line $l \in -F$ in the island \mathcal{G} , $\Delta f_l = 0$ if the unique block \mathcal{E}_k containing l satisfies both the following conditions:

- \mathcal{E}_k contains no disconnected internal line $\hat{l} \in F$; and
- For every tie line $\hat{l} \in F_{\text{tie}}$ and every participating bus $i \in \mathcal{N}$ with $\alpha_i > 0$, \mathcal{E}_k is not on a simple path in \mathcal{G} between $j(\hat{l})$ and i .

The converse in general is false because, even when both terms in (12) are nonzero, they may cancel with each other, resulting in $\Delta f_l = 0$.

V. CASE STUDIES

Theorem 7 in Part I and Theorem 5 in Part II summarize the mathematical theory that characterizes the patterns of line failure propagation in power systems. More specifically, the failure localizability depends critically on the block decomposition of a power network. In this section, we demonstrate these localizability properties through simulations using the IEEE 118-bus test network.

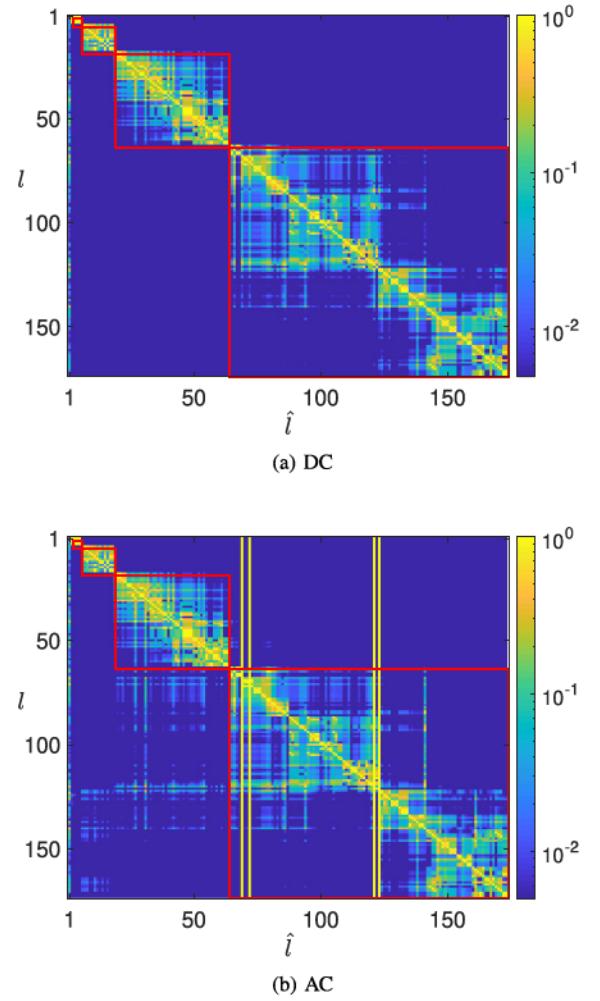


Fig. 2. The LODF matrix (reporting the absolute values of the distribution factors) of IEEE 118-bus network calculated under (a) DC and (b) AC power flow model. The four yellow solid line in (b) represent four transmission lines whose failure lead to non-convergent AC power flow equations. The red rectangles indicate blocks of the network.

A. IEEE 118-Bus Network

This IEEE test network consists of 118 nodes and 186 edges and has a block decomposition with two non-bridge blocks: one giant block with 164 edges, and a smaller one with 13 edges. There are trivial “dangling” appendages each of which connects a single node to the giant block.⁴ For a clearer demonstration, we remove these dangling bridges and replace them by power injections at the corresponding endpoints in the giant block that equal the power flows on these bridges. In addition, we switch off three transmission lines to create more non-bridge blocks to better illustrate the block diagonal structure of the LODF matrix. The resulting network is composed of 4 non-bridge blocks with 4, 13, 45 and 110 edges, connected by one bridge block and two cut vertices.

In [11], [12] we explore ways to judiciously switch off a small number of transmission lines to create more blocks for localizing

⁴It should be mentioned that many detailed models of transmission networks have a meshed core with “dangling” appendages like IEEE 118-bus system.

