



LUND UNIVERSITY

On distributed robust routing for transportation networks under local information constraints

Como, Giacomo; Savla, Ketan; Acemoglu, Daron; Dahleh, Munther A.; Frazzoli, Emilio

2011

[Link to publication](#)

Citation for published version (APA):

Como, G., Savla, K., Acemoglu, D., Dahleh, M. A., & Frazzoli, E. (2011). *On distributed robust routing for transportation networks under local information constraints*. 6290-6295. Paper presented at 50th IEEE Conference on Decision and Control and European Control Conference, 2011, Orlando, Florida, United States.

Total number of authors:

5

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Robust Distributed Routing in Dynamical Flow Networks

Giacomo Como Ketan Savla Daron Acemoglu Munther A. Dahleh Emilio Frazzoli

Abstract—Robustness of distributed routing policies is studied for dynamical flow networks, with respect to adversarial disturbances that reduce the link flow capacities. A dynamical flow network is modeled as a system of ordinary differential equations derived from mass conservation laws on a directed acyclic graph with a single origin-destination pair and a constant inflow at the origin. Routing policies regulate the way the inflow at a non-destination node gets split among its outgoing links as a function of the current particle density, while the outflow of a link is modeled to depend on the current particle density on that link through a flow function. The robustness of distributed routing policies is evaluated in terms of the network’s weak resilience, which is defined as the infimum sum of link-wise magnitude of disturbances under which the total inflow at the destination node of the perturbed dynamical flow network is positive. The weak resilience of a dynamical flow network with arbitrary routing policy is shown to be upper-bounded by the network’s min-cut capacity, independently of the initial flow conditions. Moreover, a class of distributed routing policies that rely exclusively on local information on the particle densities, and are locally responsive to that, is shown to yield such maximal weak resilience. These results imply that locality constraints on the information available to the routing policies do not cause loss of weak resilience.

I. INTRODUCTION

Flow networks provide a fruitful modeling framework for many applications of interest such as transportation, data, and production networks. They entail a fluid-like description of the macroscopic motion of *particles*, which are routed from their origins to their destinations via intermediate nodes: we refer to standard textbooks, such as [1], for a thorough treatment.

The present paper studies *dynamical flow networks*, modeled as systems of ordinary differential equations derived from mass conservation laws on directed acyclic graphs with a single origin-destination pair and a constant inflow at the origin. The rate of change of the particle density on each link of the network equals the difference between the *inflow* and the *outflow* of that link. The latter is modeled to depend on the current particle density on that link through a *flow function*. On the other hand, the way the inflow at an intermediate node gets split among its outgoing links depends on the current particle density, possibly on the whole network, through a *routing policy*. Such a routing policy is said to be *distributed* if the proportion of inflow routed to the outgoing links of a node is allowed to depend only on *local information*, consisting of the current particle densities on the

outgoing links of the same node. The inspiration for such a modeling paradigm comes from empirical findings from several application domains such as transportation networks, data networks and production networks.

Our objective is the design and analysis of distributed routing policies for dynamical flow networks that are *maximally robust* with respect to *adversarial disturbances* that reduce the link flow capacities. The robustness of distributed routing policies is evaluated in terms of the network’s *weak resilience*, which is defined as the infimum sum of link-wise magnitude of disturbances under which the total inflow at the destination node of the perturbed dynamical flow network is positive. In this paper, we prove that the maximum possible weak resilience is yielded by a class of *locally responsive* distributed routing policies, which rely only on *local information* on the current particle densities on the network, and are characterized by the property that the portion of its inflow that a node routes towards an outgoing link does not decrease as the particle density on any other outgoing link increases. These results are mainly a consequence of the particular *cooperative* structure (in the sense of [2], [3]) that the dynamical flow network inherits from locally responsive routing policies. Moreover, we show that the maximum weak resilience of dynamical flow networks with arbitrary, not necessarily distributed, routing policies equals the *min-cut capacity* of the network and hence is independent of the initial equilibrium flow.

Stability analysis of network flow control policies under non-persistent disturbances, especially in the context of internet, has attracted a lot of attention, e.g., see [4], [5], [6], [7]. Recent work on robustness analysis of static flow networks under adversarial and probabilistic persistent disturbances in the spirit of this paper include [8], [9]. It is worth comparing the distributed routing policies studied in this paper with the back-pressure policy [10], which is one of the most well-known robust distributed routing policy for queueing networks. While relying on local information in the same way as the distributed routing policies studied here, back-pressure policies require the nodes to have, possibly unlimited, buffer capacity. In contrast, in our framework, the nodes have no buffer capacity. In fact, the distributed routing policies considered in this paper are closely related to the well-known *hot-potato* or deflection routing policies [11, Sect. 5.1], where the nodes route incoming packets immediately to one of the outgoing links. However, to the best of our knowledge, the robustness properties of dynamical flow networks, where the outflow from a link is not necessarily equal to its inflow have not been studied before. Due to space limitations, we keep our presentation concise here; detailed exposition can be found in [12].

Before proceeding, we define some preliminary notation

G. Como is with the Department of Automatic Control, Lund University, Lund, Sweden. giacomo@control.lth.se

K. Savla, M. A. Dahleh and E. Frazzoli are with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology, Cambridge, MA, USA.

