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 give sufficient conditions on network parameters under which the upper bound is tight under an appropriate routing policy.**

I. Introduction

Resilience is becoming a key consideration in the design and operation of many critical infrastructure systems such as transportation, power, water, 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. 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]. 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, [4], [5]. 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 [6], [7], 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., [8]) Recently, more physically motivated dynamical models for cascading failures overcoming such limitations have been proposed and analyzed in the context of power networks [9], [10], even allowing for control in between successive failure events [11].

This paper is concerned with dynamical model for cascading failures in single-commodity flow networks, and with the characterization of maximally resilient routing policies. 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. Due to space limitations, we refer to [12] for missing proofs and technical details.

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 nonneg-

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

ative 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 stay 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 flow networks 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 [13], 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

$$\max_{\mathcal{U}} \{\lambda_{\mathcal{U}} - C_{\mathcal{U}}\} < 0, \tag{1}$$

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 $t = 0, 1, \ldots$, 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 $t = 0, 1, 2, \ldots$, 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.,

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

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

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

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

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

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

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

The sequence $(\delta(1), \delta(2), \ldots) \subseteq \mathbb{R}_+^{\mathcal{E}}$ 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

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

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

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

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

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

In reading (7), recall our notation established at the end of Section I that, for $x \in \mathbb{R}^{\mathcal{A}}$ and $y \in \mathbb{R}_+^{\mathcal{B}}$, $x \leq y$ implies $x_i \leq y_i$ 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 6.

Remark 1: (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 one which assigns flow proportional to links capacities.

The model in (2)-(5) has several salient features. 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, 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. This is in contrast to the setting of power networks, where the time scale for flow dynamics is much faster than the link failure dynamics and control action.

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 $C_i = 4$ for $i = 1, 2$, $C_i = 3$ for $i = 3, 4, 6, 7, 10$, $C_i = 1.5$ for $i = 5, 9$ and $C_8 = 0.75$. Let the arrival rate at the origin be $\lambda = 4$. We consider proportional routing policies at all the nodes, under which the initial flow on all links are given by $f_i(0) = 2$ for $i = 1, 2, 10$, $f_i(0) = 1$ for $i = 3, 4, 5$, and $f_i(0) = 0.5$ for $i = 6, 7, 8, 9$. We now consider the network dynamics under a disturbance process for which $\delta_5(1) = 0.55$, $\delta_e(t) = 0$ for all $t \geq 2$ and $\delta_i(t) \equiv 0$ for all $i \in \{1, \ldots, 10\} \setminus \{5\}$. Since $C_5(1) = 1.5 - 0.55 = 0.95 < f_1(1) = 1$, $e_5 \notin \mathcal{E}(2)$. This is followed by $3 \notin \mathcal{V}(3)$ and $e_3 \notin \mathcal{E}(4)$. Therefore,

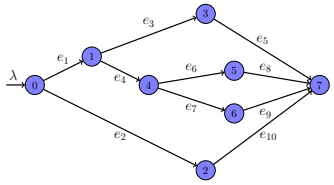

Fig. 1. A simple graph for the illustration of cascading failure under the proposed network dynamics.

$f_4(5) = 2$, $f_6(6) = f_7(6) = f_8(7) = f_9(7) = 1$. Since $C_8(7) = C_8(0) = 0.75 < f_8(7)$, $e_8 \notin \mathcal{E}(9)$. By continuing along these lines, the order of links to become inactivated is $e_5, e_3, e_8, e_6, e_9, e_7, e_4, e_1, e_{10}, e_2$. This clearly demonstrates that the links to fail successively under our proposed network dynamics are not necessarily adjacent to each other.

Remark 2: The model in (2)-(5) is to be contrasted with the dynamical flow network formulation in our previous work [14], [15] where every link has infinite buffer capacity, and hence there are no cascading effects under link overload. This feature is relaxed in our subsequent work [16], 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 cascade effect, which was proven to increase the resilience of the network with respect to the framework in [14], [15]. 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

$$\lim_{t \to +\infty} \sum_{d \in \mathcal{D}} \sum_{e \in \mathcal{E}_d^-} f_e(t) = \sum_v \lambda_v \,, \tag{8}$$

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., $\sum_{d \in \mathcal{D}} \sum_{e \in \mathcal{E}_d^-} f_e(t) = \lambda$ for $t = 0$. 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 in and outflow from the network also match asymptotically. For disturbance processes that are active only over finite time, (8) 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)): $\mathcal{D}(\delta) := \sum_{e\in\mathcal{E}} \triangle_e(\infty)$.

