

Robust Network Routing under Cascading Failures

Ketan Savla Giacomo Como Munther A. Dahleh

Abstract

We propose a dynamical model for cascading failures in single-commodity network flows. In the proposed model, the network state consists of flows *and* activation status of the links. Network dynamics is determined by a, possibly state-dependent and adversarial, disturbance process that reduces flow capacity on the links, and routing policies at the nodes that have access to the network state, but are oblivious to the presence of disturbance. Under the proposed dynamics, a link becomes irreversibly inactive either due to overload condition on itself or on all of its immediate downstream links. The coupling between link activation and flow dynamics implies that links to become inactive successively are not necessarily adjacent to each other, and hence the pattern of cascading failure under our model is qualitatively different than standard cascade models. The magnitude of a disturbance process is defined as the sum of cumulative capacity reductions across time and links of the network, and the margin of resilience of the network is defined as the infimum over the magnitude of all disturbance processes under which the links at the origin node become inactive. We propose an algorithm to compute an upper bound on the margin of resilience for the setting where the routing policy only has access to information about the local state of the network. For the limiting case when the routing policies update their action as fast as network dynamics, we identify sufficient conditions on network parameters under which the upper bound is tight under an appropriate routing policy. Our analysis relies on making connections between network parameters and monotonicity in network state evolution under proposed dynamics.

I. INTRODUCTION

Resilience is becoming a key consideration in the design and operation of many critical infrastructure systems such as transportation, power, water, gas, and data networks. Due to their increasing scale and interconnectedness, these systems tend to exhibit complex behaviors that pose several new challenges in their design and operation. In particular, while exhibiting good performance in terms of both efficiency and robustness to fluctuations in nominal operating conditions, such systems can be fragile to unexpected local disruptions that may give rise to cascading failures with potentially systemic effects. The problem is further exacerbated by the fact that local actions aimed at mitigating disruptions can increase vulnerability of the other parts of the system.

K. Savla is with the Sonny Astani Department of Civil and Environmental Engineering at the University of Southern California, Los Angeles, CA, USA. ksavla@usc.edu. G. Como is with the Department of Automatic Control, Lund University, Lund, Sweden. giacomo@control.lth.edu. M. A. Dahleh is with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology, Cambridge, MA, USA. dahleh@mit.edu

Models for cascading phenomena in infrastructure networks have been proposed in the statistical physics literature and studied mainly through numerical simulations, e.g., see [1], [2], [3], [4], [5]. Simpler models, based on percolation and other interacting particle systems describing the activation status of nodes and links as dependent on the activation status of their neighbors in the network, have lend themselves to more analytical studies, [6], [7], [8]. While largely used to model the spread of epidemics and rumors in social and economic networks, cascading failures in financial networks and in wireless networks [9], [10], [11], [12], the applicability of the latter models to the design and control of actual physical networks is severely limited because of their simplistic description of the causal relationship between failures of successive nodes and links. In particular, an inherent characteristic of such percolation- and interacting particles-based models is that the successive nodes and links to fail are constrained to be adjacent to each other, which is typically not the case in infrastructure networks (see, e.g., [13]). In particular, there is extensive work on studying cascading failures in power networks using physics-based models, e.g., see [14] for a recent overview. Recently, some of these models have also been analyzed from a control perspective, e.g., see [15], [16].

This paper is concerned with dynamical model for cascading failures in single-commodity flow networks, and with the characterization of maximally resilient routing policies. The modeling framework is inspired by transportation and data networks. When considering dynamical models for cascading failures in physical infrastructure networks, there are several possibilities for time scale separation between link inactivation dynamics under overload, flow dynamics and reaction time of routing (control) policies that can simplify the analysis. The rate of information propagation among geographically distributed routing policies relative to the dynamics can add further complexity. In this paper, we focus our analysis on the limiting case when the rate of information propagation is slow (i.e., routing policies are distributed), and the link inactivation and flow dynamics under routing policies evolve at the same and much faster time scale. Our ability to analyze the dynamical model relies on identifying conditions under which the network state evolves monotonically. Irreversibility in link inactivation in our model naturally implies monotonicity in the link activation status. However, monotonicity in the link flows requires additional restrictions on the routing policy. We study these restrictions under *flow monotonicity* and *link monotonicity* which refer to the sensitivity of the action of a distributed routing policy with respect to changes in inflow (due to changes in the upstream part of the network) and activation status of outgoing links, respectively.

The contributions of the paper are as follows. First, we propose a dynamical model for cascading failures in network flows and formally state the problem of designing maximally resilient routing policies. Second, we propose a backward propagation algorithm for computing an upper bound on the margin of resilience and to motivate the design of a maximally resilient routing policy. Third, we introduce the properties of flow and link monotonicity for distributed routing policies, and show that these are sufficient conditions for the upper bound to be tight. Finally, we identify sufficient conditions on network parameters under which these monotonicity conditions are satisfied.

The paper is organized as follows. In Section II, we present our dynamical model for cascading failures in network flow under routing policies. Section III contains main results of this paper in terms of upper bound computation, monotonicity conditions on the routing policies and sufficient conditions on network parameters to guarantee these

monotonicity properties. The proofs of the main results are collected in Section IV. Section V provides concluding remarks. The appendix contains a technical lemma which is used heavily in the proofs in Section IV.

Before proceeding, we define some preliminary notations to be used throughout the paper. Let \mathbb{R} be the set of real numbers, $\mathbb{R}_+ := \{x \in \mathbb{R} : x \geq 0\}$ be the set of nonnegative real numbers, and \mathbb{N} be the set of natural numbers. When \mathcal{A} is a finite set, $|\mathcal{A}|$ will denote the cardinality of \mathcal{A} , $\mathbb{R}^{\mathcal{A}}$ (respectively, $\mathbb{R}_+^{\mathcal{A}}$) will stand for the space of real-valued (nonnegative-real-valued) vectors whose components are indexed by elements of \mathcal{A} . For $x \in \mathbb{R}^{\mathcal{A}}$ and $y \in \mathbb{R}^{\mathcal{B}}$, x' stands for the transpose of x , and $x \leq y$ means that $x_i \leq y_i$ for all $i \in \mathcal{A} \cap \mathcal{B}$. When $\mathcal{A} = \mathcal{B}$, $x'y$ stands for the dot product of x and y . The all-one and all-zero vectors will be denoted by $\mathbf{1}$ and $\mathbf{0}$, respectively, their size being clear from the context. A directed multigraph is the pair $(\mathcal{V}, \mathcal{E})$ of a finite set \mathcal{V} of nodes, and of a multiset \mathcal{E} of links consisting of ordered pairs of nodes (i.e., we allow for parallel links between a pair of nodes). If $e = (v, w) \in \mathcal{E}$ is a link, where $v, w \in \mathcal{V}$, we shall write $\sigma_e = v$ and $\tau_e = w$ for its tail and head node, respectively. The sets of outgoing and incoming links of a node $v \in \mathcal{V}$ will be denoted by $\mathcal{E}_v^+ := \{e \in \mathcal{E} : \sigma_e = v\}$ and $\mathcal{E}_v^- := \{e \in \mathcal{E} : \tau_e = v\}$, respectively. For $x \in \mathbb{R}$, we shall use the notation $[x]^+$ to mean $\max\{0, x\}$.

II. DYNAMICAL MODEL FOR NETWORK FLOWS AND PROBLEM FORMULATION

In this section, we propose a dynamical model for cascading failure in network flows under distributed routing policies. We model network flows as finite weighted directed multi-graphs $\mathcal{N} = (\mathcal{V}, \mathcal{E}, C)$, where \mathcal{V} and \mathcal{E} stand for the sets of nodes and links, respectively, and $C \in \mathbb{R}^{\mathcal{E}}$ is the vector of link capacities, all assumed to be strictly positive. We refer to nodes with no incoming links as origin nodes and to those with no outgoing links as destination nodes. The set of destination nodes is denoted by \mathcal{D} . Nodes which are neither origin nor destination are referred to as intermediate nodes and are assumed to lie on a path from some origin to some destination.

Let an external inflow $\lambda_o \geq 0$ be associated to every origin node $o \in \mathcal{V}$, and, by convention, put $\lambda_v = 0$ for every other node v . Then, the max-flow min-cut theorem, e.g., see [17], implies that a necessary and sufficient condition for the existence of a feasible equilibrium flow is that the capacity of every cut in the network is larger than the aggregate inflow associated to the non-destination side of the cut. Here, a feasible equilibrium flow refers to a vector $f \in \mathbb{R}_+^{\mathcal{E}}$ satisfying capacity constraints $f_e < C_e$ on every link $e \in \mathcal{E}$, and mass conservation at every non-destination node, i.e, $\lambda_v + \sum_{e \in \mathcal{E}_v^-} f_e = \sum_{e \in \mathcal{E}_v^+} f_e$ for all $v \in \mathcal{V} \setminus \mathcal{D}$. On the other hand, a cut refers to a subset of non-destination nodes $\mathcal{U} \subseteq \mathcal{V} \setminus \mathcal{D}$, with $C_{\mathcal{U}} := \sum_{e \in \mathcal{E} : \sigma_e \in \mathcal{U}, \tau_e \in \mathcal{V} \setminus \mathcal{U}} C_e$ standing for its capacity and $\lambda_{\mathcal{U}} := \sum_{v \in \mathcal{U}} \lambda_v$ for the associated aggregate external inflow. Then, the necessary and sufficient condition for the existence of a feasible equilibrium flow is

$$\max_{\mathcal{U}} \{\lambda_{\mathcal{U}} - C_{\mathcal{U}}\} < 0, \quad (1)$$

with the index \mathcal{U} running over all possible cuts.

We now describe network flow dynamics, evolving in discrete time. Let $\mathcal{N} = (\mathcal{V}, \mathcal{E}, C)$ be a network as above, with inflows λ_o at the origin nodes satisfying condition (1). At every time $t = 0, 1, \dots$, the state of the system is described by a tuple $(\mathcal{V}(t), \mathcal{E}(t), f(t), C(t))$ where: $\mathcal{V}(t) \subseteq \mathcal{V} \setminus \mathcal{D}$ and $\mathcal{E}(t) \subseteq \mathcal{E}$ are the subsets of active non-destination nodes, and links, respectively; $f(t) \in \mathbb{R}_+^{\mathcal{E}}$ is the vector of link flows; and $C(t) \in \mathbb{R}^{\mathcal{E}}$, with $0 \leq C_e(t) \leq$

C_e , is the vector of residual link capacities. The initial condition $(\mathcal{V}(0), \mathcal{E}(0), C(0), f(0))$ is such that $\mathcal{V}(0) = \mathcal{V} \setminus \mathcal{D}$, $\mathcal{E}(0) = \mathcal{E}$, i.e., all non-destination nodes and all links start active, $C(0) = C$, and $f(0)$ is a feasible equilibrium flow for \mathcal{N} .

Given its current state $(\mathcal{V}(t), \mathcal{E}(t), f(t), C(t))$ at time $t = 0, 1, 2, \dots$, the network evolves as follows. All currently active links which become overloaded, i.e., whose current flow exceeds the current residual capacity, along with all those whose head node is currently inactive, become irreversibly inactive, i.e.,

$$\mathcal{E}(t+1) = \mathcal{E}(t) \setminus \{e \in \mathcal{E}(t) : f_e(t) \geq C_e(t)\} \setminus \{e \in \mathcal{E}(t) : \tau_e(t) \notin \mathcal{V}(t)\}. \quad (2)$$

All currently active nodes v that have no active outgoing link become irreversibly inactive, i.e.,

$$\mathcal{V}(t+1) = \mathcal{V}(t) \setminus \{v \in \mathcal{V}(t) : \mathcal{E}_v^+(t) = \emptyset\}. \quad (3)$$

At every currently active node $v \in \mathcal{V}(t)$, a routing policy determines how to split the current inflow $\lambda_v(t) := \lambda_v + \sum_{e \in \mathcal{E}_v^-(t)} f_e(t)$ among the set $\mathcal{E}_v^+(t)$ of its currently active outgoing links, so that

$$f_e(t+1) = G_e(\mathcal{E}_v^+(t), \lambda_v(t)), \quad e \in \mathcal{E}_v^+(t). \quad (4)$$

Finally, the residual capacity vector is reduced by a disturbance $\delta(t) \in \mathbb{R}_+^{\mathcal{E}}$ so that

$$C_e(t+1) = C_e(t) - \delta_e(t+1), \quad e \in \mathcal{E}(t). \quad (5)$$

The sequence $(\delta(1), \delta(2), \dots) \subseteq \mathbb{R}_+^{\mathcal{E}}$ of incremental flow capacity reductions is meant to represent an external, possibly adversarial and network state dependent, process that, without any loss of generality, will be assumed to satisfy

$$\Delta(t) := \sum_{1 \leq s \leq t} \delta(s) \leq C, \quad \forall t \geq 1. \quad (6)$$

Observe that, in writing (4), we have assumed that the routing at node v is determined only by the local observation of the current inflow $\lambda_v(t)$ and the currently active set of outgoing links $\mathcal{E}_v^+(t)$. In particular, the routing policies have no information about the residual link capacities, or equivalently about the disturbance process. The formal definition of distributed oblivious routing policies is as follows.