D. Acemoglu is with the Department of Economics at the Massachusetts Institute of Technology, Cambridge, MA, USA.

{[ksavla](mailto:ksavla@mit.edu), [daron](mailto:daron@mit.edu), [dahleh](mailto:dahleh@mit.edu), [frazzoli](mailto:frazzoli@mit.edu)}@mit.edu

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. Let \mathcal{A} and \mathcal{B} be finite sets. Then, $|\mathcal{A}|$ will denote the cardinality of \mathcal{A} , $\mathbb{R}^{\mathcal{A}}$ (respectively, $\mathbb{R}_+^{\mathcal{A}}$) the space of real-valued (nonnegative-real-valued) vectors whose components are indexed by elements of \mathcal{A} , and $\mathbb{R}^{\mathcal{A} \times \mathcal{B}}$ the space of matrices whose real entries indexed by pairs of elements in $\mathcal{A} \times \mathcal{B}$. The transpose of a matrix $M \in \mathbb{R}^{\mathcal{A} \times \mathcal{B}}$, will be denoted by $M^T \in \mathbb{R}^{\mathcal{B} \times \mathcal{A}}$, while $\mathbf{1}$ the all-one vector, whose size will be clear from the context. Let $\text{cl}(\mathcal{X})$ be the closure of a set $\mathcal{X} \subseteq \mathbb{R}^{\mathcal{A}}$. 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). Given a multigraph $(\mathcal{V}, \mathcal{E})$, for every node $v \in \mathcal{V}$, we shall denote by $\mathcal{E}_v^+ \subseteq \mathcal{E}$, and $\mathcal{E}_v^- \subseteq \mathcal{E}$, the set of its outgoing and incoming links, respectively. Moreover, we shall use the shorthand notation $\mathcal{R}_v := \mathbb{R}_+^{\mathcal{E}_v^+}$ for the set of nonnegative-real-valued vectors whose entries are indexed by elements of \mathcal{E}_v^+ , $\mathcal{S}_v := \{p \in \mathcal{R}_v : \sum_{e \in \mathcal{E}_v^+} p_e = 1\}$ for the simplex of probability vectors over \mathcal{E}_v^+ , and $\mathcal{R} := \mathbb{R}_+^{\mathcal{E}}$ for the set of nonnegative-real-valued vectors whose entries are indexed by the links in \mathcal{E} .

II. DYNAMICAL FLOW NETWORKS AND THEIR RESILIENCE

In this section, we introduce our model of dynamical flow networks and define the notions of transfer efficiency.

A. Dynamical flow networks

We start with the following definition of a flow network.

Definition 1: A flow network $\mathcal{N} = (\mathcal{T}, \mu)$ is the pair of a topology, described by a finite directed multigraph $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the node set and \mathcal{E} is the link multiset, and a family of flow functions $\mu := \{\mu_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+\}_{e \in \mathcal{E}}$ describing the functional dependence $f_e = \mu_e(\rho_e)$ of the flow on the density of particles on every link $e \in \mathcal{E}$. The flow capacity of a link $e \in \mathcal{E}$ is defined as

$$f_e^{\max} := \sup_{\rho_e \geq 0} \mu_e(\rho_e). \quad (1)$$

We shall use the notation $\mathcal{F}_v := \times_{e \in \mathcal{E}_v^+} [0, f_e^{\max})$ for the set of admissible flow vectors on outgoing links from node v , and $\mathcal{F} := \times_{e \in \mathcal{E}} [0, f_e^{\max})$ for the set of admissible flow vectors for the network. We shall write $f := \{f_e : e \in \mathcal{E}\} \in \mathcal{F}$, and $\rho := \{\rho_e : e \in \mathcal{E}\} \in \mathcal{R}$, for the vectors of flows and of densities, respectively, on the different links. The notation $f^v := \{f_e : e \in \mathcal{E}_v^+\} \in \mathcal{F}_v$, and $\rho^v := \{\rho_e : e \in \mathcal{E}_v^+\} \in \mathcal{R}_v$ will stand for the vectors of flows and densities, respectively, on the outgoing links of a node v . We shall compactly denote by $f = \mu(\rho)$ and $f^v = \mu^v(\rho^v)$ the functional relationships between density and flow vectors.

Throughout this paper, we shall restrict ourselves to network topologies satisfying the following:

Assumption 1: The topology \mathcal{T} contains no cycles, has a unique origin (i.e., a node $v \in \mathcal{V}$ such that \mathcal{E}_v^- is empty), and a unique destination (i.e., a node $v \in \mathcal{V}$ such that \mathcal{E}_v^+ is empty). Moreover, there exists a path in \mathcal{T} to the destination node from every other node in \mathcal{V} .

Assumption 1 implies that one can find a (not necessarily unique) topological ordering of the node set \mathcal{V} (see, e.g.,




Fig. 1. An origin/destination cut of the network: \mathcal{U} is a subset of nodes including the origin 0 but not the destination n , and $\mathcal{E}_{\mathcal{U}}^+$ is the subset of those edges with tail node in \mathcal{U} , and head node in $\mathcal{V} \setminus \mathcal{U}$.

[13]). We shall assume to have fixed one such ordering, identifying \mathcal{V} with the integer set $\{0, 1, \dots, n\}$, where $n := |\mathcal{V}| - 1$, in such a way that

$$\mathcal{E}_v^- \subseteq \bigcup_{0 \leq u < v} \mathcal{E}_u^+, \quad \forall v = 0, \dots, n. \quad (2)$$

In particular, (2) implies that 0 is the origin node, and n the destination node in the network topology \mathcal{T} . An *origin-destination cut* (see, e.g., [1]) of \mathcal{T} is a partition of \mathcal{V} into \mathcal{U} and $\mathcal{V} \setminus \mathcal{U}$ such that $0 \in \mathcal{U}$ and $n \in \mathcal{V} \setminus \mathcal{U}$. Let

$$\mathcal{E}_{\mathcal{U}}^+ := \{(u, v) \in \mathcal{E} : u \in \mathcal{U}, v \in \mathcal{V} \setminus \mathcal{U}\} \quad (3)$$

be the set of all the links pointing from some node in \mathcal{U} to some node in $\mathcal{V} \setminus \mathcal{U}$ (see Fig. 1). The *min-cut capacity* of a flow network \mathcal{N} is defined as

$$C(\mathcal{N}) := \min_{\mathcal{U}} \sum_{e \in \mathcal{E}_{\mathcal{U}}^+} f_e^{\max}, \quad (4)$$

where the minimization runs over all the origin-destination cuts of \mathcal{T} . Throughout this paper, we shall assume a constant inflow $\lambda_0 \geq 0$ at the origin node. Let us define the set of *admissible equilibrium flows* associated to an inflow λ_0 as

$$\mathcal{F}^*(\lambda_0) := \{f^* \in \mathcal{F} : \sum_{e \in \mathcal{E}_0^+} f_e^* = \lambda_0, \sum_{e \in \mathcal{E}_v^+} f_e^* = \sum_{e \in \mathcal{E}_v^-} f_e^*, \forall 0 < v < n\}.$$

