21st International Symposium on
Mathematical Theory of Networks and Systems
July 7-11, 2014. Groningen, The Netherlands

On Resilience of Oblivious Routing Policies for Networks under Cascaded Dynamics

Ketan Savla Giacomo Como Munther A. Dahleh

I. Overview

A popular class of generic models for cascading phenomena are the percolation models, e.g., see [1], [2]. The core of these models is a discrete-time dynamical description of the activation status of every node in the network, conditional on the activation status of the neighboring nodes. In the context of physical networks, these models are a gross approximation of the causal relationship between failure of successive nodes. A by-product of such simplistic models is that the successive nodes to fail are constrained to be adjacent to each other. Physics-based dynamical models for cascading failures in power networks have been proposed and analyzed in [3], [4], even allowing for control in between successive failure events [5]. In this paper, we propose dynamical models for cascading failures in network flows [6], and develop routing policies under such a framework that are provably resilient towards disturbances that reduce flow capacities of the links.

We consider network flow over directed graphs between a single origin-destination pair, where the network state consists of flows *and* activation status of the links. The evolution of the activation status of a link depends on its saturation status and the activation status of the downstream links. The flow dynamics is determined by activation status of the links and node-wise routing policies under flow balance constraints at the nodes. We formulate discrete time dynamics for the network state, where the time epochs correspond to a change in the activation status of the links, and study network resilience towards disturbances that reduce link-wise flow capacities, under distributed routing policies that are oblivious to disturbances. The margin of resilience of a routing policy is defined as the minimum, among all possible disturbances, of the link-wise sum of reductions in flow capacities, under which the links outgoing from the origin become inactive in finite time. We provide a tight characterization of the margin of resilience under such a setting in terms of network parameters for several scenarios. By bringing our attention to oblivious routing policies, we are able to give a tight characterization of margin of resilience for a wider class of networks than those in [7].

II. Formulation

A *flow network* $\mathcal{N} = (\mathcal{T}, C)$ is the pair of a *topology*, described by a finite directed multigraph $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{0, 1, \ldots, n\}$ is the node set and $\mathcal{E}$ is the link multiset, and a vector $C \in \mathbb{R}_+^{\mathcal{E}}$ describing the maximum flow capacities on the links. The unique origin and destination nodes are identified with $v = 0$ and $v = n$ respectively. We shall restrict ourselves to a class of topologies such that, for every node $v \in \mathcal{V} \setminus \{n\}$, there exists one and only one directed path from the origin node 0 to node v in $\mathcal{T}$. We shall refer to such $\mathcal{T}$ as a *polytree*. For every link $e \in \mathcal{E}$, we describe its state by $(\xi_e, f_e) \in \{0,1\} \times \mathcal{F}_e$. The binary variable ξ_e is used to denote the status of link e, i.e., link e is active if $\xi_e = 1$ and inactive otherwise. The variable f_e denotes the flow on link e. For $e = (v, w) \in \mathcal{E}$, where $v, w \in \mathcal{V}$, $\sigma_e = v$ and $\tau_e = w$ represent the tail and head node, respectively, of e. Let $\mathcal{E}_v^+ := \{e \in \mathcal{E} : \sigma_e = v\}$ and $\mathcal{E}_v^- := \{e \in \mathcal{E} : \tau_e = v\}$, respectively, denote the set of outgoing and incoming links at node $v \in \mathcal{V}$. We use the vector notations $f = \{f_e : e \in \mathcal{E}\}$ and $f^v = \{f_e : e \in \mathcal{E}_v^+\}$. ξ and ξ^v are defined in a similar fashion. For a given inflow $\lambda > 0$ at the origin, a flow vector $f = \{f_e : e \in \mathcal{E}\} \in \Pi_{e \in \mathcal{E}} [0, C_e]$ is called admissible if $\sum_{e \in \mathcal{E}_0^+} f_e = \lambda$ and $\sum_{e \in \mathcal{E}_v^+} f_e = \sum_{e \in \mathcal{E}_v^-} f_e$ for all $v \in \{1, \ldots, n-1\}$. The network suffers a disturbance $\delta \in \mathbb{R}_+^{\mathcal{E}}$ at time $t = 0$. The effect of the disturbance is to reduce the link-wise maximum flow capacities. Formally, under the application of disturbance δ, the maximum flow capacity on link e decreases to $C_e - \delta_e$. The response of the network to disturbance is through routing policies, one for every node, that determine the splitting of the inflow at a node among the links outgoing from that node. In this paper, we focus on *oblivious* routing policies that have no access to any information about the disturbance. The formal description is as follows, where $\mathcal{R}_v := \mathbb{R}_+^{\mathcal{E}_v^+}$ is the set of nonnegative-real-valued vectors whose entries are indexed by elements of $\mathcal{E}_v^+$; and for a given $\mu \geq 0$, $\mathcal{S}_v(\mu) := \{x \in \mathcal{R}_v : \sum_{e \in \mathcal{E}_v^+} x_e = \mu\}$.

