POLITECNICO DI TORINO
Repository ISTITUZIONALE

Robust distributed routing in dynamical networks-Part I: Locally responsive policies and weak resilience

Original
Robust distributed routing in dynamical networks-Part I: Locally responsive policies and weak resilience / Como, G., Savla, K., Acemoglu, D., Dahleh, M.A., Frazzoli, E.. - In: IEEE TRANSACTIONS ON AUTOMATIC CONTROL. - ISSN 0018-9286. - 58:2(2013), pp. 317-332. [10.1109/TAC.2012.2209951]

Availability:
This version is available at: 11583/2624515 since: 2015-11-30T17:29:33Z

Publisher:
IEEE

Published
DOI:10.1109/TAC.2012.2209951

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

15 November 2026

Robust Distributed Routing in Dynamical Networks – Part I: Locally Responsive Policies and Weak Resilience

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

***Abstract*—Robustness of distributed routing policies is studied for dynamical networks, with respect to adversarial disturbances that reduce the link flow capacities. A dynamical 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 total outflow at the origin. Routing policies regulate the way the total outflow at each 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 dynamical network is called partially transferring if the total inflow at the destination node is asymptotically bounded away from zero, and its weak resilience is measured as the minimum sum of the link-wise magnitude of disturbances that make it not partially transferring. The weak resilience of a dynamical 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. Fundamental properties of dynamical networks driven by locally responsive distributed routing policies are analyzed in detail, including global convergence to a unique limit flow. The derivation of these properties exploits the cooperative nature of these dynamical systems, together with an additional stability property implied by the assumption of monotonicity of the flow as a function of the density on each link.**

Index terms: dynamical networks, distributed routing policies, weak resilience, min-cut capacity, cooperative dynamical systems, monotone control systems.

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

K. Savla is with the Sonny Astani Department of Civil and Environmental Engineering at the University of Southern California. ksavla@usc.edu. This work was performed when he was with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology.

D. Acemoglu is with the Department of Economics at the Massachusetts Institute of Technology. daron@mit.edu.

M.A. Dahleh and E. Frazzoli are with the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. {dahleh,frazzoli}@mit.edu.

This work was supported in part by NSF EFRI-ARES grant number 0735956 and AFOSR grant number FA9550-09-1-0538. The work of the first author was partially supported by the Swedish Research Council through the LCCC Linnaeus Center. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the supporting organizations. A preliminary version of this paper appeared in part as [1].

I. Introduction

Network flows provide a fruitful modeling framework for transport phenomena, with many applications of interest, e.g., road traffic, 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 [2], for a thorough treatment.

The present and a companion paper [3] study dynamical formulations of flows over networks. In particular, we study *dynamical 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 total outflow 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 total outflow at a non-destination node gets split among its outgoing links depends on the current particle density, possibly on the whole network, through a *routing policy*. A routing policy is said to be *distributed* if the proportion of total outflow 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: in road traffic networks [4], the flow functions are typically referred to as *fundamental diagrams*; in data networks [5], [6], [7], flow functions model congestion-dependent *throughput* and *average delays*, while routing policies are designed in order to optimize the total throughput or other performance measures; in production networks [8], [9], flow functions correspond to *clearing functions*. As for the routing policies, in data and production networks they have to be thought as suitably designed distributed feedback controls. On the other hand, in road traffic networks routing policies are meant to describe the selfish dynamic route choice behavior of the drivers adapting to the current congestion levels of the network.

Our objective is the analysis and design of distributed routing policies for dynamical networks that are *maximally robust* with respect to *adversarial disturbances* that reduce the link flow capacities. Two notions of transfer efficiency are introduced in order to capture the extremes of the resilience of the network towards disturbances: the dynamical network

is *fully transferring* if the total inflow at the destination node asymptotically approaches the total outflow at the origin node, and *partially transferring* if the total inflow at the destination node is asymptotically bounded away from zero. The robustness of distributed routing policies is evaluated in terms of the network's *strong* and *weak resilience*, which are defined as the minimum sum of link-wise magnitude of disturbances making the perturbed dynamical network not fully transferring, and, respectively, not partially transferring. 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 total outflow that a node routes towards an outgoing link does not decrease as the particle density on any other outgoing link increases. Moreover, we show that the maximum weak resilience of dynamical networks with arbitrary, not necessarily distributed, routing policies equals the *min-cut capacity* of the network and hence is independent of the initial flow.

The contributions of this paper are as follows: (i) we formulate a novel dynamical system framework for robustness analysis of network flows under feedback routing policies that are possibly constrained in the available information; (ii) we introduce a rich class of locally responsive distributed routing policies that yield the maximum weak resilience; (iii) we provide a simple characterization of the resilience in terms of the topology and capacity of the network. In particular, the class of locally responsive distributed routing policies can be interpreted as approximate Nash equilibria in an appropriate zero-sum game setting where the objective of the adversary inflicting the disturbance is to make the network not partially transferring with a disturbance of minimum possible magnitude, and the objective of the system planner is to design distributed routing policies that yield the maximum possible resilience. The results of this paper imply that locality constraints on the information available to routing policies do not affect the maximally achievable weak resilience. In contrast, in the companion paper [3], we focus on the strong resilience properties of dynamical networks, and show that locally responsive distributed routing policies are maximally robust, but only within the class of distributed routing policies which are constrained to use only local information on the particle densities.