Definition 1. Given a network $\mathcal{N} = (\mathcal{V}, \mathcal{E}, C)$, a distributed oblivious routing policy \mathcal{G} is a family of functions

$$G^v(\mathcal{J}, \cdot) : \mathbb{R}_+ \rightarrow \mathbb{R}_+^{\mathcal{J}}, \quad v \in \mathcal{V} \setminus \mathcal{D}, \quad \emptyset \neq \mathcal{J} \subseteq \mathcal{E}_v^+,$$

such that, for every $\mu \geq 0$, $\sum_{e \in \mathcal{J}} G_e^v(\mathcal{J}, \mu) = \mu$, and, for all $\mathcal{K} \subseteq \mathcal{J} \subseteq \mathcal{E}_v^+$,

$$G^v(\mathcal{J}, \mu) \leq G^v(\mathcal{K}, \mu). \quad (7)$$

In reading (7), recall our notation established at the end of Section I that, for $x \in \mathbb{R}^{\mathcal{A}}$ and $y \in \mathbb{R}^{\mathcal{B}}$, $x \leq y$ means $x_i \leq y_i$ for all $i \in \mathcal{A} \cap \mathcal{B}$. Definition 1 implicitly implies that $G_e^v(\mathcal{J}, \mu) = 0$ for all $e \in \mathcal{E}_v^+ \setminus \mathcal{J}$. Moreover, we will assume throughout that the initial equilibrium flow $f(0)$ is consistent with the given distributed oblivious routing policy, i.e., $G_e^v(\mathcal{E}_v^+, \lambda_v(0)) = f_e(0)$ for all $e \in \mathcal{E}_v^+$, $v \in \mathcal{V} \setminus \mathcal{D}$. In other words, the initial equilibrium flow

is specified by the routing policy and, as long as there is no perturbation, i.e., $\delta(t) = 0$, the network state does not change. The term *oblivious* in distributed routing policies is meant to emphasize that routing policies have no information about the disturbance process. Hereafter, unless explicitly stated otherwise, we shall refer to a routing policy satisfying Definition 1 simply as a distributed routing policy. Equation (7) implies that, at every node, for a fixed inflow, shrinking of the set of active links results in increase in flow assigned to each of the remaining active outgoing links. We shall refer to (7) as the *link monotonicity* property. While (7) represents a natural condition for distributed routing policies, the maximally resilient routing policies designed in this paper have been found to satisfy it. Alternately, one could regard the results in this paper to be optimal within this class of distributed routing policies. We provide additional comments on this aspect in Remark 7.

Remark 1. Condition (7) is satisfied by any routing policy at a node v if $|\mathcal{E}_v^+| \leq 2$.

A simple example of a distributed routing policy is the *proportional* policy: for every $\emptyset \neq \mathcal{J} \subseteq \mathcal{E}_v^+$, $v \in \mathcal{V} \setminus \mathcal{D}$, $\mu \geq 0$:

$$G_e^v(\mathcal{J}, \mu) = \begin{cases} (C_e / \sum_{j \in \mathcal{J}} C_j) \mu & \text{if } e \in \mathcal{J}, \\ 0 & \text{if } e \notin \mathcal{J}. \end{cases} \quad (8)$$

The model in (2)-(5) has several salient features worth emphasizing. First, note that the transition from active to inactive status of a link is irreversible. Second, note that a link could become inactive either because it is overloaded or because its downstream node becomes inactive. The mismatch between flow and residual capacity of a link, which gives rise to overload condition, depends on the disturbance process and the action of a distributed routing policy. Therefore, the links to fail successively are not necessarily adjacent to each other. Finally, note that in our model, the routing policy updates its action at the same time scale as flow and link inactivation dynamics. An implication of this is that the flow vector f may not be an equilibrium flow at all time instants because of violation of flow conservation at some nodes. It is possible to extend (2)-(5) to model scenarios representing a combination of centralized and non-oblivious routing, link recovery after finite time and long range coupling between failure of links. However, the analysis presented in this paper is restricted to the model in (2)-(5).

The following example illustrates cascading failure under the dynamics in (2)-(5).

Example 1. Consider the graph topology depicted in Figure 1, where the flow capacities are given by $C_1 = 4$, $C_2 = 4$, $C_3 = 1.5$, $C_4 = 2.5$, $C_5 = 1$, $C_6 = 1$, $C_7 = 1$ and $C_8 = 1$. Let the arrival rate at the origin be $\lambda = 4$. We consider proportional routing policies, specified in (8), at all the nodes, under which the initial flow on all links is given by $f_1(0) = 2$, $f_2(0) = 2$, $f_3(0) = 3/4$, $f_4(0) = 5/4$, $f_5(0) = 3/8$, $f_6(0) = 3/8$, $f_7(0) = 5/8$ and $f_8(0) = 5/8$. Consider the network dynamics under a disturbance process for which $\delta_3(1) = 0.75$, $\delta_3(t) = 0$ for all $t \geq 2$ and $\delta_i(t) \equiv 0$ for all $i \in \{1, \dots, 8\} \setminus \{3\}$. Since $C_3(1) = 1.5 - 0.75 = 0.75 = f_3(1)$, $e_3 \notin \mathcal{E}(2)$. Therefore, $f_4(3) = 2$, $f_7(4) = 1$ and $f_8(4) = 1$. Since $C_7(4) = C_8(4) = 1$, $e_7 \notin \mathcal{E}(5)$ and $e_8 \notin \mathcal{E}(5)$. By continuing along these lines, the order of links to become inactivate is $e_3, e_7 = e_8, e_4, e_1, e_2$. In particular, the link to fail after e_3 is e_7 , demonstrating non-adjacency in successive link failures under our proposed network dynamics.

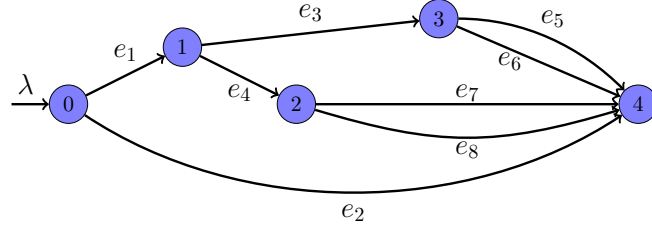


Figure 1. A simple graph for the illustration of cascading failure under the proposed network dynamics.

Remark 2. *The model in (2)-(5) is to be contrasted with the dynamical flow network formulation in our previous work [18], [19] where every link has infinite buffer capacity, and hence there are no cascading effects under link overload. These features are relaxed in our subsequent work [20], where the links are modeled to have finite buffer capacity, and the control policy at every node implements routing as well as flow control under information about the densities and the disturbances on the links incoming and outgoing from that node. Such a framework allows for backward propagation effects, which were proven to increase the resilience of the network with respect to the framework in [18], [19]. Such control policies were also shown to exhibit graceful collapse, i.e., when the inflow to the network exceed its capacity, then all the critical links saturate simultaneously. In this paper, we constrain the actions of the control policies to only routing, and under no information about the disturbance. We emphasize that, although the routing policies have no explicit information about disturbance on the links, they have information about its effect on the activation status on the local links. On the other hand, due to cascade effects, the change in the activation status of a link may not be exclusively due to disturbance on that link.*

A. Problem Formulation

In this paper, the performance criterion of interest is the ability of a network to transfer flow from the origin nodes to the destination nodes, under a wide range of disturbance processes. We formalize this notion as follows.

Definition 2. *Let \mathcal{N} be a network, λ a vector of inflows at the origin nodes, \mathcal{G} a distributed routing policy, and $(\delta(t))_{t \geq 1}$ a disturbance process. Then, the associated network flow dynamics in (2)-(5) is said to be transferring if*

$$\lim_{t \rightarrow +\infty} \sum_{d \in \mathcal{D}} \sum_{e \in \mathcal{E}_d^-} f_e(t) = \sum_v \lambda_v, \quad (9)$$

where the summation in v is over the origin nodes.

Observe that, since $f(0)$ is assumed to be a feasible equilibrium flow, one has that, at time 0, the aggregate outflow from, and inflow to, the network match, i.e., $\sum_{d \in \mathcal{D}} \sum_{e \in \mathcal{E}_d^-} f_e(0) = \lambda$. Definition 2 requires that, for a network \mathcal{N} and a distributed routing policy \mathcal{G} to be transferring under a disturbance process $(\delta(t))_{t \geq 1}$, aggregate inflow to, and outflow from the network also match asymptotically. For disturbance processes that are active only over finite time, (9) can be rephrased to require the inflow and the outflow to match at all times with the possible

exception of a finite transient. We shall use this latter formulation in Section III, where the setup allows to focus only on finite time disturbance processes without loss of generality.

The magnitude of a disturbance process δ is defined as (see (6)):

$$\mathcal{D}(\delta) := \sum_{e \in \mathcal{E}} \Delta_e(\infty).$$

Definition 3. Let \mathcal{N} be a network, λ a vector of inflows at the origin nodes, and \mathcal{G} a distributed routing policy. The margin of resilience of the network, denoted as $\mathcal{R}(\mathcal{N}, \lambda, \mathcal{G})$, is defined as the infimum of the magnitude of disturbance processes under which the associated dynamics is not transferring, i.e.,

$$\mathcal{R}(\mathcal{N}, \lambda, \mathcal{G}) := \inf_{\delta} \{ \mathcal{D}(\delta) \mid \text{network flow dynamics in (2)-(5) for } \mathcal{N}, \lambda, \mathcal{G}, \delta \text{ is not transferring} \}.$$

We are now ready to formally state the problem. Our objective in this paper is to (i) compute the margin of resilience under distributed routing policies; and (ii) identify maximally resilient distributed routing policies. Formally, we consider the following optimization problem:

$$\mathcal{R}^*(\mathcal{N}, \lambda) = \sup_{\mathcal{G}} \mathcal{R}(\mathcal{N}, \lambda, \mathcal{G}), \quad (10)$$

where the supremum is over the class of distributed routing policies. A distributed routing policy \mathcal{G} is called *maximally resilient* if $\mathcal{R}(\mathcal{N}, \lambda, \mathcal{G}) = \mathcal{R}^*(\mathcal{N}, \lambda)$.

III. MAIN RESULTS

In this section, we present our main results addressing problem (10). From now on, we will be restricted to networks $\mathcal{N} = (\mathcal{V}, \mathcal{E}, C)$ with a single origin destination pair. We will identify the node set \mathcal{V} with the integer set $\{0, 1, \dots, n\}$, with 0 and n associated with the unique origin and destination nodes, respectively. Moreover, let $\lambda > 0$ be the constant inflow at the unique origin node. While extensions to multiple destinations are straightforward, extensions to multiple origin nodes is not trivial. We start by giving simple bounds on the margin of resilience.

A. Simple Bounds

It is straightforward to obtain the following upper and lower bounds on the margin of resilience, valid for every routing policy \mathcal{G}

$$\min_{e \in \mathcal{E}} \{C_e - f_e(0)\} \leq \mathcal{R}(\mathcal{N}, \lambda, \mathcal{G}) \leq \min_{\mathcal{U}} C_{\mathcal{U}} - \lambda, \quad (11)$$

where the minimization in the upper bound is over all the cuts in \mathcal{N} . The lower bound in (11) is due to the fact that at least one link needs to become inactive to ensure non transferring of the network, possibly under cascading failure, and $\min_{e \in \mathcal{E}} \{C_e - f_e(0)\}$, which is the minimum among all link residual capacities, corresponds to the disturbance process with minimum magnitude that can cause a link to become inactive. The upper bound in (11), which is usually referred to as the network residual capacity, is obtained by noting that the network is non-transferring under a disturbance process that removes residual capacity at $t = 1$ from the links that constitute a min cut of \mathcal{N} . As it may be expected, the gap between the upper and lower bounds can be arbitrarily large in general networks.

As an illustration, in Example 1, the minimum link residual capacity is $3/18$, corresponding to links e_7 and e_8 , and the network residual capacity is 4, corresponding to, e.g., the cut $\{1, 2\}$. However the example also constructs a disturbance process of magnitude 0.75 under which the network is not transferring (under proportional routing policy).