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:

$$\mathcal{R}^*(\mathcal{N},\lambda) = \sup_{\mathcal{G}} \mathcal{R}(\mathcal{N},\lambda,\mathcal{G}), \tag{9}$$

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 (9). 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 $\{0,1,\ldots,n\}$, 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 is 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 $\mathcal{G}$

$$\min_{e\in\mathcal{E}}\{C_e - f_e(0)\} \leq \mathcal{R}(\mathcal{N},\lambda,\mathcal{G}) \leq \min_{\mathcal{U}} C_{\mathcal{U}} - \lambda\,, \tag{10}$$

where the minimization in the upper bound is over all the cuts in $\mathcal{N}$. The lower bound in (10) 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 (10), 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 0.25, corresponding to link e_8, and the network residual capacity is 2.75, corresponding to the cut $\{3,5,6,2\}$. However the example also constructs a disturbance process of magnitude 0.55 under which the network is not transferring (under proportional routing policy).

We refer to [12] for a recursive procedure to compute a sharper upper bound under a centralized routing architecture that considers the multi-stage feature of cascading failures, but does not take into account the possibility of link inactivation due to the inactivation of the corresponding head node. In Section III-B, 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 1: 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 3: 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 1 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. The Backward Propagation Algorithm (BPA)

We now describe the Backward Propagation Algorithm (BPA) to compute a tighter upper bound on the margin of resilience. 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 $\mathcal{V} = \{0, \dots, n\}$. We shall assume to have fixed one such ordering in such a way that $\mathcal{E}_v^- \subseteq \bigcup_{0 \le u < v} \mathcal{E}_u^+$ 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: Backward Propagation Algorithm (BPA)

1: $S(\mathcal{E}_n^+, r, \mu) := +\infty$ for all $r \in \mathbb{R}_+^{\mathcal{E}_n^+}$ and $\mu \ge 0$ {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 $\mu \ge 0$, $S(\emptyset, r, \mu) = 0$,

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

$$S_e(\mu) = S(e, r, \mu) := \min\left\{C_e - \mu, S(\mathcal{E}_{\tau_e}^+, \mathbf{0}, \mu)\right\} \quad \forall\, e \in \mathcal{E}_v^+. \tag{11}$$

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

$$S(\mathcal{J}, r, \mu) := \max_{x \in \mathcal{X}_v(\mathcal{J}, r, \mu)} \min_{e \in \mathcal{J}} \Big(S_e(x_e) + S\big(\mathcal{J} \setminus \{e\}, x, \mu\big)\Big) \tag{12}$$

5: **end for**

Note that r appears only in the constraint set in the right hand side of (12). The Backward Propagation Algorithm derives its name from the central feature of the algorithm, where an intermediate node collects $S(\mathcal{J}, r, \mu)$ 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. We refer to [12] for illustration of BPA on simple networks.

Complementary to the maximization in (12) is the set of corresponding maximizers:

$$g\left(\mathcal{J}, r, \mu\right) := \operatorname*{argmax}_{x \in \mathcal{X}_v(\mathcal{J}, r, \mu)} \min_{e \in \mathcal{J}} \Big(S_e(x_e) + S\left(\mathcal{J} \setminus \{e\}, x, \mu\right)\Big). \tag{13}$$

C. 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 $(\delta(t))_{t \ge 1}$ with $\mathcal{D}(\delta) \le S^*(\mathcal{N}, \lambda)$ under which the associated network flow dynamics (2)-(5) is not transferring.

Remark 4: Theorem 1 implies that $\mathcal{R}(\mathcal{N}, \lambda, \mathcal{G}) \le S^*(\mathcal{N}, \lambda)$ for all distributed routing policies $\mathcal{G}$, and hence $\mathcal{R}^*(\mathcal{N}, \lambda) \le S^*(\mathcal{N}, \lambda)$.

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

We now 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 $v \in \mathcal{V} \setminus \{n\}$, $\mu \ge 0$:

$$\begin{aligned} r^* := G^v(\mathcal{E}_v^+, \mu) &\in g\left(\mathcal{E}_v^+, \mathbf{0}, \mu\right), \\ G^v\left(\mathcal{J}, \mu\right) &\in g\left(\mathcal{J}, r^*, \mu\right), \quad \mathcal{J} \subset \mathcal{E}_v^+. \end{aligned} \tag{14}$$

BPA routing derives its name from the fact that it relies on the function $g(\mathcal{J}, r, \mu)$ from (13), which is directly related to the central computation in the BPA. However, note that the lower bound r^* in (14) 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.