The main technical assumptions in our model are that the network topology is acyclic and contains a single origin-destination pair; that there is no bound on the density on the links; and that, on every link, the flow is monotonically increasing in the density. The acyclicity assumption does not cause serious limitations to the applicability of our results as long as one is dealing, as we are, with a single origin-destination pair. However, such assumption limits the generalizability of our results to scenarios with multiple origin-destination pairs, where the absence of cycles is harder to justify. On the other hand, the absence of an *a priori* bound on the particle density prevents the occurrence of backward effects, such as the so-called *bullwhip effect* often observed in production networks (see, e.g., [10]). Finally, monotonicity of the flow function is a reasonable assumption for production networks [10], as well as data networks, in particular the internet, for the existence of TCP/IP (traffic control protocol/internet protocol) congestion control procedures [6], [7]. In contrast, this assumption constitutes a major limitation in road traffic networks, where fundamental diagrams are typically assumed to have a $\cap$-shaped graph. However, in this application context, our results can be applied, provided that the density on each link remains on the interval in which the flow function is increasing, and thus allow one to obtain possibly conservative bounds on the resilience of road traffic networks. Also, it is worth stressing out that, in road traffic networks, local responsiveness of the distributed routing policies appears to be a very natural assumption for the behavior of drivers which naturally tend to choose a link with higher frequency the less congested it is. Similarly, local responsiveness appears to be an intuitive design guideline for distributed routing policies in data and production networks.

In the course of our analysis, we prove some fundamental properties of dynamical networks driven by locally responsive distributed policies, including global convergence to a unique limit flow. These results are mainly a consequence of the *cooperativeness* property (in the sense of Hirsch [11], [12], see also the recent survey [13]) which dynamical networks inherit from local responsiveness of the distributed routing policies. In particular, our proof of global convergence to a unique limit flow exploits the acyclicity assumption on the network topology in order to treat the dynamical network as a cascade of monotone local systems (in the spirit of [14]), whose input-output characteristics are established by an *ad hoc* contraction argument which makes careful use of the local responsiveness of the distributed routing policy, as well as of the monotonicity of the link flow functions.

Stability analysis of network flow control policies under non-persistent disturbances, especially in the context of the internet, has attracted a lot of attention, e.g., see [15], [16], [17], [18]. Robustness of the internet with respect to its architecture has been studied in [19], [20]. Recent work on robustness analysis of static network flows under adversarial and probabilistic persistent disturbances in the spirit of this paper include [21], [22], [23]. Our problem setup could also be considered as a dynamical and distributed version of the *network interdiction problem*, e.g., see [24], where the objective is to find the set of links to be removed from a network to maximize the reduction in its flow capacity subject to budget constraints on link removal. It is worth comparing the distributed routing policies studied in this paper with the back-pressure policy [25], 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 [26] [5, 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 networks, where the outflow from a link is not necessarily equal to its inflow have not been studied before.

It is also worth contrasting our work with the fluid-dynamical and kinetic models of transport networks as treated, e.g., in [4], [9], and references therein. As compared to these models (typically described by partial, or integro-differential equations), ours provide a much coarser description (treating particle density and flow as homogeneous quantities on the links, representative of spatial averages), whereas it highlights the role of the feedback routing policies, with possibly different levels of information, which is typically neglected in that literature.

Finally, we wish to stress once more that the notion of resilience we deal with in this work is with respect to deterministic adversarial disturbances. Based on similar analyses for other complex networks [19], [20], it is reasonable to expect that, in some cases, stochastic perturbations may guarantee better resilience for given probabilistic models of perturbations, as opposed to worst case scenario analyzed here. This point is not addressed here, but rather left as a topic for further research.

The rest of the paper is organized as follows. In Section II, we formulate the problem by formally defining the notion of a dynamical network and its resilience, and we prove that the weak resilience of a dynamical network driven by an arbitrary, not necessarily distributed, routing policy is upper-bounded by the min-cut capacity of the network. In Section III, we introduce the class of locally responsive distributed routing policies, and state the main results on dynamical networks driven by such locally responsive distributed routing policies: Theorem 1, concerning global convergence towards a unique equilibrium flow; and Theorem 2 concerning the maximal weak resilience property. In Sections IV, and V, we state proofs of Theorem 1 and Theorem 2 respectively.

Before proceeding, we define some preliminary notation 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. When $\mathcal{A}$ and $\mathcal{B}$ are finite sets, $|\mathcal{A}|$ 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 $\mathcal{A}$, and $\mathbb{R}^{\mathcal{A}\times\mathcal{B}}$ for 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 $\mathbf{1}$ will stand for the all-one vector, whose size will be clear from the context. Let $\mathrm{cl}(\mathcal{X})$ 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 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. 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 $\mathcal{E}$.

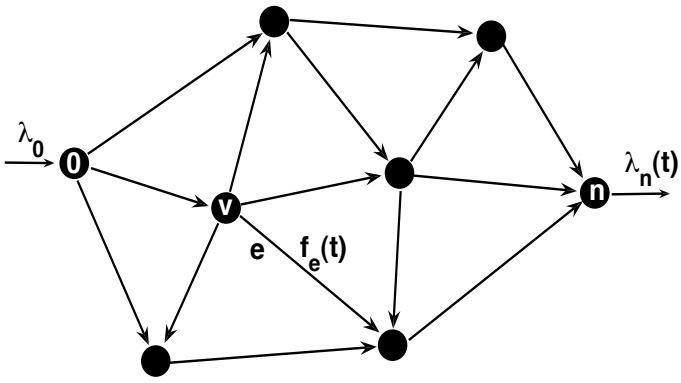


Fig. 1. A network topology satisfying Assumption 1: the nodes v are labeled by the integers between 0 (denoting the origin node) and n (denoting the destination node), in such a way that the label of the head node of each link is higher than the label of its tail node. The total outflow at the origin, λ_0, may be interpreted as the input to the dynamical network, and the total inflow at destination, $\lambda_n(t)$, as the output. For $\alpha \in [0, 1]$, the dynamical network is α-transferring if $\liminf_{t\to\infty} \lambda_n(t) \geq \alpha\lambda_0$, i.e., if at least an α-fraction of the total outflow at the origin is transferred to the destination, asymptotically.