Then, it follows from the max-flow min-cut theorem (see, e.g., [1]), that $\mathcal{F}^*(\lambda_0) \neq \emptyset$ whenever $\lambda_0 < C(\mathcal{N})$. That is, the min-cut capacity equals the maximum flow that can pass from the origin to the destination node while satisfying capacity constraints on the links, and conservation of mass at the intermediate nodes.

Throughout the paper, we shall make the following assumption on the flow functions:

Assumption 2: For every link $e \in \mathcal{E}$, the map $\mu_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuously differentiable, strictly increasing, has bounded derivative, and is such that $\mu_e(0) = 0$, and $f_e^{\max} < +\infty$.

Thanks to Assumption 2, one can define the *median density* on link $e \in \mathcal{E}$ as the unique value $\rho_e^\mu \in \mathbb{R}_+$ such that

$$\mu_e(\rho_e^\mu) = f_e^{\max}/2. \quad (5)$$

We now introduce the notion of a distributed routing policy used in this paper.

Definition 2: A *routing policy* for a flow network \mathcal{N} is a family of differentiable functions $\mathcal{G} := \{G^v : \mathcal{R} \rightarrow \mathcal{S}_v\}_{0 \leq v < n}$ describing the ratio in which the particle flow incoming in each non-destination node v gets split among its outgoing link set \mathcal{E}_v^+ , as a function of the observed current particle density. A routing policy is said to be *distributed* if, for all $0 \leq v < n$, there exists a differentiable function $\bar{G} : \mathcal{R}_v \rightarrow \mathcal{S}_v$ such that $G^v(\rho) = \bar{G}^v(\rho^v)$ for all $\rho \in \mathcal{R}$, where ρ^v is the projection of ρ on the outgoing link set \mathcal{E}_v^+ .

The salient feature in Definition 2 is that a distributed routing policy depends only on the *local information* on the particle density ρ^v on the set \mathcal{E}_v^+ of outgoing links of the non-destination node v , instead of the full vector of current particle densities ρ on the whole link set \mathcal{E} . Throughout this paper, we shall make a slight abuse of notation and write $G^v(\rho^v)$, instead of $\bar{G}^v(\rho^v)$, for the vector of the fractions in which the inflow of node v gets split into its outgoing links.

We are now ready to define a dynamical flow network.

Definition 3: A *dynamical flow network* associated to a flow network \mathcal{N} satisfying Assumption 1, a distributed routing policy \mathcal{G} , and an inflow $\lambda_0 \geq 0$, is the dynamical system

$$\frac{d}{dt} \rho_e(t) = \lambda_v(t) G_e^v(\rho(t)) - f_e(t), \quad \forall 0 \leq v < n, \quad \forall e \in \mathcal{E}_v^+, \quad (6)$$

where $f_e(t) := \mu_e(\rho_e(t))$ and

$$\lambda_v(t) := \begin{cases} \lambda_0 & \text{if } v = 0 \\ \sum_{e \in \mathcal{E}_v^-} f_e(t) & \text{if } 0 < v \leq n. \end{cases} \quad (7)$$

Equation (6) states that the rate of variation of the particle density on a link e outgoing from some non-destination node v is given by the difference between $\lambda_v(t) G_e^v(\rho(t))$, i.e., the portion of the inflow at node v which is routed to link e , and $f_e(t)$, i.e., the particle flow on link e . Observe that the (distributed) routing policy $G^v(\rho)$ induces a (local) feedback which couples the dynamics of the particle flow on the the different links. We now introduce the following notion of transfer efficiency of a dynamical flow network.

Definition 4: Consider a dynamical flow network \mathcal{N} satisfying Assumptions 1 and 2. Given some flow vector $f^\circ \in \mathcal{F}$, and $\alpha \in [0, 1]$, the dynamical flow network (6) is said to be α -*transferring* with respect to f° if the solution of (6) with initial condition $\rho(0) = \mu^{-1}(f^\circ)$ satisfies

$$\liminf_{t \rightarrow +\infty} \lambda_n(t) \geq \alpha \lambda_0. \quad (8)$$

In particular, a dynamical flow network is said to be *partially transferring* if the total inflow at the destination node is asymptotically bounded away from zero, i.e., if it is α -transferring for any $\alpha \in (0, 1]$.

Remark 1: Standard definitions in the literature are typically limited to static flow networks describing the particle flow at equilibrium via conservation of mass. In fact, they usually consist (see e.g., [1]) in the specification of a topology \mathcal{T} , a vector of flow capacities $f^{\max} \in \mathcal{R}$, and an admissible equilibrium flow vector $f^* \in \mathcal{F}^*(\lambda_0)$ for $\lambda_0 < C(\mathcal{N})$ (or, often, $f^* \in \text{cl}(\mathcal{F}^*(\lambda_0))$ for $\lambda_0 \leq C(\mathcal{N})$). In contrast, in our model, we focus on the off-equilibrium particle dynamics on a flow network \mathcal{N} , induced by a (distributed) routing policy \mathcal{G} .

We now present an illustrative application of the dynamical flow network framework in the context of transportation

networks; connections to other application domains are explained in [12].