We now describe a recursive procedure to compute a sharper upper bound. The quantity computed by this procedure can be intuitively related to the margin of resilience when there is a clear time-scale separation between the flow dynamics (fastest), speed of control action (intermediate) and the link inactivation dynamics (slowest), and the routing policy is *centralized* but oblivious to disturbance. In our current setup, it will provide an upper bound. Let $\mathcal{X}(\mathcal{J}, \lambda)$ be the set of equilibrium flow vectors when the inflow at the origin node is λ and the set of active links is $\mathcal{J} \subseteq \mathcal{E}$, i.e., $\mathcal{X}(\mathcal{J}, \lambda)$ is the set of $f \in \mathbb{R}_+^{\mathcal{J}}$ satisfying $f_e < C_e$ on every link $e \in \mathcal{J}$, $\sum_{e \in \mathcal{E}_0^+ \cap \mathcal{J}} f_e = \lambda$ and $\sum_{e \in \mathcal{E}_v^+ \cap \mathcal{J}} f_e = \sum_{e \in \mathcal{E}_v^- \cap \mathcal{J}} f_e$ for all $v \in \mathcal{V} \setminus \{n\}$. Let $S(\mathcal{J}, \lambda)$ correspond to the margin of resilience when the network starts with active link set $\mathcal{J} \subseteq \mathcal{E}$. The computation of $S(\mathcal{E}, \lambda)$, which corresponds to margin of resilience of interest, is based on values of $S(\mathcal{J}, \lambda)$ for all $\mathcal{J} \subseteq \mathcal{E}$ as follows. Starting with sets $\mathcal{J} \subseteq \mathcal{E}$ of size one, i.e., $|\mathcal{J}| = 1$, and then increasing in size, perform the following recursion: if $\mathcal{X}(\mathcal{J}, \lambda) = \emptyset$, then $S(\mathcal{J}, \lambda) = 0$, else

$$S(\mathcal{J}, \lambda) := \max_{x \in \mathcal{X}(\mathcal{J}, \lambda)} \min_{e \in \mathcal{J}} \left(C_e - x_e + S(\mathcal{J} \setminus \{e\}, \lambda) \right). \quad (12)$$

In (12), $C_e - x_e$ is the difference between the original capacity of link e and the flow on it under control action x by a (centralized) routing policy. Therefore, it represents the minimum disturbance on link e that will make it inactive. Since $S(\mathcal{J} \setminus \{e\}, \lambda)$ represents the margin of resilience once link e is removed, the whole term inside the minimum in (12) is the magnitude of the smallest disturbance required to make the network with \mathcal{J} links to become non-transferring starting with removal of link e . The minimization in (12) over all $e \in \mathcal{J}$ searches for the initial link e whose inactivation will minimize the disturbance required to make the network non-transferring. The outer maximization over the feasible action set of the (centralized) routing policy maximizes the magnitude of the worst-case disturbance process that will make the network non-transferring. The next results states the the output of the iterations in (12) is indeed an upper bound on the margin of resilience under any distributed routing policy.

Proposition 1. *Let $\mathcal{N} = (\mathcal{V}, \mathcal{E}, C)$ be a network with $\lambda > 0$ a constant total inflow at the origin node. Then, for any distributed routing policy, there exists a disturbance process $(\delta(t))_{t \geq 1}$ with $\mathcal{D}(\delta) \leq S(\mathcal{E}, \lambda)$ under which the associated network flow dynamics (2)-(5) is not transferring.*

Remark 3. (i) *It is easy to prove that $S(\mathcal{E}, \lambda)$ is a tighter upper bound than the one in (11).*

(ii) *The computation of $S(\mathcal{E}, \lambda)$ involves recursion over all sets \mathcal{J} in $2^{\mathcal{E}}$. However, for each $e \in \mathcal{J} \subseteq \mathcal{E}$ and $\lambda \geq 0$, the term inside the minimization in (12) is affine in $x \in \mathcal{X}(\mathcal{J}, \lambda)$. Hence, computing $S(\mathcal{J}, \lambda)$ is a convex optimization problem.*

For Example 1, by simulations, we find that $S(\mathcal{E}, \lambda) \approx 1.75$, which is less than 4, the value corresponding to the upper bound in (11). However, it is still greater than 0.75, the magnitude of disturbance in Example 1, under which the network is not transferring. This conservatism is because the recursions in (12) implicitly assume centralized

routing, and do not take into account the possibility of link inactivation due to the inactivation of the corresponding head node. In Section III-C, we propose an algorithm, the Backward Propagation Algorithm (BPA), that addresses these limitations to provide a tighter upper bound, and we identify conditions under which this upper bound is provably tight. The BPA is designed for network topologies satisfying the following acyclicity assumption.

Assumption 1. $(\mathcal{V}, \mathcal{E})$ contains no cycles.

A consequence of Assumption 1, the oblivious property of routing policies and the finiteness of \mathcal{V} and \mathcal{E} is that, we can assume without loss of generality that, for every $e \in \mathcal{E}$, there exists at most one $t_e \geq 0$ such that $\delta_e(t_e) > 0$, and that $\delta(t) = \mathbf{0}$ after some finite time. Therefore, it is sufficient to restrict our attention to disturbance processes δ that are non-zero only for a finite time, and hence there exists a finite time after which $(\mathcal{V}(t), \mathcal{E}(t), f(t), C(t))$ comes to a steady state under any such disturbance process δ . Let \mathcal{T} denote that finite termination time. In this case, Definition 2 simplifies as: network flow dynamics is transferring if $\lambda_n(\mathcal{T}) = \lambda$.

The formulation and analysis of the BPA implicitly relies on the following simple result showing an equivalence between a network being transferring and its origin node being active all the time.

Proposition 2. Let \mathcal{N} be a network satisfying Assumption 1 with λ a constant inflow at the origin node, \mathcal{G} a routing policy, and $(\delta(t))_{t \geq 1}$ a disturbance process. Then, the associated network flow dynamics (2)-(5) is transferring if and only if $0 \in \mathcal{V}(\mathcal{T})$. Moreover, $\lambda_n(\mathcal{T}) \in \{0, \lambda\}$.

Remark 4. The analyses of conventional models for cascading failure focus primarily on the connectivity of the residual graph $(\mathcal{V}(\mathcal{T}), \mathcal{E}(\mathcal{T}))$. For the setting of this paper, the proof of Proposition 2 can be used to easily show that there exists a directed path from 0 to n in $(\mathcal{V}(\mathcal{T}), \mathcal{E}(\mathcal{T}))$ if and only if the associated network flow dynamics is transferring.

B. Simple settings

Before describing the BPA, we present results for the maximal margin of resilience and the maximally resilient routing policy in simple settings. We use these calculations merely to motivate the key steps in the construction of BPA, and refer to Theorems 1 and 2 for their rigorous justification.

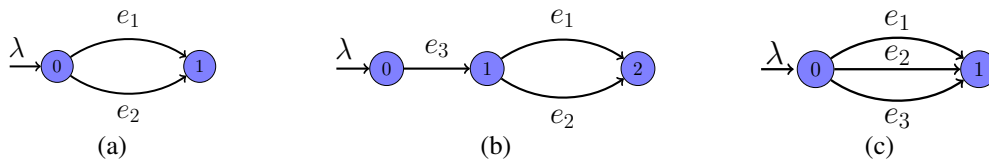


Figure 2. Illustrative simple network topologies.

Let \mathcal{N}_1 denote the flow network illustrated in Figure 2(a), with $\mathcal{E}_1 = \{e_1, e_2\}$, maximum link flow capacities C_i , $i = 1, 2$. The routing policy at node 0 is completely specified by any vector $x \in \mathcal{X}(\mathcal{E}_1, \lambda)$. Considering all the

possible outcomes of the disturbance process, the margin of resilience is given by:

$$\begin{aligned} \mathcal{R}(\lambda, \mathcal{N}_1, x) &= \min \{ [C_1 - x_1]^+ + [C_2 - \lambda]^+, [C_2 - x_2]^+ + [C_1 - \lambda]^+, [C_1 - x_1]^+ + [C_2 - x_2]^+ \} \\ &= \min \{ [C_1 - x_1]^+ + [C_2 - \lambda]^+, [C_2 - x_2]^+ + [C_1 - \lambda]^+ \}, \end{aligned} \quad (13)$$

where first term inside the min in the right hand side of (13) corresponds to the inactivation of link e_1 at $t = 2$ under $\delta(1) = [[C_1 - x_1]^+ \ 0]^T$ followed by inactivation of e_2 at $t = 3$ under $\delta(2) = [0 \ [C_2 - \lambda]^+]^T$, the second term corresponds to inactivation of e_2 at $t = 2$ followed by inactivation of e_1 at $t = 3$, and the third term corresponds to the simultaneous inactivation of links e_1 and e_2 at $t = 2$ under $\delta(1) = [[C_1 - x_1]^+ \ [C_2 - x_2]^+]^T$. Therefore, the maximum possible margin of resilience, and the corresponding maximally resilient routing policy are, respectively, given by $\mathcal{R}^*(\mathcal{N}_1, \lambda) = \max_{x \in \mathcal{X}(\mathcal{E}_1, \lambda)} \mathcal{R}(\mathcal{N}_1, \lambda, x)$ and $x^* = \operatorname{argmax}_{x \in \mathcal{X}(\mathcal{E}_1, \lambda)} \mathcal{R}(\mathcal{N}_1, \lambda, x)$, which, using simple algebra, can be computed as (see Figure 3 (a) for an illustration):

$$\mathcal{R}^*(\mathcal{N}_1, \lambda) = \begin{cases} C_1 + C_2 - 3\lambda/2 & \text{if } \lambda \in [0, \underline{C}], \\ \underline{C}/2 + \bar{C} - \lambda & \text{if } \lambda \in [\underline{C}, \bar{C}], \\ (C_1 + C_2)/2 - \lambda/2 & \text{if } \lambda \in [\bar{C}, C_1 + C_2], \\ 0 & \text{if } \lambda \geq C_1 + C_2, \end{cases} \quad (14)$$

where $\underline{C} := \min \{C_1, C_2\}$ and $\bar{C} := \max \{C_1, C_2\}$, and (see Figure 3 (b) for an illustration)

$$x_1^*(\lambda) = \begin{cases} \lambda/2 & \text{if } \lambda \in [0, \underline{C}] \\ C_1/2 + (\lambda - (C_1 + C_2))/2 \mathbb{1}_{C_1 > C_2} & \text{if } \lambda \in [\underline{C}, \bar{C}] \\ \lambda/2 + (C_1 - C_2)/2 & \text{if } \lambda \in [\bar{C}, C_1 + C_2]. \end{cases} \quad (15)$$

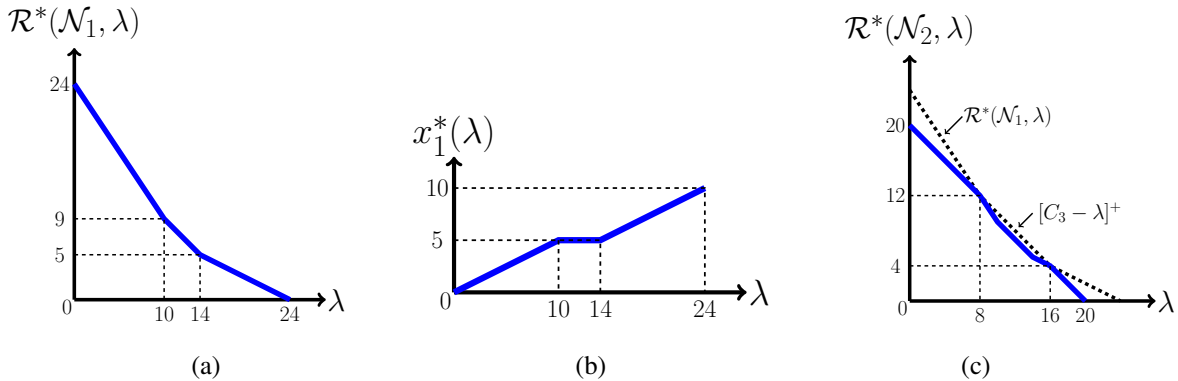


Figure 3. Illustrations of (a) $\mathcal{R}^*(\mathcal{N}_1, \lambda)$ and (b) $x_1^*(\lambda)$ for $C_1 = 10$ and $C_2 = 14$; and (c) $\mathcal{R}^*(\mathcal{N}_2, \lambda)$ for $C_3 = 20$, $C_1 = 10$ and $C_2 = 14$.

Figure 3 (b) illustrates that proportional routing policies (e.g., see (8)) where the proportionality constants are independent of the arrival rate are in general not maximally resilient. For \mathcal{N}_1 , comparing (13) and (12), we see that $\mathcal{R}^*(\mathcal{N}_1, \lambda) = S(\mathcal{E}_1, \lambda)$.

Let \mathcal{N}_2 denote the flow network illustrated in Figure 2(b), with maximum link flow capacities C_i , $i = 1, 2, 3$. Following the dynamics in (2)-(5), link e_3 can become inactivate because of saturation of link e_3 or because of

inactivation of node $\tau_{e_3} = 1$. Accordingly, the maximum margin of resilience is the minimum of $[C_3 - \lambda]^+$ and the maximum margin of resilience of the sub-network rooted at node 1 when the inflow at node 1 is λ . From our analysis of \mathcal{N}_1 , the latter is equal to $\mathcal{R}^*(\mathcal{N}_1, \lambda)$, and hence the maximally resilient routing policy at node 1 is the same as in (15). An illustration of $\mathcal{R}^*(\mathcal{N}_2, \lambda)$ is given in Figure 3 (c). In this case, $S(\mathcal{E}, \lambda)$ from (12) is in general not equal to $\mathcal{R}^*(\mathcal{N}_2, \lambda)$. This is because (12) does not take into account the possibility of e_3 becoming inactive due to inactivation of node 1.