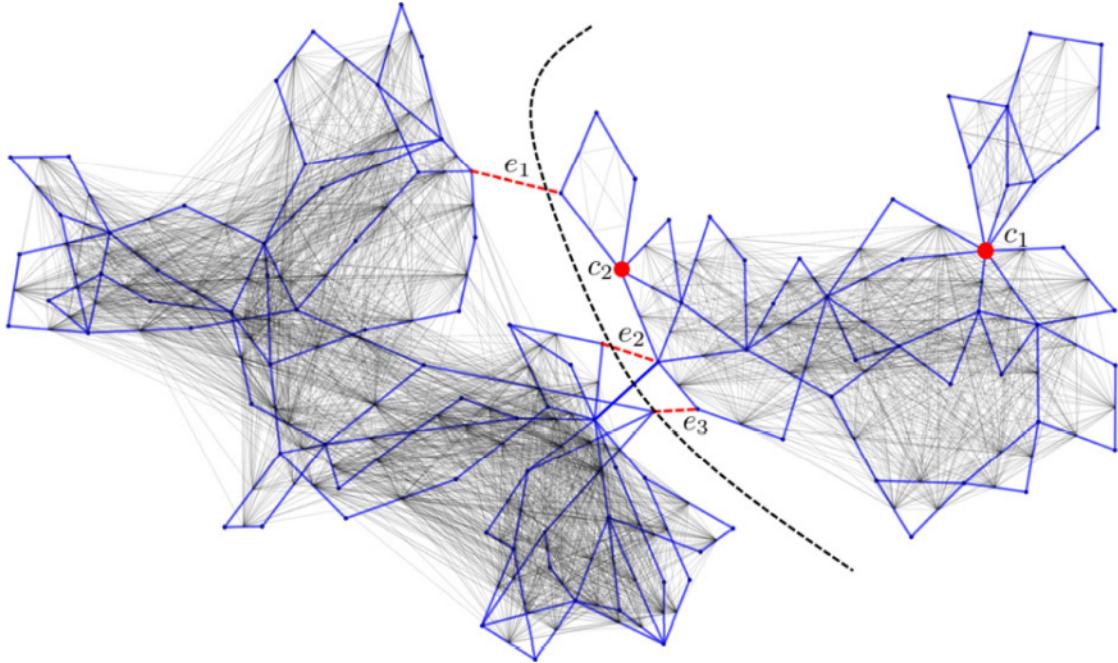


Fig. 3. Influence graphs on the IEEE 118-bus test network after switching off lines e_1 , e_2 , and e_3 . Blue edges represent physical transmission lines and grey edges represent connections in the influence graph. The black dashed line and the red solid points indicate the failure propagation boundary defined by the blocks.

failures. In this paper, however, we focus on characterizing the intrinsic properties of a power network itself rather than the design for network reconfiguration and dynamic controller.

B. LODF Matrix

In our experiments, the system parameters are taken from the Matpower Package [13] and we calculate the LODF matrix $K = (K_{\hat{l}\hat{l}}, \hat{l} \in \mathcal{E}, \hat{l} \in \mathcal{E})$ that describes the impact on other lines of a single line outage. For non-bridge outages, we directly calculate the LODF as defined in equation (4) in Part I of this paper. For bridge outages that create islands, we adopt the proportional control and assume all nodes are participating with $\alpha_k = \alpha$ for all k .

We visualize the magnitude of LODF matrix in Fig. 2(a) by means of a heatmap, in which we reorder the lines based on the block they belong to. Specifically, we stack bridge blocks first, followed by non-bridge blocks in nondecreasing order of the block size. In addition, we set a color limit for better visualization so that $|K_{\hat{l}\hat{l}}| \leq 0.005$ maps to dark blue and $|K_{\hat{l}\hat{l}}| \geq 1$ maps to yellow. In Fig. 2(b), we plot the same heatmap under the AC power flow equations (the four yellow solid line represent four transmission lines whose failure lead to non-convergent AC power flow equations), where the LODF is computed directly from its definition $K_{\hat{l}\hat{l}} = \frac{\Delta f_{\hat{l}}}{f_{\hat{l}}}$.

In both the DC and AC case, the global effect of a bridge failure is clearly visible in the first column (since only one line is a bridge) of the LODF matrices in Fig. 3(a) and (b). Since almost all the entries of that column are non-zero, almost all surviving lines will be impacted by the failure of that bridge. For non-bridge failures, the LODF matrix in the DC case exhibits a clear block diagonal structure. In the AC case, however, the cross-block entries are not necessarily zero, but they are relatively

small. Moreover, the LODF within a block are similar for both cases, indicating that the LODF computed from DC model can be a good approximation for AC model. We further remark that the LODFs within a block can be small, but they are strictly nonzero in all cases, confirming our result in Theorem 7 from Part I.