In transportation networks, particles represent drivers and distributed routing policies correspond to their local route choice behavior in response to the locally observed link congestions. A desired route choice behavior from a social optimization perspective may be achieved by appropriate incentive mechanisms. The robust distributed routing policies designed in this paper would correspond to the *ideal* node-wise route choice behavior of the drivers. The flow function $\mu_e(\rho_e)$ presented in this paper is related to the notion of fundamental diagram in traffic theory, e.g., see [14]. Note that in our formulation, we assume that the density of drivers is homogeneous over a link. One can refer to [14] for models that incorporate inhomogeneity, although such models are developed under non-feedback routing policies.

Remark 2: While there are many examples of congestion-dependent throughput functions and clearing functions that satisfy Assumption 2, typical fundamental diagrams in transportation systems have a \cap -shaped profile. While we do not study the implications of this analytically, some simulations are provided in [15] illustrating how the results of this paper could be extended to this case.

Remark 3: It is worth stressing that, while distributed routing policies depend only on local information on the current congestion, their structural form may depend on some global information on the flow network which might have been accumulated through a slower time-scale evolutionary dynamics. A two time-scale process of this sort has been analyzed in our related work [16] in the context of transportation networks. Multiple time-scale dynamical processes have also been analyzed in [17] in the context of communication networks.

B. Perturbed dynamical flow networks and resilience

We shall consider persistent perturbations of the dynamical flow network (6) that reduce the flow functions on the links, as per the following¹:

Definition 5: An *admissible perturbation* of a flow network $\mathcal{N} = (\mathcal{T}, \mu)$, satisfying Assumptions 1 and 2, is a flow network $\tilde{\mathcal{N}} = (\mathcal{T}, \tilde{\mu})$, with the same topology \mathcal{T} , and a family of perturbed flow functions $\tilde{\mu} := \{\tilde{\mu}_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+\}_{e \in \mathcal{E}}$, such that, for every $e \in \mathcal{E}$, $\tilde{\mu}_e$ satisfies Assumption 2, as well as $\tilde{\mu}_e(\rho_e) \leq \mu_e(\rho_e)$ for all $\rho_e \geq 0$. We accordingly let $\tilde{f}_e^{\max} := \sup\{\tilde{\mu}_e(\tilde{\rho}_e) : \tilde{\rho}_e \geq 0\}$. The *magnitude* of an admissible perturbation is defined as

$$\delta := \sum_{e \in \mathcal{E}} \delta_e, \quad \delta_e := \sup\{\mu_e(\rho_e) - \tilde{\mu}_e(\rho_e) : \rho_e \geq 0\}. \quad (9)$$

The *stretching coefficient* of an admissible perturbation is defined as

$$\theta := \max\{\tilde{\rho}_e^\mu / \rho_e^\mu : e \in \mathcal{E}\}, \quad (10)$$

where ρ_e^μ , and $\tilde{\rho}_e^\mu$ are the median densities associated to the unperturbed and the perturbed flow functions, respectively, on link $e \in \mathcal{E}$, as defined in (5).

Given a dynamical flow network as in Definition 3, and an admissible perturbation as in Definition 5, we shall consider

¹In this paper, we use *disturbance* and *perturbation* interchangeably.

the *perturbed dynamical flow network*

$$\frac{d}{dt}\tilde{\rho}_e(t) = \tilde{\lambda}_v(t)G_e^v(\tilde{\rho}(t)) - \tilde{f}_e(t), \quad \forall 0 \leq v < n, \quad \forall e \in \mathcal{E}_v^+, \quad (11)$$

where $\tilde{f}_e(t) := \tilde{\mu}_e(\tilde{\rho}_e(t))$ and

$$\tilde{\lambda}_v(t) := \begin{cases} \sum_{e \in \mathcal{E}_v^-} \tilde{f}_e(t) & \text{if } 0 < v < n \\ \lambda_0 & \text{if } v = 0. \end{cases} \quad (12)$$

Observe that the perturbed dynamical flow network (11) has the same structure of the original dynamical flow network (6), as it describes the rate of variation of the particle density on each link e outgoing from some non-destination node v as the difference between $\tilde{\lambda}_v(t)G_e^v(\tilde{\rho}(t))$, i.e., the portion of the perturbed inflow at node v which is routed to link e , minus the perturbed flow on link e itself. Notice that the only difference with respect to the original dynamical flow network (6) is in the perturbed flow function $\tilde{\mu}_e(\rho_e)$ on each link $e \in \mathcal{E}$, which replaces the original one, $\mu_e(\rho_e)$. In particular, the (distributed) routing policy \mathcal{G} is the same for the unperturbed and the perturbed dynamical flow networks. In this way, we model a situation in which the routers are not aware of the fact that the flow network has been perturbed, but react to this change only indirectly, in response to variations of the local density vectors $\tilde{\rho}^v(t)$.

We are now ready to define the following notion of resilience of a dynamical flow network as in Definition 3 with respect to an initial flow.

Definition 6: Let \mathcal{N} be a flow network satisfying Assumptions 1 and 2, \mathcal{G} be a distributed routing policy, and $\lambda_0 \geq 0$ be a constant inflow at the origin node. Given $\alpha \in (0, 1]$, $\theta \geq 1$ and $f^\circ \in \mathcal{F}$, let $\gamma_{\alpha, \theta}(f^\circ, \mathcal{G})$ be equal to the infimum magnitude of all the admissible perturbations of stretching coefficient less than or equal to θ for which the perturbed dynamical flow network (11) is not α -transferring with respect to f° . Also, define $\gamma_{0, \theta}(f^\circ, \mathcal{G}) := \lim_{\alpha \downarrow 0} \gamma_{\alpha, \theta}(f^\circ, \mathcal{G})$. For $\alpha \in [0, 1]$, the α -resilience with respect to f° is defined as $\gamma_\alpha(f^\circ, \mathcal{G}) := \lim_{\theta \rightarrow +\infty} \gamma_{\alpha, \theta}(f^\circ, \mathcal{G})$. The 0-resilience will be referred to as the *weak resilience*. The 1-resilience, which is the focus of our related work [15] is referred to as the *strong resilience*.