Before proceeding with our next example, we define feasible flow vectors over active local links at node v . This will be the set of feasible control actions for the routing policy at node v . For $\mathcal{J} \subseteq \mathcal{E}_v^+$, $r \in \mathbb{R}_+^{\mathcal{E}_v^+}$, $v \in \mathcal{V} \setminus \{n\}$ and $\mu \geq 0$, let

$$\mathcal{X}_v(\mathcal{J}, r, \mu) := \left\{ x \in \mathbb{R}_+^{\mathcal{J}} : r \leq x \leq C; \mathbf{1}'x = \mu \right\}. \quad (16)$$

(16) is a local version of the set $\mathcal{X}(\mathcal{J}, \lambda)$ used in (12) with two exceptions. First, (16) is defined for a generic inflow μ , since the inflow at v is time-varying. Second, as will become clear in the construction of BPA, the presence of element-wise lower bound r allows one to impose the link monotonicity property defined in (7).

Let \mathcal{N}_3 denote the flow network illustrated in Figure 2(c), with maximum link flow capacities C_i , $i = 1, 2, 3$. Let $S(\mathcal{J}, r, \lambda)$ be the margin of resilience when the set of active links at node v is $\mathcal{J} \subseteq \mathcal{E}_v^+$, and when the action of routing policy is constrained to be element-wise greater than r . For all $e \in \mathcal{E}_v^+$, $S(e, r, \lambda) = 0$ if $\lambda < r_e$ or $\lambda \geq C_e$, and $S(e, r, \lambda) = C_e - \lambda$ otherwise. For $|\mathcal{J}| \geq 2$ in \mathcal{N}_3 , one can write the following recursion:

$$S(\mathcal{J}, r, \lambda) = \max_{x \in \mathcal{X}_0(\mathcal{J}, r, \mu)} \min_{e \in \mathcal{J}} \left(C_e - x_e + S(\mathcal{J} \setminus \{e\}, x, \lambda) \right). \quad (17)$$

Inside the minimization in (17), the term $C_e - x_e$ is the difference between the capacity of link e when the flow on it is x_e , and hence represents the minimal disturbance required to make link e inactive under routing action x . The term $S(\mathcal{J} \setminus \{e\}, x, \lambda)$ is the magnitude of disturbance that is sufficient to make the network non-transferring after link e has become inactive, under the constraint that the flows on links in $\mathcal{J} \setminus \{e\}$ can not be element-wise less than the flow x on them when e was active. The margin of resilience for \mathcal{N}_3 is then $S(\mathcal{E}, \mathbf{0}, \lambda)$. The ability of (17) to incorporate link monotonicity constraints yields a sharper upper bound in comparison to (12).

The recursion in (17) can be used to derive margin of resilience for a network with arbitrary number of links between nodes 0 and 1 in Figure 2 (c). However, in order to handle networks with arbitrary number of links between nodes 1 and 2 in Figure 2 (b), we need to include the effect of inactivation of downstream nodes into (17). This is the basis of the Backward Propagation Algorithm, which we describe next.

C. The Backward Propagation Algorithm (BPA)

We now describe the Backward Propagation Algorithm (BPA) to compute a tighter upper bound on the margin of resilience in comparison to Proposition 1. The same algorithm will also motivate the design of BPA routing which will be proven to be maximally resilient under certain sufficient conditions.

Assumption 1 implies that one can find a (not necessarily unique) topological ordering of the node set $\mathcal{V} = \{0, \dots, n\}$ (see, e.g., [21]). We shall assume to have fixed one such ordering in such a way that $\mathcal{E}_v^- \subseteq \bigcup_{0 \leq u < v} \mathcal{E}_u^+$

for all $v = 1, \dots, n$. We recall that the *depth* of a graph $(\mathcal{V}, \mathcal{E})$ satisfying Assumption 1 is the length of the longest directed path in $(\mathcal{V}, \mathcal{E})$.

Algorithm 1: The Backward Propagation Algorithm (BPA)

- 1: $S(\mathcal{E}_n^+, r, \mu) := +\infty$ for all $r \in \mathbb{R}_+^{\mathcal{E}_n^+}$ and $\mu \geq 0$ {destination node}
- 2: **for** $v = n - 1, n - 2, \dots, 0$ **do** {construct a series of intermediate functions for every node starting with $n - 1$, and going backward up to the origin}
- 3: for all $r \in \mathbb{R}_+^{\mathcal{E}_v^+}$ and $\mu \geq 0$,

$$S(\emptyset, r, \mu) = 0$$

$$S(\mathcal{J}, r, \mu) := 0 \text{ if } \mathcal{X}_v(\mathcal{J}, r, \mu) = \emptyset, \quad \forall \emptyset \neq \mathcal{J} \subseteq \mathcal{E}_v^+,$$

$$S_e(\mu) = S(e, r, \mu) := \min \left\{ C_e - \mu, S(\mathcal{E}_{\tau_e}^+, \mathbf{0}, \mu) \right\} \quad \forall e \in \mathcal{E}_v^+. \quad (18)$$

- 4: iteratively compute $S(\mathcal{J}, r, \mu)$ for $\mathcal{J} \subseteq \mathcal{E}_v^+$ of increasing size, starting with sets of size 2:

$$S(\mathcal{J}, r, \mu) := \max_{x \in \mathcal{X}_v(\mathcal{J}, r, \mu)} \min_{e \in \mathcal{J}} \left(S_e(x_e) + S(\mathcal{J} \setminus \{e\}, x, \mu) \right) \quad (19)$$

- 5: **end for**
-

Note that r appears only in the constraint set in the right hand side of (19). The fundamental difference between the recursions in (19) and (17) is in the first term inside the minimization in (19). This term represents the minimum magnitude of disturbance required to make a link inactive. While it was sufficient to consider the disturbance on link e for this purpose in \mathcal{N}_3 , for general networks, (18) implies that the minimal disturbance could correspond to making the downstream node inactive. Therefore, the recursive computations at node v depend on the outcome of the computations done for nodes downstream to v . The Backward Propagation Algorithm derives its name from the central feature of the algorithm, where an intermediate node collects $S(\mathcal{E}_v^+, r, \mu)$ functions from its downstream nodes, performs updates with respect to local network parameters, and transmits it to upstream nodes. As such, the BPA can be executed in a distributed fashion.

Complementary to the maximization in (19) is the set of corresponding maximizers:

$$g(\mathcal{J}, r, \mu) := \operatorname{argmax}_{x \in \mathcal{X}_v(\mathcal{J}, r, \mu)} \min_{e \in \mathcal{J}} \left(S_e(x_e) + S(\mathcal{J} \setminus \{e\}, x, \mu) \right). \quad (20)$$

A simple implication of (20) which is used heavily in the paper is:

$$z \in g(\mathcal{J}, r, \mu) \implies z \geq r. \quad (21)$$

D. Upper bound on the margin of resilience

The quantity $S(\mathcal{E}_0^+, \mathbf{0}, \lambda)$ computed by BPA is next shown to be an upper bound on the margin of resilience under any distributed routing policy. For brevity in notation, we let $S^*(\mathcal{N}, \lambda) := S(\mathcal{E}_0^+, \mathbf{0}, \lambda)$.

Theorem 1. *Let \mathcal{N} be a network satisfying Assumption 1 and with λ a constant inflow at the origin node. Then, for any distributed routing policy \mathcal{G} , there exists a disturbance process $(\delta(t))_{t \geq 1}$ with $\mathcal{D}(\delta) \leq S^*(\mathcal{N}, \lambda)$ under which the associated network flow dynamics (2)-(5) is not transferring.*

Remark 5. (i) *Theorem 1 implies that $\mathcal{R}(\mathcal{N}, \lambda, \mathcal{G}) \leq S^*(\mathcal{N}, \lambda)$ for all distributed routing policies \mathcal{G} , and hence $\mathcal{R}^*(\mathcal{N}, \lambda) \leq S^*(\mathcal{N}, \lambda)$.*

(ii) *While the statement of Theorem 1 merely suggests the existence of a worst-case disturbance process, its proof in Section IV-D explicitly constructs one such disturbance process. Therefore, in scenarios when the upper bound in Theorem 1 is tight, this constructive procedure can also be used to identify the most vulnerable links of the network for adversarial setting.*

(iii) *The computational complexity of BPA has tradeoffs in comparison to (12). On one hand, while the recursion in (12) involves all elements in $2^{\mathcal{E}}$, BPA involves all elements only in $\cup_v 2^{\mathcal{E}_v^+}$, which is much smaller in comparison, especially when $|\mathcal{V}|$ is large. On the other hand, (12) involves computation only for a fixed λ , whereas BPA involves computations, in general, for all $\mu \in [0, \lambda]$ and $r \in \mathcal{X}_v(\mathcal{J}, \mathbf{0}, \mu)$. Moreover, whereas each recursion in (12) is a convex optimization problem (see Remark 3 (ii)), BPA does not enjoy this property in general. This is because, under (18), the expression inside the minimization in (19) is in general not affine in x , e.g., see Figure 3 (c) for an illustration.*

E. BPA routing and lower bound on the margin of resilience

In this section, we develop lower bounds for $\mathcal{R}^*(\mathcal{N}, \lambda)$. This will be done by analyzing a specific distributed routing policy, called *BPA-routing*, whose construction is inspired by the Backward Propagation Algorithm. BPA routing is a routing policy that satisfies the following for all $v \in \mathcal{V} \setminus \{n\}$, $\mu \geq 0$:

$$\begin{aligned} r^* &:= G^v(\mathcal{E}_v^+, \mu) \in g(\mathcal{E}_v^+, \mathbf{0}, \mu), \\ G^v(\mathcal{J}, \mu) &\in g(\mathcal{J}, r^*, \mu), \quad \mathcal{J} \subseteq \mathcal{E}_v^+. \end{aligned} \tag{22}$$

BPA routing derives its name from the fact that it relies on the function $g(\mathcal{J}, r, \mu)$ from (20), which is directly related to the central computation in the BPA. However, note that the lower bound r^* in (22) is independent of \mathcal{J} and μ , and is always equal to the action of the routing policy under the same inflow μ , when all local links are active, and with no lower bound constraint. Following (21), $G^v(\mathcal{J}, \mu) \geq G^v(\mathcal{E}_v^+, \mu)$ for all $\mathcal{J} \subseteq \mathcal{E}_v^+$ and $\mu \geq 0$. The following lemma formally states conditions under which BPA routing satisfies the link monotonicity in (7).

Lemma 1. *Let \mathcal{N} be a network satisfying Assumption 1 with $|\mathcal{E}_v^+| \leq 3$ for all $v \in \mathcal{V} \setminus \{n\}$. Then BPA routing defined in (22) and (20) satisfies (7), and hence is a distributed routing policy as per Definition 1.*

Proof: The only non-trivial case to prove is that, for every $j \in \mathcal{J} \subseteq \mathcal{E}_v^+$, $\mu \geq 0$, BPA routing satisfies:

$$G^v(\mathcal{E}_v^+, \mu) \leq G^v(\mathcal{E}_v^+ \setminus \{j\}, \mu).$$

This is straightforward since (21) implies $G^v(\mathcal{E}_v^+ \setminus \{j\}, \mu) \geq r^* = G^v(\mathcal{E}_v^+, \mu)$. ■

In general, BPA routing is not readily maximally resilient for general networks. This is because if a non-destination node v has multiple incoming links, then inactivation of v results in inactivation of all the incoming links. However, the BPA algorithm does not take into account such correlations between link inactivations and hence the upper bound in Theorem 1 is conservative for networks which are not trees. While it is possible to modify BPA algorithm to reduce this conservatism, this comes with additional computational complexity and additional difficulty in formulating the corresponding maximally resilient routing policy. Therefore, we make the following directed tree assumption in this paper for deriving lower bound on the margin of resilience.

Assumption 2. $(\mathcal{V} \setminus \{n\}, \mathcal{E} \setminus \mathcal{E}_n^-)$ is a directed tree¹.

We emphasize that Assumption 2 puts a tree-like assumption on the original graph *minus the destination node n , and the set of links \mathcal{E}_n^- incoming to the destination node*. In particular, under Assumption 2, the destination node n is allowed to have multiple incoming links. For example, the graph topology illustrated in Figure 1 satisfies Assumption 2. However, with a slight abuse of terminology, we refer to \mathcal{N} satisfying Assumption 2 as a tree. Note that, \mathcal{N} satisfying Assumption 2 is a tree rooted at the unique origin node.

Remark 6. *For a network satisfying Assumption 2, if λ is less than the min cut capacity, then $f(0)$ under BPA routing is an equilibrium flow. Recall that the max flow min cut theorem implies that this is also a necessary condition for the existence of an equilibrium flow.*

BPA routing is maximally resilient on flow networks which are trees and *symmetric*. Recall that a weighted rooted tree of depth one is called symmetric if all the links outgoing from the root node have equal weights. A weighted rooted tree of depth greater than one is called symmetric if all the sub-trees rooted at the children² nodes are symmetric, and identical to each other.

Proposition 3. *Let \mathcal{N} be a symmetric network satisfying Assumption 2 with $\lambda > 0$ a constant inflow at the origin node and BPA routing policy. Then, the associated network flow dynamics (2)-(5) is transferring for every disturbance process $(\delta(t))_{t \geq 1}$ with $\mathcal{D}(\delta) < S^*(\mathcal{N}, \lambda)$.*

The tree assumption is not sufficient for BPA routing to match the upper bound $S^*(\mathcal{N}, \lambda)$ given by the BPA for networks which are not symmetric, as illustrated in the following example.