II. Dynamical networks and their resilience

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

A. Dynamical networks

We start with the following definition.

Definition 1 (Network): A *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* $\mu := \{\mu_e : \mathbb{R}_+ \to \mathbb{R}_+\}_{e\in\mathcal{E}}$ 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

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

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

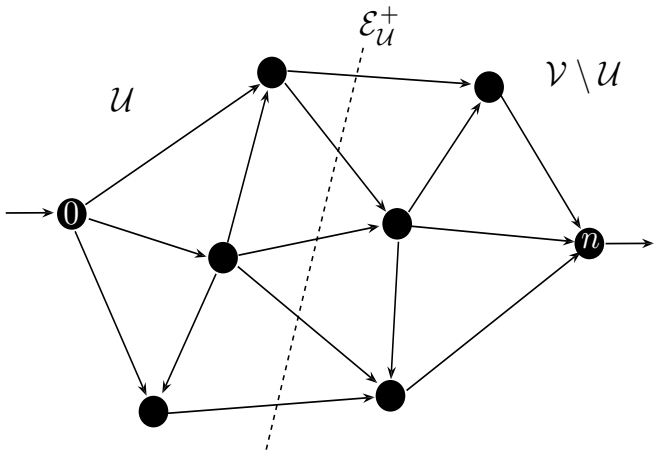


Fig. 2. An origin/destination cut of the network: $\mathcal{U}$ is a subset of nodes including the origin 0 but not the destination n, and $\mathcal{E}_{\mathcal{U}}^{+}$ is the subset of those links with tail node in $\mathcal{U}$, and head node in $\mathcal{V} \setminus \mathcal{U}$.

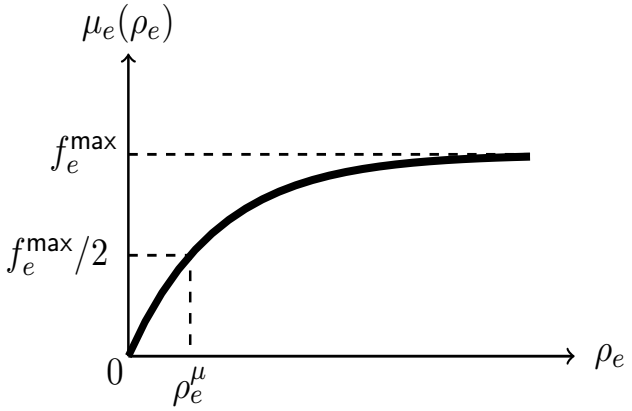


Fig. 3. Qualitative behavior of a flow function satisfying Assumption 2: $\mu_e(\rho_e)$ is differentiable, strictly increasing, has bounded derivative and such that $\mu_e(0) = 0$, and $\lim_{\rho_e \to \infty} \mu_e(\rho_e) = f_e^{\max} < \infty$. The median density ρ_e^{μ}, as defined in (5) is plotted as well.

by $f = \mu(\rho)$ and $f^v = \mu^v(\rho^v)$ the functional relationships between the corresponding 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., [27]). We shall assume to have fixed one such ordering, identifying $\mathcal{V}$ with the integer set $\{0, 1, \ldots, n\}$, where $n := |\mathcal{V}| - 1$, in such a way that

$$\mathcal{E}_v^- \subseteq \bigcup_{0 \le u < v} \mathcal{E}_u^+ , \qquad \forall v = 0, \ldots, n . \tag{2}$$

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

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

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. 2). The *min-cut capacity* of a network $\mathcal{N}$ is defined as

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

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

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

Then, it follows from the max-flow min-cut theorem (see, e.g., [2]), 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 flow at the intermediate nodes.

Throughout the paper, we shall make the following assumption on the flow functions (see also Fig. 3):

Assumption 2: For every link $e \in \mathcal{E}$, the map $\mu_e : \mathbb{R}_+ \to \mathbb{R}_+$ 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

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

Example 1 (Flow function): For every link $e \in \mathcal{E}$, let a_e and $f_e^{\max}$ be positive real constants. Then, a simple example of flow function satisfying Assumption 2 is given by

$$\mu_e(\rho_e) = f_e^{\max} \left(1 - \exp(-a_e \rho_e)\right) .$$

It is easily verified that the flow capacity is $f_e^{\max}$, while the median density for such a flow function is $\rho_e^{\mu} = a_e^{-1} \log 2$.

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

Definition 2 ((Distributed) routing policy): A *routing policy* for a network $\mathcal{N}$ is a family of differentiable functions $\mathcal{G} := \{G^v : \mathcal{R} \to \mathcal{S}_v\}_{0 \le v < n}$ 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 $\overline{G}^v : \mathcal{R}_v \to \mathcal{S}_v$ such that $G^v(\rho) = \overline{G}^v(\rho^v)$ 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 current 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 $\overline{G}^v(\rho^v)$, for the vector of the fractions in which the total outflow of node v gets split into its outgoing links.

We are now ready to define a dynamical network.