Remark 4: The notions of resilience are with respect to adversarial perturbations. Therefore, one can provide a zero-sum game interpretation as follows. Let the strategy space of the system planner be the class of distributed routing policies and the strategy space of an adversary be the set of admissible perturbations. Let the utility function of the adversary be $M\Theta - \delta$, where M is a large quantity, e.g., $\sum_{e \in \mathcal{E}} f_e^{\max}$, and Θ takes the value 1 if the network is not α -transferring under given strategies of the system planner and the adversary, and zero otherwise. Let the utility function of the system planner be $\delta - M\Theta$. As stated in Section III, a certain class of *locally responsive* distributed routing policies characterized by Definition 7, is maximally robust. This will then show that the locally responsive distributed

routing policies correspond to approximate Nash equilibria in this zero-sum game setting.

In the remainder of the paper, we shall focus on the characterization of the weak resilience of dynamical flow networks. Before proceeding, let us elaborate a bit on Definition 6. Notice that, for every $\alpha \in (0, 1]$, the α -resilience $\gamma_\alpha(f^\circ, \mathcal{G})$ is simply the infimum magnitude of all the admissible perturbations such that the perturbed dynamical network (11) is not α -transferring with respect to the equilibrium flow f° . In fact, one might think of $\gamma_\alpha(f^\circ, \mathcal{G})$ as the minimum effort required by a hypothetical adversary in order to modify the dynamical flow network from (6) to (11), and make it not α -transferring, provided that such an effort is measured in terms of the magnitude of the perturbation $\delta = \sum_{e \in \mathcal{E}} \|\mu_e(\cdot) - \tilde{\mu}_e(\cdot)\|_\infty$. For $\alpha = 0$, trivially the perturbed network flow is always 0-transferring with respect to any initial flow. For this reason, the definition of the weak resilience $\gamma_0(f^\circ, \mathcal{G})$ involves the double limit $\lim_{\theta \rightarrow +\infty} \lim_{\alpha \downarrow 0} \gamma_{\alpha, \theta}(f^\circ, \mathcal{G})$: the introduction of the bound on the stretching coefficient of the admissible perturbation is a mere technicality whose necessity will become clear in Section IV.

We conclude this section with the following result, proven in [12], providing an upper bound on the weak resilience of a dynamical flow network driven by any, not necessarily distributed, routing policy \mathcal{G} , in terms of the min-cut capacity of the network. Tightness of this bound will follow from Theorem 2 in Section III, which will show that, for a particular class of locally responsive distributed routing policies, the dynamical flow network has weak resilience equal to the min-cut capacity.

Proposition 1: Let \mathcal{N} be a flow network satisfying Assumptions 1 and 2, $\lambda_0 > 0$ a constant inflow, and \mathcal{G} an arbitrary routing policy. Then, for any initial flow f° , the weak resilience of the associated dynamical flow network satisfies $\gamma_0(f^\circ, \mathcal{G}) \leq C(\mathcal{N})$.

III. MAIN RESULTS AND DISCUSSION

In this paper, we shall be concerned with a family of *maximally robust* distributed routing policies. Such a family is characterized by the following:

Definition 7: A *locally responsive* distributed routing policy for a flow network topology $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ with node set $\mathcal{V} = \{0, 1, \dots, n\}$ is a family of continuously differentiable distributed routing functions $\mathcal{G} = \{G^v : \mathcal{R}_v \rightarrow \mathcal{S}_v\}_{v \in \mathcal{V}}$ such that, for every non-destination node $0 \leq v < n$:

- (a) $\frac{\partial}{\partial \rho_e} G_j^v(\rho^v) \geq 0, \quad \forall j, e \in \mathcal{E}_v^+, j \neq e, \rho^v \in \mathcal{R}^v;$
- (b) for every nonempty proper subset $\mathcal{J} \subsetneq \mathcal{E}_v^+$, there exists a continuously differentiable map $G^{\mathcal{J}} : \mathcal{R}_{\mathcal{J}} \rightarrow \mathcal{S}_{\mathcal{J}}$, where $\mathcal{R}_{\mathcal{J}} := \mathbb{R}_+^{\mathcal{J}}$, and $\mathcal{S}_{\mathcal{J}} := \{p \in \mathcal{R}_{\mathcal{J}} : \sum_{j \in \mathcal{J}} p_j = 1\}$ is the simplex of probability vectors over \mathcal{J} , such that, for every $\rho^{\mathcal{J}} \in \mathcal{R}_{\mathcal{J}}$, if $\rho_e^v \rightarrow +\infty, \quad \forall e \in \mathcal{E}_v^+ \setminus \mathcal{J}, \rho_j^v \rightarrow \rho_j^{\mathcal{J}}, \quad \forall j \in \mathcal{J}$, then

$$G_e^v(\rho^v) \rightarrow 0, \quad \forall e \in \mathcal{E}_v^+ \setminus \mathcal{J}, \quad G_j^v(\rho^v) \rightarrow G_j^{\mathcal{J}}(\rho^{\mathcal{J}}), \quad \forall j \in \mathcal{J}.$$

Property (a) in Definition 7 states that, as the particle density on an outgoing link $e \in \mathcal{E}_v^+$ increases while the

²It is easily seen that the limits involved in this definition always exist, as $\gamma_{\alpha, \theta}(f^\circ, \mathcal{G})$ is clearly nonincreasing in α (the higher α , the more stringent the requirement of α -transfer) and θ (the higher θ , the more admissible perturbations are considered that may potentially make the dynamical flow network to be not α -transferring).

particle density on all the other outgoing links remains constant, the fraction of inflow at node v routed to any link $j \in \mathcal{E}_v^+ \setminus \{e\}$ does not decrease, and hence the fraction of inflow routed to link e itself does not increase. In fact, Property (a) in Definition 7 is reminiscent of Hirsch's notion of *cooperative dynamical systems* [2], [3]. On the other hand, Property (b) implies that the fraction of incoming particle flow routed to a subset of outgoing links $\mathcal{K} \subset \mathcal{E}_v^+$ vanishes as the density on links in \mathcal{K} grows unbounded while the density on the remaining outgoing links remains bounded.