Example 2. *Consider the graph topology from Figure 1, with $\lambda = 2$, $C_1 = 2.5$, $C_2 = 0.17$, $C_3 = 0.6$, $C_4 = 2$, $C_5 = 2$, $C_6 = 2$, $C_7 = 0.6$ and $C_8 = 1.5$.*

The plot of $x_3^(\mu) := G_{e_3}(\mathcal{E}_1^+(0), \mu)$ vs. μ under BPA routing for these values is given in Figure 4, which shows that $x_3^*(\mu)$ is decreasing in μ over $[1.9, 2]$. Also, for these values, $S^*(\mathcal{N}, \lambda) = 0.3$. Consider a disturbance process such that $\delta_2(1) = 0.07$, $\delta_3(1) = 0.2$, $\delta_i(1) = 0$ for $i \in \{1, \dots, 8\} \setminus \{2, 3\}$ and $\delta(t) \equiv \mathbf{0}$ for all $t \geq 2$. The magnitude*

¹Recall that $(\mathcal{V}, \mathcal{E})$ is a directed tree if the undirected graph underlying $(\mathcal{V}, \mathcal{E})$ is a tree.

²In a directed tree $(\mathcal{V}, \mathcal{E})$, $u \in \mathcal{V}$ is called a children node of $v \in \mathcal{V}$ if $\mathcal{E}_u^- \cap \mathcal{E}_v^+ \neq \emptyset$.

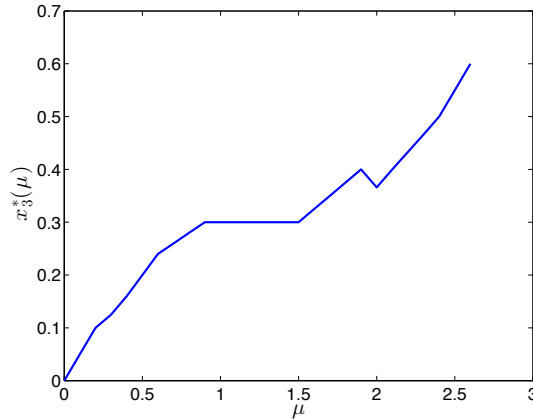


Figure 4. Plot of $x_3^*(\mu) := G_{e_3}(\mathcal{E}_1^+(0), \mu)$ vs. μ .

of such a disturbance process is 0.27, which is strictly less than $S^*(\mathcal{N}, \lambda) = 0.3$. We now describe how such a disturbance process makes the associated network flow dynamics (2)-(5) not transferring.

Under BPA routing, $f(0)$ is such that: $2 - f_2(0) = f_1(0) = 1.9$. Figure 4 then implies that $1.9 - f_4(0) = f_3(0) = 0.4$. Since $C_2(1) = 0.1 = f_2(1)$ and $C_3(1) = 0.4 = f_3(1)$, $\{e_2, e_3\} \notin \mathcal{E}(2)$. Hence $f_1(3) = 2$ and $f_4(4) = 2 = C_4$. This implies that $e_4 \notin \mathcal{E}(5)$, and hence $e_1 \notin \mathcal{E}(7)$, which leads to the dynamics being not-transferring.

On the other hand, it is easy to see that the dynamics would be transferring under this disturbance process if the routing policy at node 1 is such that $f_3(0) < 0.4$, and $f_3(0) = x_3^*(2) = 0.35$ (see Figure 4) in particular. This would correspond to the routing policy at node 1 anticipating its inflow in advance, which is not feasible under the oblivious and distributed setting for routing policies.

Example 2 suggests that the non-monotonicity in the control action of BPA routing, and hence in the evolution of flows on the links, under *point-wise* (with respect to inflow) optimization could lead to its sub optimality. This motivates consideration of the following additional constraint.

Definition 4. A distributed routing policy \mathcal{G} is called *flow-monotone* at node $v \in \mathcal{V} \setminus \{n\}$ if, for every $\mathcal{J} \subseteq \mathcal{E}_v^+$:

$$0 \leq \mu_1 \leq \mu_2 \implies G^v(\mathcal{J}, \mu_1) \leq G^v(\mathcal{J}, \mu_2), \quad (23)$$

Under a flow-monotone routing policy, if the inflow at a node increases, then the flow assigned to every active outgoing link from that node does not decrease. A routing policy which is flow monotone over all $v \in \mathcal{V} \setminus \{0, n\}$, is said to be flow monotone over \mathcal{N} . We exclude the origin node because the inflow λ at the origin node is fixed.

Remark 7. Note that, unlike the link monotonicity condition in (7), we did not include the flow monotonicity condition in (23) as part of the definition of distributed routing policies. This is because, while Example 2 illustrates that BPA routing is not necessarily flow monotone, we have not been able to find an example where link monotonicity

is violated by BPA routing with $r^* = \mathbf{0}$ in (22). However, a mathematical proof to support this observation is lacking at this point.

Under a flow monotone distributed routing policy, the network dynamics can be easily shown to possess the following simple property (which we state without proof), which simplifies the analysis considerably.

Lemma 2. *Let \mathcal{N} be a network with $|\mathcal{E}_v^+| \leq 3$ for all $v \in \mathcal{V} \setminus \{n\}$ and satisfying Assumption 2, $\lambda > 0$ a constant inflow at the origin node and BPA routing policy that is flow monotone. Then,*

$$t_1 \leq t_2 \implies f_e(t_1) \leq f_e(t_2) \quad \forall e \in \mathcal{E}(t_2).$$

The following is a key result, which, along with Theorem 1, identifies conditions under which BPA routing is maximally resilient.

Theorem 2. *Let \mathcal{N} be a network with $|\mathcal{E}_v^+| \leq 3$ for all $v \in \mathcal{V} \setminus \{n\}$ and satisfying Assumption 2, $\lambda > 0$ a constant inflow at the origin node and BPA routing policy that is flow monotone. Then, the associated network flow dynamics (2)-(5) is transferring for every disturbance process $(\delta(t))_{t \geq 1}$ with $\mathcal{D}(\delta) < S^*(\mathcal{N}, \lambda)$.*

Since BPA routing is completely specified by network parameters $(\mathcal{V}, \mathcal{E}, C)$, flow monotonicity is a condition on the network parameters. BPA routing is flow monotone at v trivially if $|\mathcal{E}_v^+| = 1$. One could perform extensive (offline) numerical tests to check flow-monotonicity of BPA routing over a given flow network \mathcal{N} . However, it is possible to identify a few flow networks over which BPA routing is provably flow-monotone. In order to characterize such networks in Proposition 4 and Remark 8 below, we need the concept of *d-expansion* of a network: given \mathcal{N} , its d-expanded version \mathcal{N}^d is obtained by creating multiple copies of the destination node in \mathcal{N} , one for each incoming link. For example, the network in Figure 5(a) is d-expanded version of the network in Figure 2 (a). It is easy to recover the original flow network from its d-expanded version.

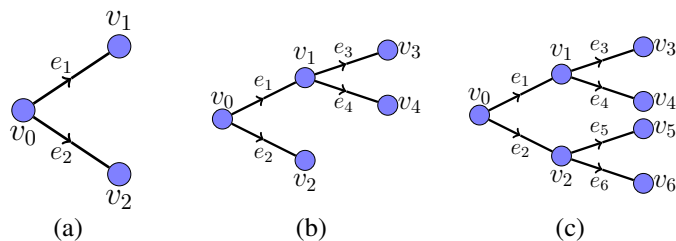


Figure 5. d-expanded flow networks that induce flow-monotonicity of BPA routing at the root node v_0 .

Proposition 4. *Let \mathcal{N} be a network satisfying Assumption 2 with $\lambda > 0$ a constant inflow at the origin node. Then, BPA routing is flow-monotone at $v \in \mathcal{V}$ if the sub-tree in the d-expanded version \mathcal{N}^d rooted at v is either (a), or (b) with $C_{e_2} \geq C_{e_1}$ or (c) with $C_{e_1} = C_{e_2}$ in Figure 5.*

Remark 8. Let $\bar{\mathcal{N}}^d$ denote the set of d -expanded versions of “simple” flow networks over which BPA routing is known to be flow monotone, either from explicit analysis as in Proposition 4, or through extensive simulations. One can use $\bar{\mathcal{N}}^d$ as a basis to form arbitrarily large networks over which BPA routing is flow monotone, using an iterative procedure as follows. Initialize \mathcal{N}^d to be an element of $\bar{\mathcal{N}}^d$. At each iteration, execute the following concatenation step. Take any destination node in \mathcal{N}^d , say v , with a single incoming link $\mathcal{E}_v^- = \{e\}$ and a member $\bar{\mathcal{N}}_i^d$ from $\bar{\mathcal{N}}^d$ whose origin node is $v_{0,i}$. If $[C_e - \mu]^+ \leq S(\mathcal{E}_{v_{0,i}}^+, \mathbf{0}, \mu)$ for all $\mu \in [0, C_e]$, then concatenate $\bar{\mathcal{N}}_i^d$ to \mathcal{N}^d at v , i.e., $v = v_{0,i}$ and the sub-network downstream of v is $\bar{\mathcal{N}}_i^d$. BPA routing on the flow network obtained at the end of every iteration of this procedure is guaranteed to be flow monotone, because (18) implies that $S_e(\mu)$ is equal to $[C_e - \mu]^+$ even in the concatenated network.

Flow monotonicity of BPA routing over a given network \mathcal{N} is maintained even after replacing any link, say $e = (u, v)$, in \mathcal{N}^d (at any iteration) with a non-branching chain e_1, \dots, e_k such that $\sigma_{e_1} = u$, $\tau_{e_m} = v$ and $\min_{i \in \{1, \dots, m\}} C_{e_i} = C_e$.

One can also devise a procedure which is counterpart to the expansion procedure described above to check if the d -expanded version of a given network \mathcal{N}^d can be decomposed into elements of $\bar{\mathcal{N}}^d$, in which case the BPA routing over \mathcal{N}^d is flow monotone. As one keeps enriching the basis $\bar{\mathcal{N}}^d$, these procedures allow to construct or to verify large networks over which BPA routing is flow monotone, and hence maximally resilient by Theorem 2.

IV. PROOFS OF THE MAIN RESULTS

In this section, we provide proofs of the main results presented in Section III. Some of the proofs rely on certain analytical properties of $S(\mathcal{J}, r, \mu)$ defined in (19), which we state first.

A. Technical Lemma

Lemma 3. Consider two networks $\mathcal{N}^1 = (\mathcal{V}, \mathcal{E}^1, C^1)$ and $\mathcal{N}^2 = (\mathcal{V}, \mathcal{E}^2, C^2)$, each satisfying Assumption 1. Let S^1 and S^2 be the functions computed by the Backward Propagation Algorithm for \mathcal{N}^1 and \mathcal{N}^2 respectively. For any $v \in \mathcal{V} \setminus \{n\}$, let \mathcal{J} be any subset of links outgoing from v and common to \mathcal{N}^1 and \mathcal{N}^2 . Then,

$$0 \leq \mu_2 \leq \mu_1, \mathbf{0} \leq r^2 \leq r^1, \mathcal{E}^1 \subseteq \mathcal{E}^2 \text{ and } C^1 \leq C^2 \implies S^1(\mathcal{K}, r^1, \mu_1) \leq S^2(\mathcal{J}, r^2, \mu_2),$$

where $\mathcal{K} := \{e \in \mathcal{J} \mid S_e^1(r_e^1) > 0\}$.

Proof: We split the lemma as follows:

- (i) $S^1(\mathcal{K}, r^1, \mu_1) = S^1(\mathcal{J}, r^1, \mu_1)$;
- (ii) $S^1(\mathcal{J}, r^1, \mu_1) \leq S^2(\mathcal{J}, r^1, \mu_1)$;
- (iii) $S^2(\mathcal{J}, r^1, \mu_1) \leq S^2(\mathcal{J}, r^1, \mu_2)$; and
- (iv) $S^2(\mathcal{J}, r^1, \mu_2) \leq S^2(\mathcal{J}, r^2, \mu_2)$.