Definition 1: (Oblivious routing policy) An *oblivious routing policy* for a network $\mathcal{N}$ is a family of functions $\mathcal{G} := \{G^v : \{0,1\}^{\mathcal{E}_v^+} \times \mathbb{R}_+ \to \mathcal{R}_v; (\xi^v, \lambda_v) \mapsto G^v(\xi^v, \lambda_v) \in \mathcal{S}_v(\lambda_v) \cup \mathbf{0}\}_{0 \leq v < n}$ describing the splitting of the incoming flow from each non-destination node v among its outgoing link set $\mathcal{E}_v^+$, as a function of the observed status on the outgoing links, and the incoming flow. We also implicitly assume that every G^v, $v \in \{0, \ldots, n-1\}$, has a priori information about the vector of link flow capacities C. Moreover, we will adopt the conventions that, for all $v \in \{0, \ldots, n-1\}$, and $\lambda_v \geq 0$: for all $e \in \mathcal{E}_v^+$, if $\xi_e = 0$ then $G_e^v(\xi^v, \lambda_v) = 0$, and $G^v(\xi^v, \lambda_v) = \mathbf{0}$ if and only if $\xi^v = \mathbf{0}$.

A simple example of a monotone oblivious routing policy is the *proportional* policy: for every $v \in \mathcal{V}$, $\xi^v \in \{0,1\}^{\mathcal{E}_v^+}$, $\mu > 0$, $G_e^v(\xi^v, \mu) = C_e \mu / (\sum_{e \in \mathcal{H}_v^0} C_e)$ if $\xi_e = 1$ and zero otherwise.

We consider the following discrete-time dynamics in this paper. For all $e \in \mathcal{E}$, $t \geq 0$,

$$\begin{array}{rcl} f_e(t+1) & = & G_e^{\sigma_e}\big(\xi^{\sigma_e}(t), \lambda_{\sigma_e}(t)\big), \\ \xi_e(t+1) & = & \chi_e(t+1) \cdot \psi_{\tau_e}(t) \cdot \xi_e(t), \end{array} \tag{1}$$

where $\chi_e(t) := \mathbf{1}_{f_e(t) < C_e - \delta_e}$ for $t \geq 1$ is a binary variable that describes the saturation status of link e, $\psi_v(t) := 1 - \Pi_{e \in \mathcal{E}_v^+}(1 - \xi_e(t))$, $t \geq 0$ is a binary variable that describes the active status of node v, and $\lambda_0(t) \equiv \lambda$, and $\lambda_v(t) := \sum_{e \in \mathcal{E}_v^-} f_e(t)$ for all $v > 0$ is the incoming flow at node $v \in \mathcal{V}$. The initial conditions are $\xi(0) = \mathbf{1}$ and $f_e(0) = G_e(\mathbf{1}, \lambda_{\sigma_e}(0))$. We emphasize that, under (1), a link e becomes irreversibly inactive, i.e., $\xi_e(t) = 0$, either as soon as the flow on it is at least equal to its flow capacity under disturbance, or if the node τ_e becomes inactive at the previous time instant. We shall say that a routing policy $\mathcal{G}$ is *feasible*, if the initial flow $f_e(0) = G_e(\mathbf{1}, \lambda_{\sigma_e}(0))$, $e \in \mathcal{E}$, induced by it is admissible. In this paper, we implicitly assume that the routing policies under consideration are feasible.

We now present an example to illustrate the phenomenon of cascading failure as modeled by (1), and distinguish it from other existing models for cascading failures [1], [2].

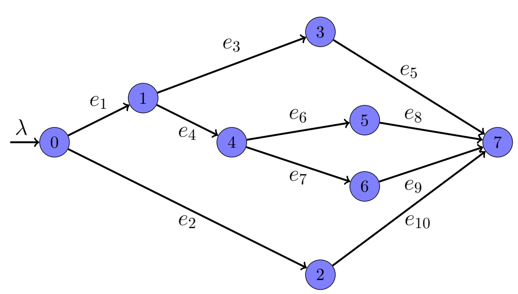

Fig. 1. A simple graph for the illustration of cascading failure.