Definition 3 (Dynamical network): A *dynamical network* associated to a network $\mathcal{N}$ satisfying Assumption 1, a distributed routing policy $\mathcal{G}$, and a total outflow $\lambda_0 \geq 0$ at the origin node, is the dynamical system

$$\frac{\mathrm{d}}{\mathrm{d}t}\rho_e(t) = \lambda_{\sigma(e)}(t) G_e^{\sigma(e)}(\rho(t)) - f_e(t)\,, \qquad \forall\, e \in \mathcal{E}\,, \quad (6)$$

where

$$f_e(t) := \mu_e(\rho_e(t))\,,$$

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

Equation (6) states that the rate of change of the particle density on a link e is given by the difference between $\lambda_{\sigma(e)}(t) G_e^{\sigma(e)}(\rho(t))$, i.e., the portion of the total outflow at the tail node $\sigma(e)$ of link e which is routed to link e, and $f_e(t)$, i.e., the particle flow on link e. Observe that the (distributed) routing policy $G^v(\rho)$ at node v induces a (local) feedback which couples the dynamics of the particle flow on the different links outgoing from node v. Notice that the differentiability assumptions on the routing policy and on the flow functions readily imply existence and uniqueness of a solution to the dynamical network (6) for every initial flow $f^\circ \in \mathcal{F}$ (or, equivalently, for every initial density $\rho^\circ \in \mathcal{R}$).

We now introduce the following notion of transfer efficiency of a dynamical network.

Definition 4 (Transfer efficiency): Consider a dynamical network $\mathcal{N}$ satisfying Assumptions 1 and 2. Given some flow vector $f^\circ \in \mathcal{F}$, and $\alpha \in [0, 1]$, the dynamical 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

$$\liminf_{t \to \infty} \lambda_n(t) \geq \alpha \lambda_0\,. \quad (7)$$

Definition 4 states that a dynamical network is α-transferring when the total inflow at destination is asymptotically not smaller than α times the total outflow at the origin. In particular, a fully transferring ($\alpha = 1$) dynamical network is characterized by the property of having total inflow at destination asymptotically equal to the total outflow at the origin, so that there is no throughput loss. On the other hand, a partially transferring dynamical network might allow for some throughput loss, provided that some fraction of the flow is still guaranteed to be asymptotically transferred.

Remark 1: Standard definitions in the literature are typically limited to static network flows describing transport of particles at equilibrium via conservation of flow at intermediate nodes. In fact, they usually consist (see e.g., [2]) 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, $f^* \in \mathrm{cl}(\mathcal{F}^*(\lambda_0))$ for $\lambda_0 \leq C(\mathcal{N})$). In contrast, in our model, we focus on the off-equilibrium particle dynamics on a network $\mathcal{N}$, induced by a (distributed) routing policy $\mathcal{G}$.

B. Examples

We now present three illustrative applications of the dynamical network framework.

(i) *Data networks*: We start by explaining how to frame data networks with TCP/IP congestion control in our setting. We shall refer to models and terminology from [5], [7]. In data networks, the particles are meant to represent data packets, the nodes are an abstraction for the combination of a modem and the corresponding local process associated with the data link control layer (for transmission of data between nodes) and the network layer (for implementing the routing protocol). Links are the channels where the packets form queues during transmission between the corresponding nodes. Hence, the notion of *density* ρ_e on a link e in our framework is directly related to the (suitably rescaled) queue length on the channels. On the other hand, the outflow $\mu_e(\rho_e) = f_e$ represents the throughput, measuring the number of packets successfully transmitted per unit of time.

Observe that, in TCP/IP, the round-trip (delay) time (RTT) is the time from sending a packet for the first time to receiving its acknowledgement from the destination (see, e.g., [7]). When only a few packets are being sent on the channel, one does not observe relevant congestion effects, and the throughput can be reasonably modeled as proportional to the density divided by the RTT. On the other hand, as the number of packets being sent increases, bandwidth limitations imply an increase in the packet drop probability, and consequently an increase in the (average) number of retransmissions per packet. As an effect, as the packet density increases, the throughput tends to saturate approaching the channel capacity. Flow functions $f_e = \mu_e(\rho_e)$ satisfying Assumption 2 provide effective models for such behavior.

(ii) *Production networks*: In the context of multi-stage production networks, the particles represent goods that need to be processed by a series of production modules located on links. The nodes represent abstractions of routing policies that route goods from one stage to the next. The density corresponds to work-in-process. It is known, e.g., see [8], that the throughput of a single-stage production system follows a nonlinear relationship with respect to its work-in-process. This relationship, which is commonly referred to as *clearing*

function, satisfies Assumption 2 of our model. Therefore, such production networks have a clear analogy with our setup where ρ_e represents the work-in-progress and $\mu_e(\rho_e)$ represents the clearing function. Notice that, in our formulation, we assume infinite capacity for work-in-process because of which our model can not generate *bullwhip* effects commonly observed in supply chains, e.g., see [28], [9].

(iii) *Traffic networks*: In road traffic networks, particles represent cars and distributed routing policies correspond to the local route choice behavior of the drivers in response to the locally observed link densities. The distributed routing policies correspond to the node-wise route choice behavior of the drivers. In that respect, observe that, in road traffic networks, locally responsive policies as characterized by Definition 7 are particularly natural as they model the behavior of drivers myopically preferring routes which appear to be locally less congested.

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 [4]. As already pointed out, however, Assumption 2 on the monotonicity of the flow functions poses a potential limitation to the applicability of our results, as typical fundamental diagrams in road traffic theory have a $\cap$-shape form. Some simulations are provided in [3] illustrating how the results of this paper could be extended to this case. On the other hand, the analysis presented in this paper continues to provide insight on the local behavior of dynamical networks with flow functions having $\cap$-shaped graph, in the region where the flow is increasing in the density.[1]

Remark 2: It is worth stressing that, while distributed routing policies depend only on local information on the current density, their structural form may depend on some global information on the network. Such global information 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 [29] in the context of traffic networks. Multiple time-scale dynamical processes have also been analyzed in [30] in the context of data networks. When not directly designable, desired route choice behaviors from a social optimization perspective may be achieved by appropriate incentive mechanisms. While we do not address the issue of mechanism design in this paper, the companion paper [3] discusses the use of tolls in influencing the long-term global route choice behavior of drivers to get a desired equilibrium flow in traffic networks.