Out of these, (iv) is trivial, and hence we omit its proof. We prove (i)-(iii) by double induction, on the number of nodes $n+1$ and the cardinality of \mathcal{J} . It is immediate to verify that the claim holds true for $n = 1$ and $|\mathcal{J}| = 1$, since

with $\mathcal{J} = \{e\}$, $S^i(e, r, \mu) = S_e^i(\mu) = [C_e^i - \mu]^+$ for all $r, i = 1, 2$. Assume that the claim is true for $|\mathcal{V}| = n + 1$ and for $|\mathcal{J}| \leq k$. Consider a \mathcal{J} of cardinality $k + 1$. (19) and (20) imply that, for all $z \in g(\mathcal{K} \cup \{j\}, r^1, \mu_1)$, $j \in \mathcal{J} \setminus \mathcal{K}$:

$$S^1(\mathcal{K} \cup \{j\}, r^1, \mu_1) \leq S_j^1(z_j) + S^1(\mathcal{K}, z, \mu_1) \leq S^1(\mathcal{K}, r^1, \mu_1), \quad (24)$$

where the second inequality follows from the fact that $z \geq r^1$ implies $0 \leq S_j^1(z_j) \leq S_j^1(r_j^1) = 0$, and using (iv) from induction. On the other hand, consider $y \in \mathcal{X}_{n+1}(\mathcal{J}, r^1, \mu_1)$ such that $y_e = r_e^1$ for all $e \in \mathcal{K}$ and $y_j = \mu_1 - \sum_{e \in \mathcal{K}} r_e^1 \geq r_j^1$, where the inequality follows from the feasibility of y . For such a y , (19) implies that there exists $i \in \mathcal{K} \cup \{j\}$ such that

$$S^1(\mathcal{K} \cup \{j\}, r^1, \mu_1) \geq S_i^1(y_i) + S^1(\mathcal{K} \cup \{j\} \setminus \{i\}, y, \mu_1) \quad (25)$$

Consider (25) under two cases. (a) $i = j$. In this case, since $S_j^1(y_j) = 0$, (25) gives $S^1(\mathcal{K} \cup \{j\}, r^1, \mu_1) \geq S^1(\mathcal{K} \cup \{j\} \setminus \{j\}, y, \mu_1) = S^1(\mathcal{K}, r^1, \mu_1)$, where the last equality follows from the fact that the components of r^1 and y along \mathcal{K} are the same. (b) $i \neq j$. In this case, recalling that $y_e = r_e^1$ for all $e \in \mathcal{K}$, and applying (i) from induction to the second term in (25), we get that

$$S^1(\mathcal{K} \cup \{j\}, r^1, \mu_1) \geq S_i^1(r_i^1) + S^1(\mathcal{K} \setminus \{i\}, y, \mu_1) = S_i^1(r_i^1) + S^1(\mathcal{K} \setminus \{i\}, r^1, \mu_1). \quad (26)$$

For every $z \in g(\mathcal{K}, r^1, \mu_1)$, we have

$$S_i^1(r_i^1) + S^1(\mathcal{K} \setminus \{i\}, r^1, \mu_1) \geq S_i^1(z_i) + S^1(\mathcal{K} \setminus \{i\}, z, \mu_1) \geq S^1(\mathcal{K}, r^1, \mu_1), \quad (27)$$

where the first inequality follows from $z \geq r^1$ and (iv) from induction. Combining (26) and (27), we arrive at the same conclusion as case (a), i.e.,

$$S^1(\mathcal{K} \cup \{j\}, r^1, \mu_1) \geq S^1(\mathcal{K}, r^1, \mu_1). \quad (28)$$

Combining (24) and (28), we establish (i) when $|\mathcal{J} \setminus \mathcal{K}| = 1$. The proof for arbitrary $|\mathcal{J} \setminus \mathcal{K}|$ follows from repetitive application of this procedure.

The proof for (ii) easily follows from induction since, for every x , and $e \in \mathcal{J}$,

$$S_e^1(x_e) + S^1(\mathcal{J} \setminus \{e\}, x, \mu_1) \leq S_e^2(x_e) + S^2(\mathcal{J} \setminus \{e\}, x, \mu_1).$$

In order to prove (iii), using the fact that $S_e^2(\mu)$ is non-increasing in μ and $S^2(\mathcal{J} \setminus \{e\}, r, \mu)$ non-increasing in r from induction, one gets that, for all r^1 such that $\mathbf{1}'r^1 \leq \mu_2$,

$$\max_{\substack{x \geq r^1 \\ \mathbf{1}'x \geq \mu_2}} \min_{e \in \mathcal{J}} \{S_e^2(x_e) + S^2(\mathcal{J} \setminus \{e\}, x, \mu_2)\} = \max_{\substack{x \geq r^1 \\ \mathbf{1}'x = \mu_2}} \min_{e \in \mathcal{J}} \{S_e^2(x_e) + S^2(\mathcal{J} \setminus \{e\}, x, \mu_2)\} = S^2(\mathcal{J}, r^1, \mu_2).$$

Specifically, to see why the first equality above holds true, let the maximum in the rightmost side be achieved in some $x \geq r^1$ such that $\mathbf{1}'x \geq \mu_2$, and let $\theta \in [0, 1]$ be such that $(1 - \theta)\mathbf{1}'x + \theta\mathbf{1}'r^1 = \mu_2$ (such θ exists since $\mathbf{1}'x \geq \mu_2$ and $\mathbf{1}'r^1 \leq \mu_2$). Then, $y := (1 - \theta)x + \theta r^1$ satisfies $x \geq y \geq r^1$, $\mathbf{1}'y = \mu_2$, and $S_e^2(y_e) + S^2(\mathcal{J} \setminus \{e\}, y, \mu_2) \geq$

$S_e^2(x_e) + S^2(\mathcal{J} \setminus \{e\}, x, \mu_2)$ for all e . Then, for $\mu_1 \geq \mu_2$,

$$\begin{aligned} S^2(\mathcal{J}, r^1, \mu_1) &= \max_{\substack{x \geq r^1 \\ \mathbf{1}'x = \mu_1}} \min_{e \in \mathcal{J}} \{S_e^2(x_e) + S^2(\mathcal{J} \setminus \{e\}, x, \mu_1)\} \\ &\leq \max_{\substack{x \geq r^1 \\ \mathbf{1}'x \geq \mu_2}} \min_{e \in \mathcal{J}} \{S_e^2(x_e) + S^2(\mathcal{J} \setminus \{e\}, x, \mu_2)\} = S^2(\mathcal{J}, r^1, \mu_2). \end{aligned}$$

This concludes the proof for $|\mathcal{V}| = n + 1$ and $|\mathcal{J}| \leq k + 1$. A similar argument allows one to extend the validity of the result to $|\mathcal{V}| = n + 2$ with $|\mathcal{J}| = 1$. The lemma then follows by induction. \blacksquare

B. Proof of Proposition 1

We first show that, for any $\lambda \geq 0$, $\mathcal{K} \subseteq \mathcal{J} \subseteq \mathcal{E}$,

$$S(\mathcal{K}, \lambda) \leq S(\mathcal{J}, \lambda). \quad (29)$$

It suffices to show that $S(\mathcal{J} \setminus \{e\}, \lambda) \leq S(\mathcal{J}, \lambda)$ for all $e \in \mathcal{J}$. It is trivially true for $|\mathcal{J}| = 1$. Assume it to be true for all $|\mathcal{J}| \leq k$ for some $k \geq 1$. (12) implies that, for all $x \in \mathcal{X}(\mathcal{J}, \lambda)$:

$$\begin{aligned} S(\mathcal{J}, \lambda) &\geq \min_{j \in \mathcal{J}} (C_j - x_j + S(\mathcal{J} \setminus \{j\}, \lambda)) \\ &= \min \left(C_e - x_e + S(\mathcal{J} \setminus \{e\}, \lambda), \min_{j \in \mathcal{J} \setminus \{e\}} (C_j - x_j + S(\mathcal{J} \setminus \{j\}, \lambda)) \right) \\ &\geq \min \left(S(\mathcal{J} \setminus \{e\}, \lambda), \min_{j \in \mathcal{J} \setminus \{e\}} (C_j - x_j + S(\mathcal{J} \setminus \{j\}, \lambda)) \right) \\ &\geq \min \left(S(\mathcal{J} \setminus \{e\}, \lambda), \min_{j \in \mathcal{J} \setminus \{e\}} (C_j - x_j + S(\mathcal{J} \setminus \{j, e\}, \lambda)) \right), \end{aligned} \quad (30)$$

where the second inequality follows the fact that $x_e \leq C_e$, whereas the third inequality follows from the inductive argument on $\mathcal{J} \setminus \{i\}$. For $x \in \arg\max_{z \in \mathcal{X}(\mathcal{J} \setminus \{e\}, \lambda)} \min_{j \in \mathcal{J} \setminus \{e\}} (C_j - z_j + S(\mathcal{J} \setminus \{j, e\}, \lambda))$ we get $S(\mathcal{J} \setminus \{e\}, \lambda) = \min_{j \in \mathcal{J} \setminus \{e\}} (C_j - x_j + S(\mathcal{J} \setminus \{j, e\}, \lambda))$, which when used in (30), finishes the proof for $|\mathcal{J}| = k + 1$. By induction, (29) is then true for all $\mathcal{J} \subseteq \mathcal{E}$.

We now prove the proposition by induction on $|\mathcal{J}|$. When $\mathcal{J} = \{e\}$, (12) implies that $S(\mathcal{J}, \lambda) \leq [C_e - \lambda]^+$. It is easy to see that under a disturbance process $\delta_e(1) = [C_e - \lambda]^+$ and $\delta_e(t) = 0$ for all $t \geq 1$, the associated network dynamics will be non-transferring. Assume that the proposition is true for all $\mathcal{J} \subseteq \mathcal{E}$ with $|\mathcal{J}| \leq k$ for some $k \geq 1$. Let $\mathcal{E}(0) = \mathcal{J}$, with $|\mathcal{J}| = k + 1$. Pick $e \in \arg\min_{j \in \mathcal{J}} (C_j - f_j(0) + S(\mathcal{J} \setminus \{j\}, \lambda))$. Therefore,

$$S(\mathcal{J}, \lambda) \geq C_e - f_e(0) + S(\mathcal{J} \setminus \{e\}, \lambda). \quad (31)$$

Consider a disturbance process such that $\delta_e(1) = C_e - f_e(0)$ and $\delta_j(1) = 0$ for all $j \in \mathcal{J} \setminus \{e\}$. Under this disturbance, link e becomes inactive, followed by a possible cascading failure. Let the network state come to a steady state after a finite time \mathcal{T}_1 . Since, $|\mathcal{E}(\mathcal{T}_1)| \leq k$, one can use induction to extend δ after \mathcal{T}_1 to ensure that the network dynamics is not transferring. By induction, the total magnitude of δ is then upper bounded as $\mathcal{D}(\delta) \leq \delta_e(1) + S(\mathcal{E}(\mathcal{T}_1), \lambda)$. Since $\mathcal{E}(\mathcal{T}_1) \subsetneq \mathcal{J}$, using (29), this can be further upper bounded as $\mathcal{D}(\delta) \leq \delta_e(1) + S(\mathcal{J} \setminus \{e\}, \lambda) = C_e - f_e(0) + S(\mathcal{J} \setminus \{e\}, \lambda)$, which combined with (31) implies that $\mathcal{D}(\delta) \leq S(\mathcal{J}, \lambda)$.

C. Proof of Proposition 2

If $0 \notin \mathcal{V}(\mathcal{T})$, then the flow across links outgoing from every cut \mathcal{U} in $(\mathcal{V}(\mathcal{T}), \mathcal{E}(\mathcal{T}))$, and $\mathcal{V}(\mathcal{T}) \setminus \{n\}$ in particular, is zero at \mathcal{T} . This implies that $\lambda_n(\mathcal{T}) = 0$. This also proves that $\lambda_n(\mathcal{T}) \in \{0, \lambda\}$.

If $0 \in \mathcal{V}(\mathcal{T})$, then $\mathcal{J} := \mathcal{E}(\mathcal{T}) \cap \mathcal{E}_0^+$ is non-empty, and $\sum_{e \in \mathcal{J}} f_e(\mathcal{T}) = \lambda$. Let $\mathcal{U} = 0 \cup \{v \in \mathcal{V}(\mathcal{T}) : (0, v) \in \mathcal{J}\}$. It is clear that $\mathcal{U} \subseteq \mathcal{V}(\mathcal{T})$, and that \mathcal{U} is in fact a cut in $(\mathcal{V}(\mathcal{T}), \mathcal{E}(\mathcal{T}))$. Since $f(\mathcal{T})$ is an equilibrium flow, the total flow across the links outgoing from \mathcal{U} is λ . One can continue along these lines to claim that the flow across links outgoing from any cut in $(\mathcal{V}(\mathcal{T}), \mathcal{E}(\mathcal{T}))$, and $\mathcal{V}(\mathcal{T}) \setminus \{n\}$ in particular, is equal to λ . That is, $\lambda_n(\mathcal{T}) = \lambda$.

D. Proof of Theorem 1

Theorem 1 is a corollary of the following lemma, where we allow the possibility that $(\mathcal{V}(0), \mathcal{E}(0)) \neq (\mathcal{V}, \mathcal{E})$.

Lemma 4. *Consider a node v in a network \mathcal{N} with initial condition $(\mathcal{V}(0), \mathcal{E}(0)) \subseteq (\mathcal{V}, \mathcal{E})$, satisfying Assumption 1, with a constant inflow $\mu \geq 0$, and operating under a distributed routing policy satisfying $G^v(\mathcal{E}_v^+(0), \mu) \geq r$ for some $r \in \mathbb{R}_+^{\mathcal{E}_v^+}$. Then, for any $h \in \mathbb{N}$, there exists a finite $\mathcal{T}_v \geq h$, and a disturbance process $(\delta^v(t))_{t \geq h}$ satisfying $\mathcal{D}(\delta^v) \leq S(\mathcal{E}_v^+(0), r, \mu)$, under which $v \notin \mathcal{V}(\mathcal{T}_v)$.*

Proof: It is sufficient to prove the lemma for $h = 1$. For brevity in notation, we let $f(t) := G^v(\mathcal{E}_v^+(t), \mu)$ be the action of the control policy at $t \geq 0$. By assumption and the link monotonicity property of routing policy in (7), $f(t_2) \geq f(t_1) \geq r$ for all $t_2 \geq t_1 \geq 0$. We follow the convention that, unless specified otherwise, $\delta_e^v(t) = 0$ for all $e \in \mathcal{E}$ and $t \geq 1$.