C. Influence Graph

We further visualize the transmission line failure propagation patterns using an *influence graph*. Despite being similar in concept to that of [14], [15], unlike these influence graphs, which are based on a probabilistic failure models, our influence graph is simply a visualization of the LODFs $K_{\hat{l}\hat{l}}$ that we superimpose on the original network topology. The IEEE 118-bus network topology is depicted in blue in Fig. 3. The corresponding influence graph has these transmission lines (in blue) as nodes and connect any two transmission lines \hat{l} and \hat{l}' in the influence graph (with a grey edge) if the corresponding LODF satisfies $|K_{\hat{l}\hat{l}'}| \geq 0.005$. As Fig. 3 shows, the impact of non-bridge outages are “blocked” by cut-vertices or bridges. Specifically, there are no edges connecting transmission lines that belong to different blocks, as predicted by our theory.

VI. CONCLUSION

In Part II of this work, we make use of the spectral representation of power redistribution developed in Part I to provide a characterization of line failure localizability when the initial failure disconnects the original network. This, together with our results in Part I, establishes a mathematical theory that covers all initial failure scenarios and reveals how a general power system responds to such disturbances. A case study on the IEEE 118-bus test network corroborates the block-diagonal structure predicted

by our theory, even when the system is under full AC power flow equations.

This work can be built upon in several ways: (a) Some practical power systems have been operating in partitioned mode. However, their tie-lines are high voltage DC (HVDC) transmission lines. One major feature of HVDC lines is that they are sensitive to voltage disturbances [16] so outages inside a block can cause HVDC outages, and further triggering outages in other blocks. More study is needed to fully understand line failure localizability for such systems. (b) The network block decomposition has potential application in power systems planning. It can also be used as a corrective action immediately after a contingency, similar to controlled islanding, to prevent failure propagation. It would be interesting to understand the tradeoffs of post-contingency corrective block configuration and controlled islanding and how they can be synergistically integrated for failure mitigation. (c) Our model builds upon DC power flow dynamics, which are accurate for small deviations but less so under large disruptions. Yet, our preliminary simulations suggest a strong underlying structure that connects the gap between DC and AC models in the context of line failures. It would be interesting to understand this structure and develop bounds on the distance between DC and AC predictions. (d) It is possible to integrate fast-timescale frequency control into our framework to provide a control strategy with provably optimal localization and mitigation properties. It would be useful to see what the best way to deploy such a strategy in practical systems would be (see [11], [12] for more details).

APPENDIX: PROOFS

A. Proof of Theorem 1

The proof is similar to that for the Simple Loop Criterion in Part I. Recall from Part I that the PTDF $D_{l,kj}$ at participating buses in (9) is given by:

$$\frac{B_l}{\sum_{H \in \mathcal{T}_E} \beta(H)} \left(\sum_{H \in \mathcal{T}(ik, jj)} \beta(H) - \sum_{H \in \mathcal{T}(ij, jk)} \beta(H) \right)$$

Note that the spanning forests H in \mathcal{T}_E , $\mathcal{T}(ik, jj)$ and $\mathcal{T}(ij, jk)$ are spanning forests of the pre-contingency network, not just the post-contingency island.

Suppose $D_{l,kj} \neq 0$. Then either $\mathcal{T}(ik, jj)$ or $\mathcal{T}(ij, jk)$ or both is nonempty. Suppose without loss of generality that there is a spanning forest $H \in \mathcal{T}(ik, jj)$. Then there is a path in H from buses \hat{j} to j and another vertex-disjoint path from buses i to k . Joining these two paths with the line $l = (i, j)$ creates a simple path from \hat{j} to k that contains line l , as desired.