C. Perturbed dynamical networks and resilience

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

Definition 5 (Admissible perturbation): An *admissible perturbation* of a network $\mathcal{N} = (\mathcal{T}, \mu)$, satisfying Assumptions 1 and 2, is a network $\tilde{\mathcal{N}} = (\mathcal{T}, \tilde{\mu})$, with the same topology $\mathcal{T}$,

[1] In particular, in that context, possibly conservative bounds on the resilience of dynamical networks driven by locally responsive distributed routing policies can be obtained by using the analysis presented in this paper and the monotone properties of such dynamical systems.

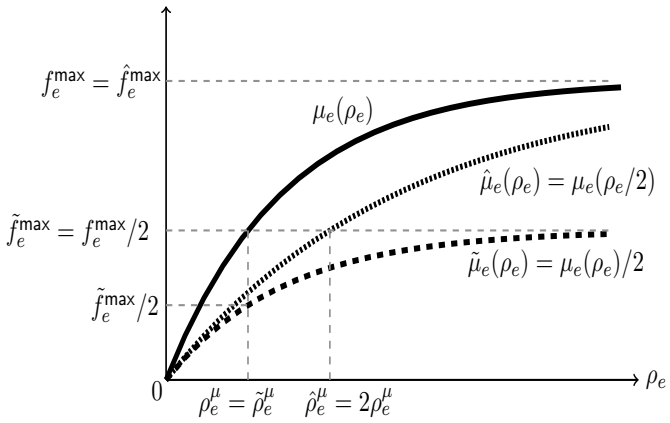


Fig. 4. Perturbed flow functions for Example 2, with $\varepsilon = 1/2$. For the first perturbation $\tilde{\mu}_e(\rho_e) = \mu_e(\rho_e)/2$ (dashed plot), the magnitude is $\delta_e = f_e^{\max}/2$, while the stretching coefficient is $\theta = \tilde{\rho}_e^\mu/\rho_e^\mu = 1$. For the second perturbation $\hat{\mu}_e(\rho_e) = \mu_e(\rho_e/2)$ (densely dotted plot), the magnitude is $\delta_e = \sup\{\mu_e(\rho_e) - \mu_e(\rho_e/2) : \rho_e \geq 0\}$, while the stretching coefficient is $\hat{\theta} = \hat{\rho}_e^\mu/\rho_e^\mu = 1/2$.

and a family of perturbed flow functions $\tilde{\mu} := \{\tilde{\mu}_e : \mathbb{R}_+ \to \mathbb{R}_+\}_{e\in\mathcal{E}}$, such that, for every $e \in \mathcal{E}$, $\tilde{\mu}_e$ satisfies Assumption 2, as well as

$$\tilde{\mu}_e(\rho_e) \leq \mu_e(\rho_e)\,, \qquad \forall \rho_e \geq 0\,.$$

We accordingly let $\tilde{f}_e^{\max} := \sup\{\tilde{\mu}_e(\rho_e) : \rho_e \geq 0\}$ and $\tilde{\mathcal{F}} := \times_{e\in\mathcal{E}}[0, \tilde{f}_e^{\max})$. The *magnitude* of an admissible perturbation is defined as

$$\delta := \sum_{e\in\mathcal{E}} \delta_e\,, \qquad \delta_e := \sup_{\rho_e \geq 0} \{\mu_e(\rho_e) - \tilde{\mu}_e(\rho_e)\}\,. \tag{8}$$

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

$$\theta := \max\{\tilde{\rho}_e^\mu/\rho_e^\mu : e \in \mathcal{E}\}\,, \tag{9}$$

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).

Observe that the magnitude of an admissible perturbation is defined as the sum, over all links, of the infinity norm of the unperturbed minus the perturbed flow functions and is therefore an aggregate measure of the changes on the ordinate of the flow function graphs. In contrast, the stretching coefficient is a measure of the maximal change of the median of the flow functions, which is measured on the abscissa of their graphs. In fact, we shall regard the former as the most informative measure of the perturbation, while the latter is introduced mostly for technical reasons which will be made clear in the sequel (see Remark 4).

Example 2: Fix $\varepsilon \in (0, 1]$, and consider the perturbed networks with flow functions $\tilde{\mu}_e(\rho_e) = \varepsilon\mu_e(\rho_e)$, and $\hat{\mu}_e(\rho_e) = \mu_e(\varepsilon\rho_e)$, respectively, for $e \in \mathcal{E}$. Then, the first perturbation has magnitude $\delta = (1-\varepsilon)\sum_{e\in\mathcal{E}} f_e^{\max}$, and stretching coefficient $\theta = 1$, while the second one has magnitude $\hat{\delta} = \sum_{e\in\mathcal{E}} \sup\{\mu_e(\rho_e) - \mu_e(\varepsilon\rho_e) : \rho_e \geq 0\}$, and stretching coefficient $\hat{\theta} = 1/\varepsilon$. The case $\varepsilon = 1/2$ is reported in Figure 2.