Example 1: Let η_v , for $0 \leq v < n$, and a_e , for $e \in \mathcal{E}$, be positive constants. Define the routing policy \mathcal{G} by

$$G_e^v(\rho) = \frac{a_e \exp(-\eta_v \rho_e)}{\sum_{j \in \mathcal{E}_v^+} a_j \exp(-\eta_v \rho_j)}, \quad \forall e \in \mathcal{E}_v^+, \forall 0 \leq v < n. \quad (13)$$

Clearly, \mathcal{G} is distributed, as it uses only information on the particle density on the links outgoing from a node v in order to compute how the inflow at node v gets split among its outgoing links. It is easy to check that (13) satisfies the properties of locally responsive routing policies. In the context of transportation networks, the example in (13) is a variant of the logit function from discrete choice theory emerging from utilization maximization perspective of drivers, where the utility associated with link e is the sum of $-\rho_e + \log a_e / \eta_v$ and a double exponential random variable with parameter η_v (see, e.g., [18]).

We are now ready to state our main results.

Theorem 1: Let \mathcal{N} be a flow network satisfying Assumptions 1 and 2, $\lambda_0 \geq 0$ a constant inflow, and \mathcal{G} a locally responsive distributed routing policy. Then, there exists a unique limit flow $f^* \in \text{cl}(\mathcal{F})$ such that, for every initial condition $\rho(0) \in \mathcal{R}$, the dynamical flow network (6) satisfies $\lim_{t \rightarrow +\infty} f(t) = f^*$.

Theorem 1, which is proven in [12], states that, when the routing policy is distributed and locally responsive, there is a unique globally attractive limit flow f^* . Such a limit flow may be in \mathcal{F} , in which case it is not hard to see that it is necessarily an equilibrium flow, i.e., $f^* \in \mathcal{F}^*(\lambda_0)$ (because of the continuity of the right hand side of (6)); or belong to $\text{cl}(\mathcal{F}) \setminus \mathcal{F}$, i.e., it satisfies the capacity constraint on one link with equality, in which case it is not an equilibrium flow. The global convergence result mainly relies on Assumption 2 on monotonicity of the flow function, and Property (a) of Definition 7 of locally responsive distributed routing policies, from which the dynamical flow network (6) inherits a cooperative property. It is worth mentioning that we shall not use general results for cooperative dynamical systems [2], [3], [19], but rather exploit some other structural properties of (6) which in fact allow us to prove stronger results.

Our second main result stated below, whose proof is in Section IV, shows that locally responsive distributed routing policies are maximally robust, as the resilience of the induced dynamical flow network coincides with the min-cut capacity of the network.

Theorem 2: Let \mathcal{N} be a flow network satisfying Assumptions 1 and 2, $\lambda_0 > 0$ a constant inflow, and \mathcal{G} a locally responsive distributed routing policy such that $G_e^v(\rho^v) > 0$ for all $0 \leq v < n$, $e \in \mathcal{E}_v^+$, and $\rho^v \in \mathcal{R}_v$. Then, for every $f^\circ \in \mathcal{F}$, the associated dynamical flow network is partially

transferring with respect to f° and has weak resilience $\gamma_0(f^\circ, \mathcal{G}) = C(\mathcal{N})$.

Theorem 2, combined with Proposition 1, shows that locally responsive distributed routing policies achieve the maximal weak resilience possible on a given flow network \mathcal{N} . A consequence of this result is that locality constraints on the feedback information available to routing policies do not reduce the achievable weak resilience. It is also worth observing that such maximal weak resilience coincides with min-cut capacity of the network, and is therefore independent of the initial flow f° . This is in sharp contrast with the results on the strong resilience of dynamical flow networks presented in [15]. There, it is shown that the strong resilience depends on the limit flow of the unperturbed system, and local information constraints reduce the maximal strong resilience achievable on a given flow network.

IV. PROOF OF THEOREM 2

To start with, let us recall that, in this case, Theorem 1 implies the existence of a globally attractive limit flow $\tilde{f}^* \in \text{cl}(\mathcal{F})$ for the perturbed dynamical flow network associated to any admissible perturbation $\tilde{\mathcal{N}}$. Define $\tilde{\lambda}_0^* = \lambda_0$, and $\tilde{\lambda}_v^* = \sum_{e \in \mathcal{E}_v^-} \tilde{f}_e^*$, for $0 < v \leq n$.

Lemma 1: Consider a dynamical flow network \mathcal{N} satisfying Assumptions 1 and 2, with locally responsive distributed routing policy \mathcal{G} such that $G_e^v(\rho^v) > 0$ for all $0 \leq v < n$, $e \in \mathcal{E}_v^+$, and $\rho^v \in \mathcal{R}_v$. Then, for every $\theta \geq 1$, there exists $\beta_\theta \in (0, 1)$ such that, if $\tilde{\mathcal{N}}$ is an admissible perturbation of \mathcal{N} with stretching coefficient less than or equal to θ , and $\tilde{f}^* \in \text{cl}(\tilde{\mathcal{F}})$ is the limit flow vector of the corresponding perturbed dynamical flow network (11), then $\tilde{f}_e^* \geq \beta_\theta \tilde{\lambda}_v^*$, for every non-destination node $0 \leq v < n$, and every link $e \in \mathcal{E}_v^+$ for which $\tilde{f}_e^* \leq \tilde{f}_e^{\max}/2$.