Conversely suppose there is a simple path in \mathcal{G} from \hat{j} to a participating bus $k' \in \mathcal{N}$ that contains $l = (i, j)$. We will show that $g(B) := \sum_k D_{l,kj}(B) \alpha_k$ in (9) as a polynomial in the susceptances B of all lines in the pre-contingency network is not identically zero, and hence $\mu(g(B + \omega) = 0) = 0$. Since the pre-contingency network is connected the simple path in the island \mathcal{G} from buses \hat{j} to k' that contains $l = (i, j)$ can be extended into a spanning tree of the pre-contingency network. Suppose without loss of generality that, on this spanning tree,

the path from \hat{j} to i contains j , i.e., the path from k' to \hat{j} is of the form $k' \rightsquigarrow i \rightarrow j \rightsquigarrow \hat{j}$. Then removing the edge l from this spanning tree creates a spanning forest H in $\mathcal{T}(ik', jj)$ that contains exactly two trees, denoted by $H_{ik'}$ connecting buses i, k' and H_{jj} connecting buses j, \hat{j} . Following a similar argument in our proof for the Simple Cycle Criterion, for $H \notin \mathcal{T}(ij, jk')$ the term $\beta(H)$ is not canceled by a negative term in $D_{l,kj}$. Moreover we claim that $H \notin \mathcal{T}(ij, jk)$ for all other participating buses k and therefore $\beta(H)$ is not canceled by negative terms from other $D_{l,kj}$ in $g(B)$ either. To see this, if $H \in \mathcal{T}(ij, jk)$ for a participating bus k then there must be a path in H connecting buses i to \hat{j} , but this path then connects tree $H_{ik'}$ to tree H_{jj} into a spanning tree. This contradicts that the spanning forest $H \in \mathcal{T}(ik', jj)$ consists of two distinct trees. Hence $g(B)$ is not a zero polynomial and $K_{l\hat{l}} \neq 0$ μ -almost surely. \square

B. Proof of Lemma 4

Let $\bar{\mathcal{C}}$ be the block that contains l . Since \mathcal{E}_k is a participating region, we know there exists a bus within \mathcal{C} , say n_1 , that participates the power balance and is not a cut vertex. Recall that any path from $j(\hat{l})$ to \mathcal{C} must go through a common cut vertex in \mathcal{C} [17], say n_e . Now by adding an edge between n_e and n_1 (if such edge did not originally exist), the resulting block \mathcal{C}' is still 2-connected. Thus there exists a simple cycle in \mathcal{C}' that contains the edge (n_e, n_1) and $l := (i, j)$, which implies we can find two disjoint paths P_1 and P_2 connecting the endpoints of these two edges. Without loss of generality, assume P_1 connects n_e to i and P_2 connects n_1 to j . By concatenating the path from $j(\hat{l})$ to n_e , we can extend P_1 to a path \tilde{P}_1 from $j(\hat{l})$ to i , which is still disjoint from P_2 . Now, by adjoining \hat{e} to \tilde{P}_1 and P_2 , we can construct a path from $j(\hat{l})$ to n_1 that passes through l . The Simple Path Criterion then implies $\mu(\Delta f_{\hat{e}} \neq 0) = 1$. \square

C. Proof of Theorem 5

The first part of Theorem 5 is proved by analyzing the post-contingency network $(\mathcal{N}, \mathcal{E} \setminus F)$ using the matrix A_{-F} , while the second part leverages the relation between the matrices A for the pre- and post-contingency networks. Similar to our analysis of GLODF, taking the difference between (5b) and (2b) shows that the branch flow changes $\Delta f_{-F} := \tilde{f}_{-F} - f_{-F}$ satisfy the post-contingency DC power flow equations:

$$\begin{aligned} C_F f_F + (\Delta p_{\alpha} - \Delta p_{\text{tie}}) &= C_{-F} \Delta f_{-F}, \\ \Delta f_{-F} &= B_{-F} C_{-F}^T (\tilde{\theta} - \theta). \end{aligned}$$

Recall that A_{-F} is defined in terms of the inverse of the reduced Laplacian matrix $L_{-F} = C_{-F} B_{-F} C_{-F}^T$ for the post-contingency network. The branch flow changes are given by

$$\begin{aligned} \Delta f_{-F} &= B_{-F} C_{-F}^T A_{-F} (C_F f_F + \Delta p_{\alpha} - \Delta p_{\text{tie}}) \\ &= \underbrace{B_{-F} C_{-F}^T A_{-F} C_F}_{K^F} f_F \\ &\quad + \underbrace{B_{-F} C_{-F}^T A_{-F}}_{\hat{D}^F} (\Delta p_{\alpha} - \Delta p_{\text{tie}}), \end{aligned}$$

where the GLODF $K^F = B_{-F} C_{-F}^T A_{-F} C_F$ is defined by Theorem 8 in Part I and \hat{D}^F is defined in (11a). From (4) and (1) it follows that

$$\Delta p_\alpha - \Delta p_{\text{tie}} = \sum_{\hat{i} \in F_{\text{tie}}} f_{\hat{i}} \sum_{k \in \mathcal{N}} \alpha_k (e_k - e_{j(\hat{i})}). \quad (13)$$