In general, BPA routing is not readily maximally resilient for general networks which are not *directed trees* [1]. 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 tree.

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 5: 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 children nodes are symmetric, and identical to each other.

Proposition 2: 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 $(\delta(t))_{t \ge 1}$ 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 $\lambda = 2$, $C_{e_1} = 2.5$, $C_{e_i} = 3$ for $i = 2, 3$, $C_{e_i} = 2$ for $i = 4, 7$, $C_{e_i} = 0.6$ for $i = 5, 6$, $C_{e_8} = 0.75$, $C_{e_9} = 1.5$ and $C_{e_{10}} = 0.17$. 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 2, 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 $\delta_5(1) = 0.2$, $\delta_{10}(1) = 0.07$, $\delta_i(1) = 0$ for $i \in \{1, \dots, 15\} \setminus \{2, 4, 5\}$ and $\delta(t) \equiv \mathbf{0}$ for all $t \ge 2$. The

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

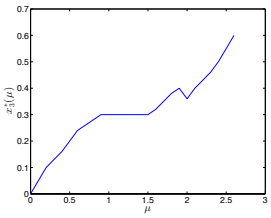

Fig. 2. Plot of $x_3^*(\mu) := G_{e_3}(\mathcal{E}_1^+(0), \mu)$ vs. μ.

magnitude 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 2 then implies that $1.9 - f_4(0) = f_3(0) = f_5(0) = 0.4$. Therefore, under the given disturbance process, $\{e_{10}, e_5\} \notin \mathcal{E}(2)$, and $\{e_2, e_3\} \notin \mathcal{E}(3)$. Hence $f_1(4) = 2$ and $f_4(5) = 2 = C_{e_4}$. This implies that $e_4 \notin \mathcal{E}(6)$, and hence $e_1 \notin \mathcal{E}(8)$, 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 2) 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 $\mathcal{J} \subseteq \mathcal{E}_v^+$:

$$0 \leq \mu_1 \leq \mu_2 \implies G^v(\mathcal{J}, \mu_1) \leq G^v(\mathcal{J}, \mu_2), \qquad (15)$$

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 6: Note that, unlike the link monotonicity condition in (7), we did not include the flow monotonicity condition in (15) 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 (14).

We refer to [12] for sufficient conditions on network parameters that guarantee flow monotonicity. 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 $\mathcal{D}(\delta) < S^*(\mathcal{N}, \lambda)$.

IV. 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 local inflow and active status of outgoing links, and in particular no information about the disturbance process. We quantified margin of resilience and presented an algorithm that provides an upper bound 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, etc.. We also plan to extend our formulation to the physics of infrastructure networks such as transportation, power, gas, water and supply chains.

References

[1] D. J. Watts, "A simple model of global cascades on random networks," *PNAS*, vol. 99, no. 9, pp. 5766–5771, 2002.

[2] A. Motter, "Cascade-based attacks on complex networks"," *Phys. Rev. E; Physical Review E*, vol. 66, no. 6, 2002.

[3] P. Crucitti, V. Latora, and M. Marchiori, "Model for cascading failures in complex networks," *Physical Review E*, vol. 69, no. 4, 2004.

[4] T. Liggett, *Interacting particle systems*. Springer-Verlag, 1985.

[5] G. Grimmett, *Percolation*. Springer, 1999.

[6] M. Draief and L. Massoulié, *Epidemics and rumors in complex networks*. Cambridge University Press, 2010.

[7] 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.

[8] 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.

[9] I. Dobson, B. A. Carreras, and D. E. Newman, "A loading-dependent model of probabilistic cascading failure," *Probability in the Engineering and Informational Sciences*, vol. 19, no. 01, pp. 15–32, 2005.

[10] 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.

[11] 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.

[12] K. Savla, G. Como, and M. A. Dahleh, "Robust network routing under cascading failures," *IEEE Transactions on Network Science and Engineering*, 2014, under review. Available at `http://arxiv.org/abs/1407.3518`.

[13] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[14] 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.

[15] ——, "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.

[16] 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, to appear. Available at `http://arxiv.org/abs/1308.1993`.