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

$$\frac{\mathrm{d}}{\mathrm{d}t}\tilde{\rho}_e(t) = \tilde{\lambda}_{\sigma(e)}(t)G_e^{\sigma(e)}(\tilde{\rho}(t)) - \tilde{f}_e(t)\,, \qquad \forall e \in \mathcal{E}\,, \quad (10)$$

where

$$\begin{aligned} &\tilde{f}_e(t) := \tilde{\mu}_e(\tilde{\rho}_e(t))\,, \\ &\tilde{\lambda}_v(t) := \begin{cases} \sum_{e\in\mathcal{E}_v^-} \tilde{f}_e(t) & \text{if} \quad 0 < v < n \\ \lambda_0 & \text{if} \quad v = 0\,. \end{cases} \end{aligned} \quad (11)$$

Observe that the perturbed dynamical network (10) has the same structure as the original dynamical network (6), as it describes the rate of change of the particle density on each link e as the difference between $\tilde{\lambda}_{\sigma(e)}(t)G_e^{\sigma(e)}(\tilde{\rho}(t))$, i.e., the portion of the total outflow at the tail node $\sigma(e)$ of link e which is routed to link e, and the perturbed flow on link e itself. Notice that the only difference with respect to the original dynamical 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 networks. In this way, we model a situation in which the routers are not aware of the fact that the network has been perturbed, but react to this change only indirectly, in response to variations of the local density vectors $\tilde{\rho}^v(t)$.

Remark 3: Observe that admissible perturbations as characterized by Definition 5 do not include complete link shutdowns, as these would correspond to perturbed flow functions $\tilde{\mu}_e(\rho_e) \equiv 0$, which clearly violate the strict monotonicity required by Assumption 2. Nevertheless, complete link shutdowns can be approximated arbitrarily closely by admissible perturbations, e.g., by considering perturbed flow functions of the form $\tilde{\mu}_e(\rho_e) = \varepsilon\mu_e(\rho_e)$ with arbitrarily small but positive ε. In fact, the analysis presented in this work could be suitably extended so as to include complete link shutdowns. The authors' choice not to do that is in the interest of simplicity and brevity of the exposition.

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

Definition 6: (Resilience of a dynamical network) Let $\mathcal{N}$ be a network satisfying Assumptions 1 and 2, $\mathcal{G}$ be a distributed routing policy, and $\lambda_0 \geq 0$ be a constant total outflow 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 network (10) 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[2]

$$\gamma_{\alpha}(f^\circ,\mathcal{G}) := \lim_{\theta\to\infty} \gamma_{\alpha,\theta}(f^\circ,\mathcal{G})\,.$$

The 1-resilience will be referred to as the *strong resilience*, while the 0-resilience will be referred to as the *weak resilience*.

Remark 4: For $\alpha = 0$, the perturbed network 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 $\lim_{\theta\to\infty}\lim_{\alpha\downarrow 0}\gamma_{\alpha,\theta}(f^\circ,\mathcal{G})$: the introduction of the bound on the stretching coefficient of the admissible perturbation is a mere technicality whose necessity will become clear in Section V.

In the remainder of the paper, we shall focus on the characterization of the weak resilience of dynamical networks, while the strong resilience will be addressed in the companion paper [3]. 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 (10) 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 network from (6) to (10), and make it not α-transferring, provided that such an effort is measured in terms of the magnitude of the perturbation $\delta = \sum_{e\in\mathcal{E}} ||\mu_e(\,\cdot\,) - \tilde{\mu}_e(\,\cdot\,)||_\infty$.

Remark 5 (Zero-sum game interpretation): 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., $\sum_{e\in\mathcal{E}} f_{\mathrm{e}}^{\max}$, 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 optimal in terms of both weak and strong resilience. This will then show that the locally responsive distributed routing policies correspond to approximate Nash equilibria in this zero-sum game setting.

We conclude this section with the following result, providing an upper bound on the weak resilience of a dynamical 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 network has weak resilience equal to the min-cut capacity.

Proposition 1: Let $\mathcal{N}$ be a network satisfying Assumptions 1 and 2, $\lambda_0 > 0$ a constant total outflow at the origin node,

[2]It 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 network to be not α-transferring).

and $\mathcal{G}$ an arbitrary routing policy. Then, for any initial flow f°, the weak resilience of the associated dynamical network satisfies

$$\gamma_0(f^\circ, \mathcal{G}) \leq C(\mathcal{N})\,.$$

Proof: We shall prove that, for every $\alpha \in (0,1]$, and every $\theta \geq 1$,

$$\gamma_{\alpha,\theta}(f^\circ, \mathcal{G}) \leq C(\mathcal{N}) - \frac{\alpha}{2}\lambda_0\,. \tag{12}$$

Observe that (12) immediately implies that

$$\begin{aligned}\gamma_0(f^\circ, \mathcal{G}) &= \lim_{\theta\to\infty}\lim_{\alpha\downarrow 0}\gamma_{\alpha,\theta}(f^\circ, \mathcal{G})\\ &\leq \lim_{\theta\to\infty}\lim_{\alpha\downarrow 0}(C(\mathcal{N}) - \alpha\lambda_0/2)\\ &= C(\mathcal{N})\,,\end{aligned}$$

thus proving the claim.

Consider a minimal origin-destination cut, i.e., some $\mathcal{U} \subseteq \mathcal{V}$ such that $0 \in \mathcal{U}$, $n \notin \mathcal{U}$, and $\sum_{e\in\mathcal{E}^+_{\mathcal{U}}} f^{\max}_e = C(\mathcal{N})$. Define $\varepsilon := \alpha\lambda_0/(2C(\mathcal{N}))$, and consider an admissible perturbation such that $\tilde{\mu}_e(\rho_e) = \varepsilon\mu_e(\rho_e)$ for every $e \in \mathcal{E}^+_{\mathcal{U}}$, and $\tilde{\mu}_e(\rho_e) = \mu_e(\rho_e)$ for all $e \in \mathcal{E}\setminus\mathcal{E}^+_{\mathcal{U}}$. It is readily verified that the magnitude of such perturbation satisfies

$$\delta = (1-\varepsilon)\sum_{e\in\mathcal{E}^+_{\mathcal{U}}} f^{\max}_e = (1-\varepsilon)C(\mathcal{N}) = C(\mathcal{N}) - \frac{\alpha}{2}\lambda_0\,,$$

while its stretching coefficient is 1.