Hence the branch flow changes are the sum of the impacts of internal line outages in F and tie line outages in F_{tie} :

$$\Delta f_{-F} = \underbrace{K^F f_F}_{\text{int. line } F \text{ outage}} + \underbrace{\sum_{\hat{i} \in F_{\text{tie}}} f_{\hat{i}} \sum_{k \in \mathcal{N}} \alpha_k \hat{D}^F (e_k - e_{j(\hat{i})})}_{\text{tie line } F_{\text{tie}} \text{ outage}}$$

This proves identity (10) of Theorem 5.

We now show that (11b) is an alternative representation of (11a). To do so, we express the matrix \hat{D}^F in terms of the matrices of the pre-contingency network. In particular, we relate A_{-F} with A using matrix inversion lemma as follows:

$$\begin{aligned} \hat{D}^F &= B_{-F} C_{-F}^T A_{-F} \\ &= B_{-F} C_{-F}^T A \\ &\quad + B_{-F} C_{-F}^T A C_F (I - B_F C_F^T A C_F)^{-1} B_F C_F A \\ &= B_{-F} C_{-F}^T A + K^F B_F C_F A. \end{aligned}$$

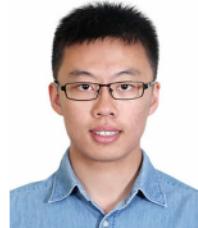
□

REFERENCES

- [1] L. Guo, C. Liang, A. Zocca, S. H. Low, and A. Wierman, "Line Failure Localization of Power Networks Part I: Non-cut outages," *IEEE Trans. Power Syst.*, vol. 36, no. 5, pp. 4140–4151, Sep. 2021, doi: [10.1109/TPWRD.2021.3066336](https://doi.org/10.1109/TPWRD.2021.3066336).
- [2] S. Soltan, D. Mazařic, and G. Zussman, "Analysis of failures in power grids," *IEEE Control Netw. Syst.*, vol. 4, no. 2, pp. 288–300, Jun. 2017.
- [3] L. Guo, C. Liang, and S. H. Low, "Monotonicity properties and spectral characterization of power redistribution in cascading failures," in *Proc. 55th Annu. Allerton Conf.*, 2017, pp. 918–925.
- [4] H. Li, G. Rosenwald, J. Jung, and C.-C. Liu, "Strategic power infrastructure defense," *Proc. IEEE*, vol. 93, no. 5, pp. 918–933, May 2005.
- [5] A. J. Wood, B. F. Wollenberg, and G. B. Sheblé, *Power Generation, Operation, and Control*, 3rd ed. Hoboken, NJ, USA: John Wiley Sons, Inc., 2014.
- [6] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman, "Power grid vulnerability to geographically correlated failures: Analysis and control implications," in *Proc. IEEE INFOCOM IEEE Conf. Comput. Commun.*, 2014, pp. 2634–2642.
- [7] D. Bienstock and A. Verma, "The $n - k$ problem in power grids: New models, formulations, and numerical experiments," *SIAM J. Optim.*, vol. 20, no. 5, pp. 2352–2380, 2010.
- [8] D. Bienstock, "Optimal control of cascading power grid failures," in *Proc. 50th IEEE Conf. Decis. Control Eur. Control Conf.*, Dec. 2011, pp. 2166–2173.
- [9] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman, "Power grid vulnerability to geographically correlated failures 2014: Analysis and control implications," in *Proc. IEEE INFOCOM IEEE Conf. Comput. Commun.*, Apr. 2014, pp. 2634–2642.
- [10] A. Zocca and B. Zwart, "Optimization of stochastic lossy transport networks and applications to power grids," in *Stochastic Syst.*, 2020, vol. 11, no. 1, pp. 34–59, *arXiv:1712.07411v3*.
- [11] C. Liang, L. Guo, A. Zocca, S. Yu, S. Low, and A. Wierman, "An Integrated Approach for Failure Mitigation & Localization in Power Systems," *Electric Power Syst. Res.*, vol. 190, Jun./Jul. 2021, 106613, ISSN 0378-7796.
- [12] A. Zocca, L. Guo, C. Liang, S. H. Low, and A. Wierman, "A Spectral Representation of Power Systems with Applications to Adaptive Grid Partitioning and Cascading Failure Localization," 2021, *arXiv:2105.05234*.
- [13] R. D. Zimmerman, C. E. Murillo-Sánchez, and R. J. Thomas, "MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education," *IEEE Trans. Power Syst.*, vol. 26, no. 1, pp. 12–19, Feb. 2011.
- [14] P. D. Hines, I. Dobson, and P. Rezaei, "Cascading power outages propagate locally in an influence graph that is not the actual grid topology," *IEEE Trans. Power Syst.*, vol. 32, no. 2, pp. 958–967, Mar. 2017.
- [15] P. D. Hines, I. Dobson, E. Cotilla-Sánchez, and M. Eppstein, "Dual graph and random chemistry methods for cascading failure analysis," in *Proc. IEEE 46th Hawaii Int. Conf. Syst. Sci.*, 2013, pp. 2141–2150.
- [16] D. L. H. Aik and G. Andersson, "Analysis of voltage and power interactions in multi-infeed HVDC systems," *IEEE Trans. Power Del.*, vol. 28, no. 2, pp. 816–824, Apr. 2013.
- [17] F. Harary, *Graph Theory*, MA, USA: Addison-Wesley, 1969.