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 $i = 3, 4, 10$, $C_i = 2$ for $i = 6, 7$, $C_i = 1.5$ for $i = 5, 9$ and $C_8 = 0.75$. Let the arrival rate 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 cascade dynamics under a disturbance δ such that $\delta_5 = 0.75$ and $\delta_i = 0$ for all $i \in \{1, \ldots, 10\} \setminus \{5\}$. For these values, we have that $\chi_{e_5}(0) = 0$. This implies that $\xi_{e_5}(0) = 0$, and hence $\psi_3(0) = 0$. Therefore, $\xi_{e_3}(1) = 0$, which implies that $f_{e_4}(2) = 2$, $f_{e_6}(3) = f_{e_7}(3) = 1$ and $f_{e_8}(4) = f_{e_9}(4) = 1$. Since $C_8 = 0.75 < 1 = f_{e_8}(4)$, this implies that $\chi_{e_8}(4) = 0 = \xi_{e_8}(4) = 0$. By continuing the arguments along similar lines, one can conclude that the sequence of links that fail one after another is $e_5, e_3, e_8, e_6, e_9, e_7, e_4, e_1, e_{10}, e_2$. In particular, the links to fail successively are not adjacent to each other, as would be implied by the cascading failure models, e.g., in [1], [2].

The effect of cascading failures is that the network flow vector $f(t)$ may not remain feasible for all $t \geq 0$. We formalize the corresponding dichotomy in the network behavior by the notion of transfer efficiency as follows. The network is called *transferring* under disturbance δ if there exists $T \geq 0$ such that $\lambda_n(t) = \lambda$ for all $t \geq T$. The following result shows that the network being transferring is equivalent to the origin node being active all the time.

Proposition 2: Consider a dynamical flow network with a distributed routing policy $\mathcal{G}$, and $\lambda > 0$ a constant total outflow at the origin node. The network is transferring under disturbance δ if and only if $\psi_0(t) = 1$ for all $t \geq 0$.

We define the *magnitude of a disturbance* $\delta \in \mathcal{F}$ to be its 1-norm, i.e., $\|\delta\|_1 = \sum_{e \in \mathcal{E}} \delta_e$. The *margin of resilience* of the network, denoted as $\mathcal{R}(\lambda, \mathcal{N}, \mathcal{G})$, is defined as the infimum of the magnitude of disturbances under which the network is not transferring. In this paper, we consider the following optimization problem: $\mathcal{R}^*(\lambda, \mathcal{N}) = \sup_{\mathcal{G}} \mathcal{R}(\lambda, \mathcal{N}, \mathcal{G})$, where the supremum is taken over the class of oblivious monotone distributed routing policies. The objective is to find a policy $\tilde{\mathcal{G}}$ such that $\mathcal{R}(\lambda, \mathcal{N}, \tilde{\mathcal{G}})$ matches $\mathcal{R}^*(\lambda, \mathcal{N})$ as closely as possible. In general, finding maximally resilient oblivious routing policies is difficult. Therefore, we shall restrict ourselves to a reasonable class of policies, known as *monotone* oblivious routing policies, which are the policies that satisfy the following two properties, where for $i = 1, 2$, $\mathcal{H}_{v,i} := \{e \in \mathcal{E}_v^+ \mid \xi_e^{v,i} = 1\}$, and $G_{\mathcal{H}_{v,i}}^v$ is the projection of G on $\mathcal{H}_{v,i}$:

a) for every $\xi^{v,1} \in \{0,1\}^{\mathcal{E}_v^+}$, $0 \leq \mu_1 < \mu_2 \implies G_{\mathcal{H}_{v,1}}^v(\xi^{v,1}, \mu_1) \preceq G_{\mathcal{H}_{v,1}}^v(\xi^{v,1}, \mu_2)$;
b) for every $\mu \geq 0$, $\xi^{v,1} \preceq \xi^{v,2} \implies G_{\mathcal{H}_{v,1}}^v(\xi^{v,2}, \lambda_v) \preceq G_{\mathcal{H}_{v,1}}^v(\xi^{v,1}, \lambda_v)$.

Property (a) implies that if the inflow at a node increases, then the flow assigned to every active outgoing link from that node does not decrease. Property (b) implies that for the same arrival rate, shrinking of the set of active links results in increase in outflow on each of the remaining active outgoing links.