As a consequence of Lemma 1, we now prove the following result showing that the dynamical flow network is partially transferring and providing a lower bound on its weak resilience:

Lemma 2: Let \mathcal{N} be a flow network satisfying Assumptions 1 and 2, $\lambda_0 \geq 0$ a constant inflow, and \mathcal{G} a locally responsive distributed routing policy such that $G_e^v(\rho^v) > 0$ for all $0 \leq v < n$, $e \in \mathcal{E}_v^+$, and $\rho^v \in \mathcal{R}_v$. Then, the associated dynamical flow network is partially transferring, and, for every $\theta \geq 1$, and $\alpha \in (0, \beta_\theta^n]$, its resilience satisfies $\gamma_{\alpha, \theta}(f^*, \mathcal{G}) \geq C(\mathcal{N}) - 2|\mathcal{E}|\lambda_0\beta_\theta^{1-n}\alpha$, where $\beta_\theta \in (0, 1)$ is as in Lemma 1.

Proof Consider an arbitrary admissible perturbation $\tilde{\mathcal{N}}$ of magnitude

$$\delta \leq C(\mathcal{N}) - 2|\mathcal{E}|\lambda_0\beta_\theta^{1-n}\alpha, \quad (14)$$

and stretching coefficient less than or equal to θ . We shall iteratively select a sequence of nodes $0 =: v_0, v_1, \dots, v_k := n$ such that, for every $1 \leq j \leq k$,

$$\exists i \in \{0, \dots, j-1\} \text{ s.t. } (v_i, v_j) \in \mathcal{E}, \tilde{f}_{(v_i, v_j)}^* \geq \lambda_0 \alpha \beta_\theta^{j-n}. \quad (15)$$

Since $v_k = n$, and $\beta_\theta^{k-n} \geq 1$, the above with $j = k \leq n$ will immediately imply that

$$\lim_{t \rightarrow +\infty} \tilde{\lambda}_n(t) = \tilde{\lambda}_n^* = \sum_{e \in \mathcal{E}_n^-} \tilde{f}_e^* \geq \alpha \lambda_0 \beta_\theta^{k-n} \geq \alpha \lambda_0, \quad (16)$$

so that the perturbed dynamical flow network is α -transferring. For $0 < \alpha \leq \beta_\theta^{n-1}/(2|\mathcal{E}|\lambda_0)$, one could chose a trivial perturbation $\tilde{\mathcal{N}} = \mathcal{N}$ so that (16) would imply the partial transferring property of the original dynamical flow network. Moreover, the rest of the claim will then readily follow from the arbitrariness of the considered admissible perturbation.

First, let us consider the case $j = 1$. Assume by contradiction that $\tilde{f}_e^* < \lambda_0 \alpha \beta_\theta^{1-n}$, for every link $e \in \mathcal{E}_0^+$. Since $\alpha \leq \beta_\theta^n$, this would imply that $\tilde{f}_e^* < \beta_\theta \lambda_0$ and hence, by Lemma 1, that $\tilde{f}_e^{\max} \leq 2\tilde{f}_e^*$ for all $e \in \mathcal{E}_0^+$, so that $\sum_{e \in \mathcal{E}_0^+} \tilde{f}_e^{\max} \leq 2 \sum_{e \in \mathcal{E}_0^+} \tilde{f}_e^* < 2\alpha |\mathcal{E}_0^+| \beta_\theta^{1-n} \lambda_0 \leq 2\alpha |\mathcal{E}| \beta_\theta^{1-n} \lambda_0$. Combining the above with the inequality $C(\mathcal{N}) \leq \sum_{e \in \mathcal{E}_0^+} \tilde{f}_e^{\max}$, one would get $\delta \geq \sum_{e \in \mathcal{E}_0^+} (f_e^{\max} - \tilde{f}_e^{\max}) > C(\mathcal{N}) - 2\alpha |\mathcal{E}| \beta_\theta^{1-n} \lambda_0$, thus contradicting the assumption (14). Hence, necessarily there exists $e \in \mathcal{E}_0^+$ such that $\tilde{f}_e^* \geq \lambda_0 \alpha \beta_\theta^{1-n}$, and choosing v_1 to be the unique node in \mathcal{V} such that $e \in \mathcal{E}_{v_1}^-$, one sees that (15) holds true with $j = 1$.

Now, fix some $1 < j^* \leq k$, and assume that (15) holds true for every $1 \leq j < j^*$. Then, by choosing i as in (15), $\tilde{\lambda}_{v_j}^* = \sum_{e \in \mathcal{E}_{v_j}^+} \tilde{f}_e^* \geq \tilde{f}_{(v_i, v_j)}^* \geq \lambda_0 \alpha \beta_\theta^{j-n} \geq \lambda_0 \alpha \beta_\theta^{j^*-1-n}$, $\forall 1 \leq j < j^*$.

Moreover,

$$\tilde{\lambda}_{v_0}^* = \lambda_0 > \lambda_0 \alpha \beta_\theta^{-n} \geq \lambda_0 \alpha \beta_\theta^{j^*-1-n}. \quad (18)$$