The proof is by double induction, on the number of nodes $n+1$ and the cardinality of $|\mathcal{E}_v^+(0)|$. The proof is easy to verify when $n = 1$ and $|\mathcal{E}_v^+(0)| = 1$, since in that case, with $\mathcal{E}_v^+(0) = \{e\}$, $S_e(\mu) = [C_e - \mu]^+$, and therefore, one can apply $\delta_e^v(1) = [C_e - \mu]^+$, under which $v \notin \mathcal{V}(2)$ and $\mathcal{D}(\delta^v) = [C_e - \mu]^+$. Assume the lemma to be true for arbitrary acyclic networks with $n+1$ nodes and $|\mathcal{E}_v^+(0)| \leq k$ for some $k \geq 1$.

For $|\mathcal{E}_v^+(0)| = k+1$, pick e in $\operatorname{argmin}_{j \in \mathcal{E}_v^+(0)} ([C_j - f_j(0)]^+ + S(\mathcal{E}_v^+(0) \setminus \{j\}, y, \mu))$. Consider a disturbance process δ^v such that $\delta_e^v(1) = [C_e - f_e(0)]^+$, in which case

$$S(\mathcal{E}_v^+(0), r, \mu) \geq [C_e - y_e]^+ + S(\mathcal{E}_v^+(0) \setminus \{e\}, f(0), \mu). \quad (32)$$

Let the times at which links fail simultaneously be $2 = t_1 \leq \dots \leq t_m$. Let $S^t(\mathcal{J}, r, \mu)$ denote the functions computed by the BPA in (19) for the residual graph $(\mathcal{V}(t), \mathcal{E}(t), C(t))$ at $t \geq 0$. By convention, we set $S^0 \equiv S$. Since links in $\mathcal{E}_v^+(t_{i-1}) \setminus \mathcal{E}_v^+(t_i)$ fail simultaneously at t_i , $S_j^{t_i-1}(f_j(t_i-1)) = 0$ for all $j \in \mathcal{E}_v^+(t_{i-1}) \setminus \mathcal{E}_v^+(t_i)$. Therefore, using Lemma 3, for all $i \in \{2, \dots, m\}$:

$$\begin{aligned} S^{t_i}(\mathcal{E}_v^+(t_i), f(t_i), \mu) &\leq S^{t_i}(\mathcal{E}_v^+(t_i), f(t_i-1), \mu) \leq S^{t_i-1}(\mathcal{E}_v^+(t_i), f(t_i-1), \mu) \\ &= S^{t_i-1}(\mathcal{E}_v^+(t_i-1), f(t_i-1), \mu) \leq S^{t_i-1}(\mathcal{E}_v^+(t_{i-1}), f(t_{i-1}), \mu), \end{aligned} \quad (33)$$

where we have used the fact that $f(t_i) \geq f(t_i-1) = f(t_{i-1})$ and $\mathcal{E}(t_i-1) \subseteq \mathcal{E}(t_{i-1})$. Using the same arguments, since $S_e^1(f_1(1)) = 0$, we have that

$$S^{t_1}(\mathcal{E}_v^+(t_1), f(t_1), \mu) \leq S^1(\mathcal{E}_v^+(1), f(1), \mu) = S^1(\mathcal{E}_v^+(0) \setminus \{e\}, f(1), \mu) \leq S(\mathcal{E}_v^+(0) \setminus \{e\}, r, \mu). \quad (34)$$

Combining (33) and (34), we get that

$$S^{t_m}(\mathcal{E}_v^+(t_m), f(t_m), \mu) \leq S(\mathcal{E}_v^+(0) \setminus \{e\}, r, \mu). \quad (35)$$

Using induction on the residual graph at $t = t_m$, where $|\mathcal{E}_v^+(t_m)| < |\mathcal{E}_v^+(0)|$, one can construct a disturbance process $(\tilde{\delta}^v(t))_{t \geq t_m}$ such that $v \notin \mathcal{V}(\mathcal{T})$, and

$$D(\tilde{\delta}^v) \leq S^{t_m}(\mathcal{E}_v^+(t_m), f(t_m), \mu). \quad (36)$$

Augmenting δ^v with $\tilde{\delta}^v$, i.e., $(\delta^v(t))_{t \geq t_m} = (\tilde{\delta}^v(t))_{t \geq t_m}$, and applying (32), (35) and (36), we get that

$$\mathcal{D}(\delta^v) = [C_e - f_e(0)]^+ + \mathcal{D}(\tilde{\delta}^v) \leq S(\mathcal{E}_v^+(0), f(0), \mu) \leq S(\mathcal{E}_v^+(0), r, \mu),$$

which proves the lemma for acyclic networks with $n + 1$ nodes and $|\mathcal{E}_v^+(0)| = k + 1$. The proof can be easily extended to acyclic networks with $n + 2$ nodes and $|\mathcal{E}_v^+(0)| = k$, after which the lemma follows from induction. ■

Theorem 1 follows by combining Lemma 4 for $v = 0$, $h = 1$, $\mu = \lambda$, $r = \mathbf{0}$ and $(\mathcal{V}(0), \mathcal{E}(0)) = (\mathcal{V}, \mathcal{E})$ with Proposition 2.

E. Proof of Proposition 3 and Theorem 2

Proposition 3 follows from Theorem 2 by recalling that, for symmetric flow networks that are directed trees, BPA routing satisfies flow monotonicity. Theorem 2 follows from the following lemmas. The following simple property of the functions $S(\mathcal{J}, r, \mu)$ computed in (19) will be useful in the proofs in this subsection. We recall the definition of $g(\mathcal{J}, r, \mu)$ from (20).

Lemma 5. For any $x \in g(\mathcal{J}, r, \mu)$, $\mathcal{J} \subseteq \mathcal{E}_v^+$, $r \in \mathbb{R}_+^{\mathcal{E}_v^+}$, $v \in \mathcal{V} \setminus \{n\}$ and $\mu \geq 0$,

$$S(\mathcal{J}, r, \mu) \leq \sum_{e \in \mathcal{K}} S_e(x_e) + S(\mathcal{J} \setminus \mathcal{K}, x, \mu) \quad \forall \mathcal{K} \subseteq \mathcal{J}.$$

Proof: The lemma is trivially true from (19) and (20) for $|\mathcal{K}| = 1$. Assume it to be true for all $\mathcal{K} \subset \mathcal{J}$ with $|\mathcal{K}| \leq k$ for some $k \geq 1$. Consider the case $|\mathcal{K}| = k + 1$. Induction implies that, for any $j \in \mathcal{K}$:

$$S(\mathcal{J}, r, \mu) \leq \sum_{e \in \mathcal{K} \setminus \{j\}} S_e(x_e) + S(\mathcal{J} \cup \{j\} \setminus \mathcal{K}, x, \mu). \quad (37)$$

Applying induction to the second term in (37), for every $z \in g(\mathcal{J} \cup \{j\} \setminus \mathcal{K}, x, \mu)$ we get

$$S(\mathcal{J} \cup \{j\} \setminus \mathcal{K}, x, \mu) \leq S_j(z_j) + S(\mathcal{J} \setminus \mathcal{K}, z, \mu). \quad (38)$$

$z \in g(\mathcal{J} \cup \{j\} \setminus \mathcal{K}, x, \mu)$ implies $z \geq x$, which in turn implies $S_j(z_j) \leq S_j(x_j)$ and $S(\mathcal{J} \setminus \mathcal{K}, z, \mu) \leq S(\mathcal{J} \setminus \mathcal{K}, x, \mu)$ by Lemma 3. Combining this with (37) and (38) establishes the lemma for $|\mathcal{K}| = k + 1$, and hence for all $\mathcal{K} \subseteq \mathcal{J}$ by induction. ■

For the next lemma, we again allow the possibility that $(\mathcal{V}(0), \mathcal{E}(0)) \neq (\mathcal{V}, \mathcal{E})$.

Lemma 6. Consider a node $v \in \mathcal{V} \setminus \{n\}$ with $|\mathcal{E}_v^+| \leq 3$ in a flow network \mathcal{N} with initial condition $(\mathcal{V}(0), \mathcal{E}(0)) \subseteq (\mathcal{V}, \mathcal{E})$, satisfying Assumption 2, and operating under BPA routing. Let the inflow $\lambda_v(t)$ be non-decreasing and

satisfy $\max_{t \geq 0} \lambda_v(t) = \mu \geq 0$. If $\lambda_v(t) \equiv \mu$, or if BPA routing at v is flow monotone, then $v \in \mathcal{V}(\mathcal{T})$ under any disturbance process δ^v satisfying $\mathcal{D}(\delta^v) < S(\mathcal{E}_v^+(0), r^*, \mu)$, with $r^* = G^v(\mathcal{E}_v^+, \mu)$.

Proof: Let $\mathcal{E}_v^+ = \{e_1, e_2, e_3\}$ and $f_i^* := G_{e_i}^v(\mathcal{E}_v^+(0), \mu)$ for all $e_i \in \mathcal{E}_v^+(0)$. We consider three possible scenarios for $|\mathcal{E}_v^+(0)|$ separately. We prove by backward induction on v in $\{0, \dots, n-1\}$. First consider $v = n-1$. When $\mathcal{E}_v^+(0) = \{e_1\}$, (18) implies that $S_{e_1}(\mu) = C_1 - \mu \leq C_1 - \lambda_v(t)$. Therefore, for all $t \geq 0$, $\Delta_{e_1}(t) \leq \Delta_{e_1}(\mathcal{T}) \leq \mathcal{D}(\delta^v) < S_{e_1}(\mu) = C_1 - \mu \leq C_1 - \lambda_v(t)$. That is, $f_{e_1}(t) = \lambda_v(t) < C_1 - \Delta_{e_1}(t) = C_1(t)$ for all $t \geq 0$, and hence $v \in \mathcal{V}(\mathcal{T})$.

When $\mathcal{E}_v^+(0) = \{e_1, e_2\}$, let $t_1 := \min\{t \geq 0 \mid \mathcal{E}_v^+(t) \neq \mathcal{E}_v^+(0)\}$. To avoid triviality, assume $t_1 < \infty$, and let $\mathcal{K} := \mathcal{E}_v^+(0) \setminus \mathcal{E}_v^+(t_1)$ be the set of links to become inactive simultaneously at t_1 . Then, necessarily $\Delta_e(t_1) \geq [C_e - f_e(t_1)]^+ \geq [C_e - f_e^*]^+$ for all $e \in \mathcal{K}$, where the second inequality follows from flow monotonicity. Therefore,

$$\sum_{e \in \mathcal{K}} \Delta_e(\mathcal{T}) \geq \sum_{e \in \mathcal{K}} \Delta_e(t_1) \geq \sum_{e \in \mathcal{K}} [C_e - f_e^*]^+. \quad (39)$$

Let $\tilde{\delta}^{\mathcal{K}}$ be such that $\tilde{\delta}_e^{\mathcal{K}}(t) \equiv 0$ for all $e \in \mathcal{K}$ and $\tilde{\delta}_e^{\mathcal{K}}(t) \equiv \delta_e^v(t)$ for all $e \in \mathcal{E}_v^+(t_1)$. Therefore,

$$\mathcal{D}(\delta^v) = \mathcal{D}(\tilde{\delta}^{\mathcal{K}}) + \sum_{e \in \mathcal{K}} \Delta_e(\mathcal{T}). \quad (40)$$

Combining (39) with Lemma 4, where we note that $f^* := G^v(\mathcal{E}_v^+(0), \mu) \in g(\mathcal{E}_v^+(0), r^*, \mu)$ (from (22)), we get

$$S(\mathcal{E}_v^+(0), r^*, \mu) \leq \sum_{e \in \mathcal{K}} \Delta_e(\mathcal{T}) + S(\mathcal{E}_v^+(t_1), f^*, \mu) \quad (41)$$

Combining (41) and (40) with $\mathcal{D}(\delta^v) < S(\mathcal{E}_v^+(0), r^*, \mu)$, we get

$$\mathcal{D}(\tilde{\delta}^{\mathcal{K}}) < S(\mathcal{E}_v^+(t_1), f^*, \mu). \quad (42)$$

If $|\mathcal{E}_v^+(t_1)| = 0$, then (42) is a contradiction, and if $|\mathcal{E}_v^+(t_1)| = 1$, the proof is then completed by using the $|\mathcal{E}_v^+(0)| = 1$ case since $S(\mathcal{E}_v^+(t_1), f^*, \mu) = S(\mathcal{E}_v^+(t_1), \mathbf{0}, \mu)$.

For $\mathcal{E}_v^+(0) = \{e_1, e_2, e_3\}$, one follows the same argument as before to arrive at (42). If $|\mathcal{E}_v^+(t_1)| \leq 1$, then we use the same arguments as before. If $|\mathcal{E}_v^+(t_1)| = 2$, then necessarily $\mathcal{E}_v^+(0) = \mathcal{E}_v^+$, in which case $f^* = r^*$. Therefore, one can continue with the $|\mathcal{E}_v^+(0)| = 2$ case to complete the proof. This proves the lemma for $v = n-1$.