Observe that

$$\tilde{\lambda}_{\mathcal{U}}(t) := \sum_{e\in\mathcal{E}^+_{\mathcal{U}}}\tilde{f}_e(t) \leq \sum_{e\in\mathcal{E}^+_{\mathcal{U}}}\tilde{f}^{\max}_e = \varepsilon\sum_{e\in\mathcal{E}^+_{\mathcal{U}}} f^{\max}_e = \frac{\alpha\lambda_0}{2}\,, \tag{13}$$

for all $t \geq 0$. Now, let $\mathcal{W} := \mathcal{V}\setminus\mathcal{U}$ be the set of nodes on the destination side of the cut, and observe that, for every $w \in \mathcal{W}\setminus\{n\}$,

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\sum_e\tilde{\rho}_e(t) &= \sum_e\sum_j\tilde{f}_j(t)G^w_e(\tilde{\rho}(t)) - \sum_e\tilde{f}_e(t)\\ &= \sum_j\tilde{f}_j(t) - \sum_e\tilde{f}_e(t)\,,\end{aligned} \tag{14}$$

where the summation indices e and j run over $\mathcal{E}^+_w$, and $\mathcal{E}^-_w$, respectively. Define

$$\mathcal{A} := \bigcup_{w\in\mathcal{W}}\mathcal{E}^+_w\,, \qquad \mathcal{B} := \bigcup_{w\in\mathcal{W}}\mathcal{E}^-_w\,,$$

and let

$$\zeta(t) := \sum_{e\in\mathcal{A}}\tilde{\rho}_e(t)\,.$$

From (14), the identity $\mathcal{A}\cup\mathcal{E}^+_{\mathcal{U}} = \mathcal{B}$, and (13), one gets

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\zeta(t) &= \sum_{w\in\mathcal{W}}\sum_{e\in\mathcal{E}^+_w}\frac{\mathrm{d}}{\mathrm{d}t}\tilde{\rho}_e(t)\\ &= \sum_{e\in\mathcal{B}}\tilde{f}_e(t) - \sum_{e\in\mathcal{E}^-_n}\tilde{f}_e(t) - \sum_{e\in\mathcal{A}}\tilde{f}_e(t)\\ &= \sum_{e\in\mathcal{E}^+_{\mathcal{U}}}\tilde{f}_e(t) - \sum_{e\in\mathcal{E}^-_n}\tilde{f}_e(t)\\ &\leq \alpha\lambda_0/2 - \tilde{\lambda}_n(t)\,.\end{aligned} \tag{15}$$

Now assume, by contradiction, that

$$\liminf_{t\to\infty}\tilde{\lambda}_n(t) \geq \alpha\lambda_0\,.$$

Then, there would exist some $s \geq 0$ such that $\tilde{\lambda}_n(t) \geq 3\alpha\lambda_0/4$ for all $t \geq s$. For all $t \geq s$, it would then follow from (15) that $\mathrm{d}\zeta(t)/\mathrm{d}t \leq -\alpha\lambda_0/4 < 0$, which would contradict the fact that $\zeta(t) \geq 0$ for all $t \geq 0$. Then, necessarily

$$\liminf_{t\to\infty}\tilde{\lambda}_n(t) < \alpha\lambda_0\,,$$

so that the perturbed dynamical network is not α-transferring. This implies (12), and therefore the claim. ∎

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

(a) for every $\rho^v \in \mathcal{R}_v$,

$$\frac{\partial}{\partial\rho_e}G^v_j(\rho^v) \geq 0\,, \qquad \forall j, e \in \mathcal{E}^+_v\,,\ j \neq e\,;$$

(b) for every nonempty proper subset $\mathcal{J} \subsetneq \mathcal{E}^+_v$, there exists a continuously differentiable map

$$G^{\mathcal{J}} : \mathcal{R}_{\mathcal{J}} \to \mathcal{S}_{\mathcal{J}}\,,$$

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

$$\rho^v_e \to \infty\,,\ \ \forall e \in \mathcal{E}^+_v\setminus\mathcal{J}\,, \qquad \rho^v_j \to \rho^{\mathcal{J}}_j\,,\ \ \forall j \in \mathcal{J}\,,$$

then

$$\begin{aligned}&G^v_e(\rho^v) \to 0\,, && \forall e \in \mathcal{E}^+_v\setminus\mathcal{J}\,,\\ &G^v_j(\rho^v) \to G^{\mathcal{J}}_j(\rho^{\mathcal{J}})\,, && \forall j \in \mathcal{J}\,.\end{aligned}$$

Property (a) in Definition 7 states that, as the particle density on an outgoing link $e \in \mathcal{E}^+_v$ increases while the particle density on all the other outgoing links remains constant, the fraction of total outflow at node v routed to any link $j \in \mathcal{E}^+_v\setminus\{e\}$ does not decrease, and hence the fraction of total outflow 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* [11], [12]. 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. It is worth observing that, when the routing policy models some selfish behavior of the particles (e.g., in road traffic networks), then Property (a) and (b) are very natural assumptions on such behavior as they capture some sort of greedy local behavior.