III. Results

We start with an algorithm to compute a quantity, which shall be related to $\mathcal{R}^*(\lambda, \mathcal{N})$. In Algorithm 1,

$$\mathcal{X}_{\mathcal{J}}(\mu) := \Big\{ x \in \prod_{i \in \mathcal{E}_v^+} [0, C_i] \Big| \; x_i = 0 \quad \forall i \in \mathcal{E}_v^+ \setminus \mathcal{J}; \sum_{i \in \mathcal{E}_v^+} x_i = \mu \Big\}.$$

Algorithm 1: Backward Propagation Algorithm

$R_n(\mu) := +\infty$ for all $\mu \geq 0$ {destination node}
for $v = n-1, n-2, \ldots, 0$ **do** {construct a series of intermediate functions for every node starting with $n-1$, and going backward up to the origin}
  put $S_{\emptyset}(\mu) \equiv 0$ and

$$S_e(\mu) := \min \left\{ [C_e - \mu]^+, R_{\tau_e}(\mu) \right\} \qquad \forall \mu \geq 0,\ e \in \mathcal{E}_v^+.$$

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

$$S_{\mathcal{J}}(\mu) = 0 \quad \text{if } \mu \geq \sum_{j \in \mathcal{J}} C_j \,,$$

  and, if $\mu < \sum_{j \in \mathcal{J}} C_j$, then

$$S_{\mathcal{J}}(\mu) := \max_{x \in \mathcal{X}_{\mathcal{J}}(\mu)} \min_{e \in \mathcal{J}} \left\{ S_e(x_e) + S_{\mathcal{J} \setminus \{e\}}(\mu) \right\} \tag{2}$$

  $R_v(\mu) := S_{\mathcal{E}_v^+}(\mu)$.
end for

Theorem 3 (Upper Bound): Consider a flow network $\mathcal{N} = (\mathcal{T}, C)$ where $\mathcal{T}$ is a polytree, and $\lambda > 0$ is a constant total outflow at the origin node. Then, for any monotone oblivious routing policy, there exists a disturbance $\delta \in \mathbb{R}_+^{\mathcal{E}}$ with $\|\delta\|_1 \leq R_0(\lambda)$ under which the network is not transferring.

Theorem 3 implies that $R_0(\lambda)$ is a uniform upper bound on the margin of resilience over the class of oblivious monotone distributed routing policies.

We next we develop lower bounds for $R^*(\lambda, \mathcal{N})$. This will be done by analyzing a specific oblivious routing policy, called *BPA-routing*, denoted as $\mathcal{G}^*$, whose construction is inspired by the Backward Propagation Algorithm. We first extend (2) as follows. For $\emptyset \subseteq \mathcal{L} \subseteq \mathcal{J}$, $\mathcal{J} \subseteq \mathcal{E}_v^+$, $v \in \{1, \ldots, n\}$, $\mu \geq \nu \geq 0$, let

$$\tilde{S}\left(\nu, \mu, \mathcal{L}, \mathcal{J}\right) := \max_{x \in \mathcal{X}_{\mathcal{L}}(\nu)} \min_{e \in \mathcal{L}} \left\{ S_e(x_e) + S_{\mathcal{J} \setminus \{e\}}(\mu) \right\}. \tag{3}$$

(2) and (3) imply that $\tilde{S}\left(\mu, \mu, \mathcal{J}, \mathcal{J}\right) = S_{\mathcal{J}}(\mu)$. The formal description of BPA routing is in Algorithm 2, where $\mathcal{H}_v^0(\xi^v) := \{e \in \mathcal{E}_v^+ \mid \xi_e^v = 1\}$.

Algorithm 2: BPA Routing

Require: $\mu \geq 0$, $\xi^v \in \{0,1\}^{\mathcal{E}_v^+}$; compute $\mathcal{H}_v^0(\xi^v, \mu)$.
let $\mathbf{0} = x^* \in \mathbb{R}_+^{\mathcal{E}_v^+}$ {initialize flow on all links as zero}
initialize $\nu = \mu$ and $\mathcal{L} = \mathcal{H}_v^0(\xi^v, \mu)$
while $\mathcal{L} \neq \emptyset$ **do** {compute link flows by iterative solution of nested optimization problems}
  let $\mathcal{M}$ be the set of links e in $\mathcal{L}$, for which there exists $y_e \in [0, \Gamma_e]$ that satisfies the following two conditions:

$$S_e(y_e) = \tilde{S}(\nu, \mu, \mathcal{L}, \mathcal{H}_v^0) - S_{\mathcal{H}_v^0 \setminus \{e\}}(\mu), \tag{4}$$

$$S_e(y_e) \leq \tilde{S}(\nu - y_e, \mu, \mathcal{L} \setminus \{e\}, \mathcal{H}_v^0) - S_{\mathcal{H}_v^0 \setminus \{e\}}(\mu). \tag{5}$$

  for all $e \in \mathcal{M}$, set x_e^* to be equal to $y_e \in [0, \Gamma_e]$ that satisfies (4)-(5).
  $\nu \leftarrow \nu - \sum_{e \in \mathcal{M}} y_e$ {remaining flow to be assigned to links in $\mathcal{L} \setminus \mathcal{M}$}
  $\mathcal{L} \leftarrow \mathcal{L} \setminus \mathcal{M}$
end while
set $G_e^*(\xi^v, \mu) = x_e^*$ for all $e \in \mathcal{E}_v^+$

The specification of $\mathcal{G}^*$ involves solution of nested optimization problems, as given by (3) and (4)-(5). This is to be contrasted with (2) in the specification of the Backward Propagation Algorithm, which involves only one optimization problem. This is because, in general, (2) does not have a unique maximizer. On the other hand, the nested optimization structure of BPA routing always ensures uniqueness.

The following result, where $\Gamma_e := \sup\{\mu \geq 0 \mid S_e(\mu) > 0\}$, states useful properties of BPA-routing.

Lemma 4: $\mathcal{G}^*(\xi^v, \mu)$ satisfies the following at every $v \in \mathcal{V} \setminus \{n\}$:

(i) For every $\xi^v \in \{0,1\}^{\mathcal{E}_v^+}$ and $\mu \in \left[0, \sum_{e \in \mathcal{H}_v^0} \Gamma_e\right]$, $G^{v,*}(\xi^v, \mu)$ is a singleton, and it is a maximizer in (2), i.e.,

$$G^{v,*}(\xi^v, \mu) \in \underset{x \in \mathcal{X}_{\mathcal{H}_v^0}(\mu)}{\operatorname{argmax}} \min_{e \in \mathcal{H}_v^0} \left\{ S_e(x_e) + S_{\mathcal{H}_v^0 \setminus \{e\}}(\mu) \right\};$$

(ii) For every $\xi^v \in \{0,1\}^{\mathcal{E}_v^+}$, $G^*_{\mathcal{E}_v^+}(\xi^v, \mu)$ is continuous in μ over $\left(0, \sum_{e\in\mathcal{H}_v^0} \Gamma_e\right)$.

In light of Lemma 4 (i), the action of the BPA-routing is unique for all combinations of active links and the corresponding feasible inflows, and hence the dynamics (1) of network flow under routing policy $\mathcal{G}^*$ is well-defined.

Theorem 5 (Lower Bound): Consider a flow network $\mathcal{N} = (\mathcal{T}, C)$ where $\mathcal{T}$ is a polytree, $\lambda > 0$ a constant total outflow at the origin node, and operating under BPA routing policy $\mathcal{G}^*$. If $\mathcal{G}^*$ is monotone, then the network is transferring for every $\delta \in \mathcal{F}(\lambda)$ such that $0 \le \|\delta\|_1 < R_0(\lambda)$.

In general, one could perform extensive (offline) numerical analysis to test the monotonicity properties of the BPA routing $\mathcal{G}^*$ over a given network $\mathcal{N}$. We now present analytical sufficient conditions on $\mathcal{N}$, under which $\mathcal{G}^*$ is guaranteed to be monotone. In order to state the results, we need to introduce a few more concepts. We start with the definition of a *symmetric tree*, where the *weights* on the links correspond to flow capacities.

A weighted directed tree of depth[1] one is called symmetric if all the links outgoing from the root node have equal weights. A weighted directed tree of depth greater than one is called symmetric if all the sub-trees rooted at the children[2] nodes are symmetric, and identical to each other.

For a directed tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, a directed chain $(u = v_0, v_1, \dots, v_k = w)$ from $u \in \mathcal{V}$ to $w \in \mathcal{V}$ is called non-branching if $|\mathcal{E}_u^+| = 1$, $|\mathcal{E}_{v_i}^+| = |\mathcal{E}_{v_i}^-| = 1$ for all $i = 1, \dots, k-1$ and $|\mathcal{E}_w^-| = 1$.