Let $\mathcal{U} := \{v_0, v_1, \dots, v_{j^*-1}\}$ and $\mathcal{E}_\mathcal{U}^+ \subseteq \mathcal{E}$ be the set of links with tail node in \mathcal{U} and head node in $\mathcal{V} \setminus \mathcal{U}$. Assume by contradiction that $\tilde{f}_e^* < \lambda_0 \alpha \beta_\theta^{j^*-n}$ for all $e \in \mathcal{E}_\mathcal{U}^+$. Thanks to (17) and (18), this would imply that, $\tilde{f}_e^* < \beta_\theta \lambda_{v_j}^*$, for every $0 \leq j < j^*$ and $e \in \mathcal{E}_{v_j}^+ \cap \mathcal{E}_\mathcal{U}^+$. Then, since $\mathcal{E}_\mathcal{U}^+ = \cup_{j=0}^{j^*-1} (\mathcal{E}_{v_j}^+ \cap \mathcal{E}_\mathcal{U}^+)$, Lemma 1 would imply that $\tilde{f}_e^{\max} \leq 2\tilde{f}_e^*$, for every $e \in \mathcal{E}_\mathcal{U}^+$. This would yield $\sum_{e \in \mathcal{E}_\mathcal{U}^+} \tilde{f}_e^{\max} \leq \sum_{e \in \mathcal{E}_\mathcal{U}^+} 2\tilde{f}_e^* < 2 \sum_{e \in \mathcal{E}_\mathcal{U}^+} \lambda_0 \alpha \beta_\theta^{j^*-n} \leq 2|\mathcal{E}| \lambda_0 \alpha \beta_\theta^{1-n}$. From this and the inequality $C(\mathcal{N}) \leq \sum_{e \in \mathcal{E}_\mathcal{U}^+} \tilde{f}_e^{\max}$, one would get $\delta \geq \sum_{e \in \mathcal{E}_\mathcal{U}^+} (f_e^{\max} - \tilde{f}_e^{\max}) > C(\mathcal{N}) - 2\alpha |\mathcal{E}| \beta_\theta^{1-n} \lambda_0$, thus contradicting the assumption (14). Hence, necessarily there exists $e \in \mathcal{E}_\mathcal{U}^+$ such that $\tilde{f}_e^* \geq \lambda_0 \alpha \beta_\theta^{1-n}$, and choosing v_{j^*} to be the unique node in \mathcal{V} such that $e \in \mathcal{E}_{v_{j^*}}^-$ one sees that (15) holds true with $j = j^*$. Iterating this argument until $v_{j^*} = n$ proves the claim. ■

It is now easy to see that Lemma 2 implies that $\lim_{\alpha \downarrow 0} \gamma_{\alpha, \theta} \geq C(\mathcal{N})$ for every $\theta \geq 1$, thus showing that $\gamma_0(f^\circ, \mathcal{G}) \geq C(\mathcal{N})$. Combined with Proposition 1, this shows that $\gamma_0(f^\circ, \mathcal{G}) = C(\mathcal{N})$, thus completing the proof of Theorem 2.

V. CONCLUSION

In this paper, we studied robustness properties of dynamical flow networks. We proposed a class of locally responsive distributed routing policies that rely only on local information about the network's current particle densities and yield the maximum weak resilience with respect to adversarial disturbances that reduce the flow functions of

the links of the network. We also showed that the weak resilience of the network in that case is equal to min-cut capacity of the network, and that it is independent of the local information constraint and the initial flow. Results for an alternate notion of resilience have been reported in our other work [15]. These findings stand to provide important guidelines for robust real-time operation of several large scale critical infrastructure systems.

In future, we plan to rigorously study the robustness properties of the network with finite link-wise capacity for the densities. We also plan to consider other general models for disturbances, including sequential disturbances than just one-shot disturbance considered in this paper. We also plan to consider more general graph topologies, e.g., graphs having cycles and multiple origin-destination pairs.

REFERENCES

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [2] M. W. Hirsch, "Systems of differential equations that are competitive or cooperative. I: Limit sets," *SIAM Journal on Mathematical Analysis*, vol. 13, pp. 167–179, 1982.
- [3] M. W. Hirsch, "Systems of differential equations that are competitive or cooperative II: Convergence almost everywhere," *SIAM Journal on Mathematical Analysis*, vol. 16, pp. 423–439, 1985.
- [4] G. Vinnicombe, "Robust congestion control for the internet," in *SIGCOMM*, 2002.
- [5] F. Paganini, "A global stability result in network flow control," *Systems and Control Letters*, vol. 46, no. 3, pp. 165–172, 2002.
- [6] S. H. Low, F. Paganini, and J. C. Doyle, "Internet congestion control," *IEEE Control Systems Magazine*, vol. 22, no. 1, pp. 28–43, 2002.
- [7] X. Fan, M. Arcak, and J. T. Wen, "Robustness of network flow control against disturbances and time-delay," *Systems and Control Letters*, vol. 53, no. 1, pp. 13–29, 2004.
- [8] Y. Vardi and C. Zhang, "Measures of network vulnerability," *IEEE Signal Processing Letters*, vol. 14, pp. 313–316, May 2007.
- [9] M. Sengoku, S. Shinoda, and R. Yatsuboshi, "On a function for the vulnerability of a directed flow network," *Networks*, vol. 18, no. 1, pp. 73–83, 1988.
- [10] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936–1948, 1992.
- [11] D. Bertsekas and R. Gallager, *Data Networks*. Prentice Hall, 2nd ed., 1992.
- [12] G. Como, K. Savla, D. Acemoglu, M. A. Dahleh, and E. Frazzoli, "Robust distributed routing in dynamical flow networks. Part I: locally responsive policies and weak resilience," *IEEE Transactions on Automatic Control*, 2011. Conditionally accepted, Available at <http://arxiv.org/abs/1102.1107>.
- [13] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*. MIT Press, 2nd ed., 2001.
- [14] M. Garavello and B. Piccoli, *Traffic Flow on Networks*. American Institute of Mathematical Sciences, 2006.
- [15] G. Como, K. Savla, D. Acemoglu, M. A. Dahleh, and E. Frazzoli, "Robust distributed routing in dynamical flow networks. part II: strong resilience, equilibrium selection and cascaded failures," *IEEE Transactions on Automatic Control*, 2011. Conditionally accepted, Available at <http://arxiv.org/abs/1103.4893>.
- [16] G. Como, K. Savla, D. Acemoglu, M. A. Dahleh, and E. Frazzoli, "Stability analysis of transportation networks with multiscale driver decisions," *SIAM Journal on Control and Optimization*, 2011. Submitted, Available at <http://arxiv.org/abs/1101.2220>.
- [17] V. S. Borkar and P. R. Kumar, "Dynamic Cesaro-Wardrop equilibration in networks," *IEEE Transactions on Automatic Control*, vol. 48, no. 3, pp. 382–396, 2003.
- [18] M. Ben-Akiva and S. Lerman, *Discrete Choice Analysis: Theory and application to travel demand*. MIT press, 1985.
- [19] H. L. Smith, *Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems*. American Mathematical Society, 1995.