Assume that the lemma is true for all $v \geq \ell$ for some $\ell \geq 1$. Let $v = \ell - 1$. We recall Lemma 2 for monotonicity of $f_i(t)$. As for the $v = n-1$, we consider three cases depending on the value of $|\mathcal{E}_v^+(0)|$ and provide proof using similar arguments. We provide a few details only for $\mathcal{E}_v^+(0) = \mathcal{E}_v^+ = \{e_1, e_2, e_3\}$. Lemma 2 implies that $f_i(t) \leq f_i^* = G_{e_i}^v(\mathcal{E}_v^+, \mu)$ for all $i = 1, 2, 3$. Let $\delta^v = \sum_{i=0}^3 \delta_i^v$, where, for $i = 1, 2, 3$, δ_i^v is the component of δ^v on links consisting of e_i and the sub-tree rooted at τ_{e_i} , and δ_0^v is the component on the rest of the links in the network. Let $\Delta_i(t)$, $i = 0, 1, 2, 3$, be defined accordingly. Let $t_1 := \min\{t \geq 0 \mid \mathcal{E}_v^+(t) \neq \mathcal{E}_v^+\}$. To avoid triviality, assume $t_1 < \infty$, and let $\mathcal{K} := \mathcal{E}_v^+ \setminus \mathcal{E}_v^+(t_1)$ be the set of links to become inactive simultaneously at t_1 . Then, necessarily $\Delta_e(t_1) \geq S_e(f_e(t_1)) \geq S_e(f_e^*)$ for all $e \in \mathcal{K}$, where the second inequality follows from Lemma 3. Following similar arguments as before, we arrive at (42), after which we use the relevant case depending on the value of

$|\mathcal{E}_v^+(t_1)| \leq 2$. This establishes the proof for $v = \ell - 1$, and hence by backward induction for all $v \in \{0, \dots, n - 1\}$. \blacksquare

Theorem 2 is obtained from Lemma 6 by substituting $v = 0$, $(\mathcal{V}(0), \mathcal{E}(0)) = (\mathcal{V}, \mathcal{E})$, $\lambda_v(t) \equiv \lambda$, and noting that $S(\mathcal{E}_0^+, r^*, \lambda) = S(\mathcal{E}_0^+, \mathbf{0}, \lambda) = S^*(\mathcal{N}, \lambda)$.

F. Proof of Proposition 4

The BPA routing for case (a) is explicitly computed in (15), which readily implies flow-monotonicity. The proof for case (c) follows from the constructs used in the proof of case (b). Therefore, we provide details only for case (b).

For brevity in notation, let $y(\mu) \equiv G^{v_0}(\mathcal{E}_{v_0}^+, \mu)$ be the flow under BPA routing. For brevity in notation, and since the lower bound constraints imposed by r are redundant in this case, we drop the dependence of $S(\cdot, r, \cdot)$ on r . Following Figure 3(b), the general relationship between $S_{e_1}(\mu)$ and $S(\mathcal{E}_{v_1}^+, \mu)$ can be written as:

$$S_{e_1}(\mu) = \begin{cases} S(\mathcal{E}_{v_1}^+, \mu) & \text{if } \mu \in [\bar{\mu}_1, \bar{\mu}_2], \\ [C_1 - \mu]^+ & \text{if } \mu \in [0, \bar{\mu}_1] \cup [\bar{\mu}_2, C_1], \end{cases} \quad (43)$$

where $\bar{\mu}_1 = 2(C_3 + C_4 - C_1)$, $\bar{\mu}_2 = 2C_1 - (C_3 + C_4)$ and one can write an expression for $S(\mathcal{E}_{v_1}^+, \mu)$ similar to (14). We prove flow-monotonicity by showing that $\frac{d}{d\mu}y(\mu) \geq 0$. Let $\Gamma_i = \sup\{\mu : S_{e_i}(\mu) > 0\}$ for $i = 1, 2$ be the effective flow capacity of link e_i . When $y(\mu)$ is on the boundary of the feasible set $\mathcal{X}_{v_0}(\mathcal{E}_{v_0}^+, \mu)$, without loss of generality, assume that $y_1(\mu) = [\mu - \Gamma_2]^+$ and $y_2(\mu) = \min\{\mu, \Gamma_2\}$, which is trivially flow-monotone. When $y(\mu)$ is in the interior of the feasible set, y_2 satisfies $S_{e_2}(y_2) + S_{e_1}(\mu) = S_{e_1}(\mu - y_2) + S_{e_2}(\mu)$. Therefore, by the implicit function theorem,³ we have that

$$\frac{d}{d\mu}y_2(\mu) = \frac{S'_{e_2}(\mu) + S'_{e_1}(\mu - y_2) - S'_{e_1}(\mu)}{S'_{e_2}(y_2) + S'_{e_1}(\mu - y_2)}, \quad (44)$$

where $S'_{e_i}(y) \equiv \frac{d}{dy}S_{e_i}(y)$, $i = 1, 2$. The strictly decreasing property of S_{e_i} , $i = 1, 2$ from Lemma 3 implies that the denominator of (44) is negative for all $\mu \in (0, \Gamma_1 + \Gamma_2)$. (14) and (43) imply that $S'_{e_1}(\mu - y_2) - S'_{e_1}(\mu) \leq 1/2$ if $\mu < C_1$ and equal to zero if $\mu > C_1$. This combined with $S_{e_2}(\mu) \equiv [C_2 - \mu]^+$ and the assumption that $C_2 \geq C_1$ implies that the numerator of (44) is non positive for all $\mu < \Gamma_1 + \Gamma_2$. Hence, $\frac{d}{d\mu}y_2(\mu) \geq 0$. The proof for $\frac{d}{d\mu}y_1(\mu) \geq 0$ follows along similar lines.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a dynamical model for cascading failures in single-commodity network flows, where the network dynamics is governed by a deterministic and possibly adversarial disturbance process which incrementally reduces flow capacity on the links, and distributed oblivious routing policies that have information only about the

³ $S_{e_1}(\mu)$ and $S_{e_2}(\mu)$ are continuous piecewise linear functions, e.g., see Figure 3 (c). Therefore, one can find analytic functions that can approximate $S_{e_1}(\mu)$ and $S_{e_2}(\mu)$, as well as their derivatives arbitrarily closely at almost all points. We implicitly assume that $S_{e_1}(\mu)$ and $S_{e_2}(\mu)$ are replaced with such analytic approximations.

local inflow and active status of outgoing links, and in particular no information about the disturbance process. The salient feature of this model is to couple the flow dynamics with the link and node inactivation dynamics. An immediate outcome of this coupling is that, links and nodes to fail successively are not necessarily adjacent to each other. We quantified margin of resilience to be the minimum cumulative capacity reductions across time and links of the network, under which the network loses its transferring property. We presented an algorithm that provides an upper bound on the margin of resilience for directed acyclic graphs between a single origin-destination pair. The same algorithm motivates a routing policy which provably matches the upper bound for networks which are tree like, have out-degree at most 3, and induce monotonicity in the flow dynamics.

In future, we plan to extend our analysis to networks with general graph topologies, multi-commodity flows, non-oblivious routing policies with possibly multi hop information, stochastic disturbance processes, reversible link activation dynamics under finite time link recovery, and exogenous coupling between failure and recovery of distant links due to coupling between the given network and other exogeneous networks. We also plan to investigate computationally efficient, and possibly distributed, algorithms for (approximate) computation of the margin of resilience. Finally, we plan to develop distributed robust control strategies suitable for physics of other physical infrastructure networks such as power, gas, water and supply chains.

ACKNOWLEDGMENTS

Ketan Savla thanks Qin Ba for help with the simulations in the paper. Giacomo Como was partially supported by the Swedish Research Council through the Junior Research Grant Information Dynamics in Large Scale Networks and the Excellence Centers LCCC and ELLIT.

REFERENCES

- [1] R. Cohen, K. Erez, D. Ben-Avraham, and S. Havlin, “Resilience of the internet to random breakdowns,” *Physical review letters*, vol. 85, no. 21, p. 4626, 2000.
- [2] D. J. Watts, “A simple model of global cascades on random networks,” *PNAS*, vol. 99, no. 9, pp. 5766–5771, 2002.
- [3] A. Motter, “Cascade-based attacks on complex networks,” *Phys. Rev. E; Physical Review E*, vol. 66, no. 6, 2002.
- [4] P. Crucitti, V. Latora, and M. Marchiori, “Model for cascading failures in complex networks,” *Physical Review E*, vol. 69, no. 4, 2004.
- [5] A. Barrat, M. Barthelemy, and A. Vespignani, *Dynamical Processes on Complex Networks*. Cambridge University Press, 2008.
- [6] T. Liggett, *Interacting particle systems*. Springer-Verlag, 1985.
- [7] G. Grimmett, *Percolation*. Springer, 1999.
- [8] J. Adler, “Bootstrap percolation,” *Physica A: Statistical Mechanics and its Applications*, vol. 171, no. 3, pp. 453–470, 1991.
- [9] F. Vega-Redondo, *Complex social networks*. Cambridge University Press, 2006.
- [10] M. Draief and L. Massoulié, *Epidemics and rumors in complex networks*. Cambridge University Press, 2010.
- [11] M. Lelarge, “Diffusion and cascading behavior in random networkx,” *Games and Economic Behavior*, vol. 75, pp. 752–775, 2012.
- [12] Z. Kong and E. M. Yeh, “Resilience to degree-dependent and cascading node failures in random geometric networks,” *Information Theory, IEEE Transactions on*, vol. 56, no. 11, pp. 5533–5546, 2010.
- [13] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman, “Power grid vulnerability to geographically correlated failures – analysis and control implications,” Columbia University, Electrical Engineering, <http://arxiv.org/pdf/1206.1099v1.pdf>, Tech. Rep., 2011.

- [14] R. Baldick, B. Chowdhury, I. Dobson, Z. Dong, B. Gou, D. Hawkins, H. Huang, M. Joung, D. Kirschen, F. Li *et al.*, “Initial review of methods for cascading failure analysis in electric power transmission systems IEEE PES CAMS task force on understanding, prediction, mitigation and restoration of cascading failures,” in *Power and Energy Society General Meeting-Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE*. IEEE, 2008, pp. 1–8.
- [15] D. Bienstock, “Optimal control of cascading power grid failures,” in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*. IEEE, 2011, pp. 2166–2173.
- [16] C. Lai and S. H. Low, “The redistribution of power flow in cascading failures,” in *51st Annual Allerton Conference on Communication, Control, and Computing*, 2013, pp. 1037–1044.
- [17] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [18] G. Como, K. Savla, D. Acemoglu, M. A. Dahleh, and E. Frazzoli, “Robust distributed routing in dynamical networks – part I: Locally responsive policies and weak resilience,” *IEEE Trans. on Automatic Control*, vol. 58, no. 2, pp. 317–332, 2013.
- [19] —, “Robust distributed routing in dynamical networks – part II: Strong resilience, equilibrium selection and cascaded failures,” *IEEE Trans. on Automatic Control*, vol. 58, no. 2, pp. 333–348, 2013.
- [20] G. Como, E. Lovisari, and K. Savla, “Throughput optimality and overload behavior of dynamical flow networks under monotone distributed routing,” *IEEE Transactions on Control of Networked Systems*, 2014, under revision. Available at <http://arxiv.org/abs/1308.1993>.
- [21] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 2nd ed. MIT Press, 2001.



Ketan Savla (M’12) is an Assistant Professor at the Sonny Astani Department of Civil and Environmental Engineering, University of Southern California, Los Angeles, CA. He received his Ph.D. in Electrical Engineering from the University of California, Santa Barbara in 2007. From 2007 to 2012, he was with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology, first as a postdoctoral associate and then as a research scientist. His current research interests are in systems and control, dynamical networks, optimization, multi-agent systems and queueing systems.



Giacomo Como (M’12) is an Associate Professor at the Department of Automatic Control, Lund University, Sweden. He received the B.Sc., M.S., and Ph.D. degrees in Applied Mathematics from Politecnico di Torino, Italy, in 2002, 2004, and 2008, respectively. In 2006–2007, he was a Visiting Assistant in Research at the Department of Electrical Engineering, Yale University. From 2008 to 2011, he was a Postdoctoral Associate at the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology. His current research interests are in control, information, game theory, and networks.



Munther A. Dahleh (S'84–M'87–SM'97–F'01) was born in 1962. He received the B.S. degree in electrical engineering from Texas A&M University, College Station, in 1983 and the Ph.D. degree in electrical engineering from Rice University, Houston, TX, in 1987. Since then, he has been with the Department of Electrical Engineering and Computer Science, MIT, Cambridge, where he is currently a Full Professor. He is the coauthor (with I. Diaz-Bobillo) of the book *Control of Uncertain Systems: A Linear Programming Approach* (Prentice-Hall, 1995) and the coauthor (with N. Elia) of the book *Computational Methods for Controller Design* (Springer, 1998). In addition to methodology development, he has been interested in the application of distributed control in the future electric grid and the future transportation system. He is also interested in various problems in network science and in the interface between systems theory and neurobiology, and, in particular, in providing an anatomically consistent model of the motor control. His research interests include problems at the interface of robust control, filtering, information theory, and computation.