For a directed weighted tree $\mathcal{N} = (\mathcal{V}, \mathcal{E}, C)$, its meta-tree is a weighted tree $\mathcal{N}^m = (\mathcal{V}^m, \mathcal{E}^m, C^m)$ defined as: $\mathcal{V}^m := \mathcal{V} \setminus \{v \in \mathcal{V} \mid |\mathcal{E}_v^-| = |\mathcal{E}_v^+| = 1\}$, $\mathcal{E}^m := \{(u,w) \in \mathcal{V}^m \times \mathcal{V}^m \mid \exists$ directed non-branching chain from u to w in $\mathcal{N}\}$, and $C^m_{(u,w)} = \min_{i=1}^k C_{(v_{i-1}, v_i)}$, where $(u = v_0, v_1, \dots, v_k = w)$ is the directed non-branching chain from u to w in $\mathcal{N}$. Note that, if there exists a directed non-branching chain from u to w in $\mathcal{N}$, then it is unique.

For a directed weighted tree $\mathcal{N} = (\mathcal{V}, \mathcal{E}, C)$, a node $v \in \mathcal{V}$ is called *pseudo-leaf* if $|\mathcal{E}_v^-| = 1$, and with $\mathcal{E}_v^- = \{e\}$, we have that $[C_e - \mu]^+ \le R_v(\mu)$ for all $\mu \in [0, C_e]$. For a node v in $\mathcal{N}$, its cascade-relevant sub-tree is the maximal sub-tree of $\mathcal{N}$ which is rooted at v, whose leaf nodes are pseudo-leaves in $\mathcal{N}$, and whose non-leaf nodes are not pseudo-leaves in $\mathcal{N}$.

Theorem 6: The BPA routing $\mathcal{G}^*$ is monotone if :

[1]The depth of a directed tree is the length of the longest directed path in the tree.

[2]In a directed tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, the children nodes of $v \in \mathcal{V}$ are the set of nodes $u \in \mathcal{V}$ such that $\mathcal{E}_u^- \cap \mathcal{E}_v^+ \neq \emptyset$.

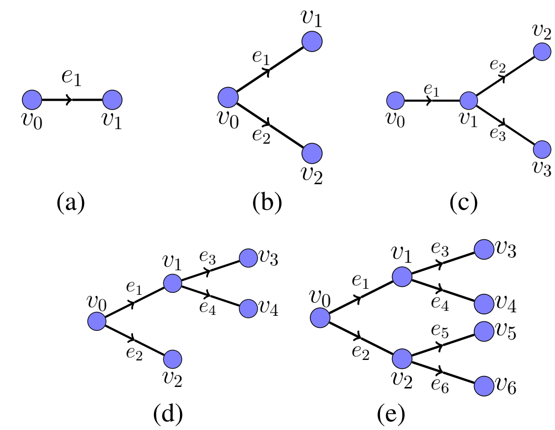

Fig. 2. Cascade-relevant subtree topologies that induce monotonicity of BPA routing at the root node.

(i) $\mathcal{N}$ is symmetric; or

(ii) the cascade-relevant sub-tree of every node in meta-tree $\mathcal{N}^m$ is either (a), or (b), or (c), or (d) with $C_{e_2} \ge C_{e_1}$ or (e) with $C_{e_1} = C_{e_2}$ in Figure 2.

Conclusion

We proposed a model for cascading failures in network flows under disturbances that reduce link-wise flow capacities. We proposed an algorithm to compute an upper bound on the margin of resilience for such networks under monotone oblivious routing policies, and we identified several scenarios under which this bound is tight. In future, we plan to extend our formulation to relax the polytree assumption, consider multi-commodity flows, stochastic disturbance generation, finite time link recovery, and exogenous coupling between failure and recovery of distant links due to interdependence of the given network with other networks.

References

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

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

[3] 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.

[4] C. Lai and S. H. Low, "The redistribution of power flow in cascading failures," in *51st Annual Allerton Conference on Communication, Control, and Computing*, pp. 1037–1044, 2013.

[5] D. Bienstock, "Optimal control of cascading power grid failures," in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pp. 2166–2173, IEEE, 2011.

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

[7] K. Savla, G. Como, M. A. Dahleh, and E. Frazzoli, "On resilience of distributed routing in networks under cascade dynamics.," in *IEEE Conference on Decision and Control*, (Florence, Italy), 2013.