Lingqi Guo received the B.Sc. degree in mathematics and the B.Eng. degree in information engineering from The Chinese University of Hong Kong, Hong Kong, in 2014, and the Ph.D. degree in computing and mathematical sciences from the California Institute of Technology, Pasadena, CA, USA, in 2019. His research interests include the control and optimization of networked systems, with focus on power system frequency regulation, cyber-physical network design, distributed load-side control, synthetic state estimation, and cascading failure analysis.



Chen Liang (Student Member, IEEE) received the B.E. degree in automation from Tsinghua University, Beijing, China, in 2016. He is currently working toward the Ph.D. degree in computing and mathematical sciences with the California Institute of Technology, Pasadena, CA, USA. His research interests include graph theory, mathematical optimization, control theory, and their applications in cascading failures of power systems.



Alessandro Zocca received the B.Sc. degree in mathematics from the University of Padua, Padua, Italy, in 2010, the M.A.St. degree in mathematics from the University of Cambridge, Cambridge, U.K., in 2011, and the Ph.D. degree in mathematics from the University of Eindhoven, Eindhoven, The Netherlands, in 2015. From 2016 to 2017, he was a Postdoctoral Researcher with CWI Amsterdam, Amsterdam, The Netherlands, and from 2017 to 2019, with the California Institute of Technology, Pasadena, CA, USA, where he was supported by his personal NWO Rubicon grant. Since 2019, he has been a tenure-track Assistant Professor with the Department of Mathematics, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands. His work include applied probability and optimization, but has deep ramifications in areas as diverse as operations research, graph theory, algorithm design, statistical physics, and control theory. His research interests include dynamics and rare events on large-scale networked systems affected by uncertainty, drawing motivation from applications to power systems, and wireless networks.



Steven H. Low (Fellow, IEEE) received the B.S. degree in electrical engineering from Cornell University, Ithaca, NY, USA, and the Ph.D. degree in electrical engineering from Berkeley. He is currently the F. J. Gilloon Professor with the Department of Computing and Mathematical Sciences, and the Department of Electrical Engineering, California Institute of Technology, Pasadena, CA, USA. Before that, he was with AT&T Bell Laboratories, Murray Hill, NJ, USA, and the University of Melbourne, Melbourne, VIC, Australia. He has held honorary/chaired professorship in Australia, China, and Taiwan. He was the co-recipient of IEEE best paper awards and is a Fellow of ACM. He was known for pioneering a mathematical theory of Internet congestion control and semidefinite relaxations of optimal power flow problems in smart grid.



Adam Wierman received the B.Sc., M.Sc., and Ph.D. degrees in computer science from Carnegie Mellon University, Pittsburgh, PA, USA, in 2007, 2004, and 2001, respectively. He is currently a Professor with the Department of Computing and Mathematical Sciences, California Institute of Technology (Caltech), Pasadena, CA, USA. Since 2017, he has been a Faculty with Caltech. He was the recipient of multiple awards, including the ACM SIGMETRICS Rising Star Award, IEEE Communications Society William R. Bennett Prize, and multiple teaching awards. He is the co-author of papers which received best paper awards at a wide variety of conferences across computer science, power engineering, and operations research, including ACM Sigmetrics, IEEE INFOCOM, IFIP Performance, and IEEE PES.