Example 3 (Locally responsive distributed routing): 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

$$G_e^v(\rho) = \frac{a_e \exp(-\eta_v \rho_e)}{\sum_{i \in \mathcal{E}_v^+} a_i \exp(-\eta_v \rho_i)}, \tag{16}$$

for every $0 \leq v < n$ and $e \in \mathcal{E}_v^+$. 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 total outflow at node v gets split among its outgoing links. Moreover, for all $0 \leq v < n$, and $e \in \mathcal{E}_v^+$, $G_e^v(\rho)$ is clearly differentiable, and computing partial derivatives one gets

$$\frac{\partial}{\partial \rho_j} G_e^v(\rho) = \eta_v \frac{a_e a_j \exp(-\eta_v \rho_e) \exp(-\eta_v \rho_j)}{\left(\sum_{i \in \mathcal{E}_v^+} a_i \exp(-\eta_v \rho_i)\right)^2} \geq 0, \tag{17}$$

for every $j \in \mathcal{E}_v^+$, $j \neq e$, and $\frac{\partial}{\partial \rho_j} G_e(\rho) = 0$ for all $j \in \mathcal{E} \setminus \mathcal{E}_v^+$. This implies that Property (a) of Definition 7 holds true. Property (b) is also easily verified. Therefore, $\mathcal{G}$ is a locally responsive distributed routing policy. In the context of road traffic networks, the example in (16) is a variant of the logit function from discrete choice theory emerging from utility 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., [31]).

We are now ready to state our main results. The first one shows that, when the distributed routing policy $\mathcal{G}$ is locally responsive, the dynamical network (6) always admits a unique, globally attractive limit flow vector.

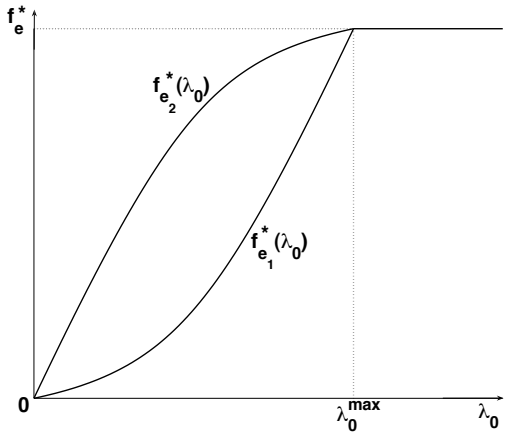
Theorem 1 (Existence of globally attractive limit flow): Let $\mathcal{N}$ be a network satisfying Assumptions 1 and 2, $\lambda_0 \geq 0$ a constant total outflow at the origin node, and $\mathcal{G}$ a locally responsive distributed routing policy. Then, there exists a unique limit flow $f^* \in \mathrm{cl}(\mathcal{F})$ such that, for every initial condition $\rho(0) \in \mathcal{R}$, the dynamical network (6) satisfies

$$\lim_{t \to \infty} f(t) = f^*.$$

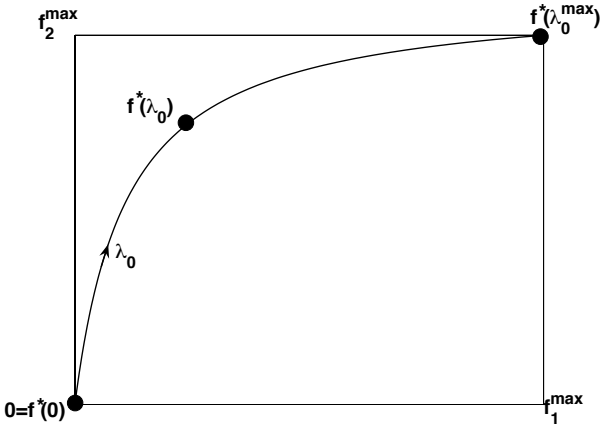
Moreover, the limit flow f^* is such that, if $f_e^* = f_e^{\max}$ for some link $e \in \mathcal{E}_v^+$ outgoing from a nondestination node $0 \leq v < n$, then $f_e^* = f_e^{\max}$ for every outgoing link $e \in \mathcal{E}_v^+$ on that node.

Proof: See Section IV. ∎

Theorem 1 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., $f^* \in \mathcal{F}^*(\lambda_0)$; or it belongs to $\mathrm{cl}(\mathcal{F}) \setminus \mathcal{F}$, i.e., it satisfies the capacity constraint on at least one link with equality, in which case it is not an equilibrium flow. In the latter case, it satisfies the additional property that, on all the links outgoing from the same node, the capacity constraints are satisfied with equality. Such additional property will prove particularly useful in our companion paper [3], when characterizing the strong resilience of dynamical networks. As it will become clear in Section IV, the global convergence result mainly relies on Assumption 2 on monotonicity of the flow function, and Property (a) in Definition 7 of locally responsive distributed routing policies, from which the dynamical network (6) inherits a cooperative property. It is worth mentioning that we shall not use general results for cooperative dynamical systems [11], [12], [32], but rather exploit some other structural properties of (6) which in fact allow us to prove stronger results. The additional property of the limit flow follows instead mainly from Property (b) of Definition 7.



(a)

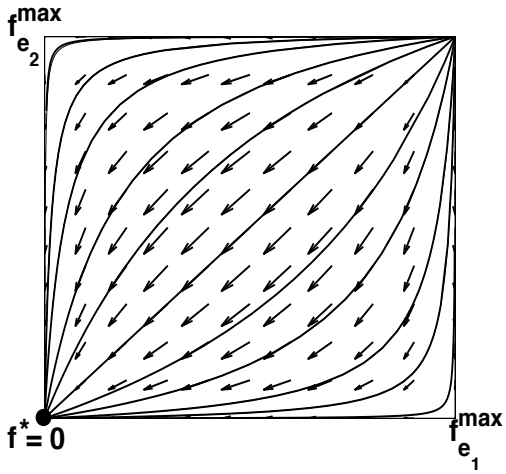


(b)

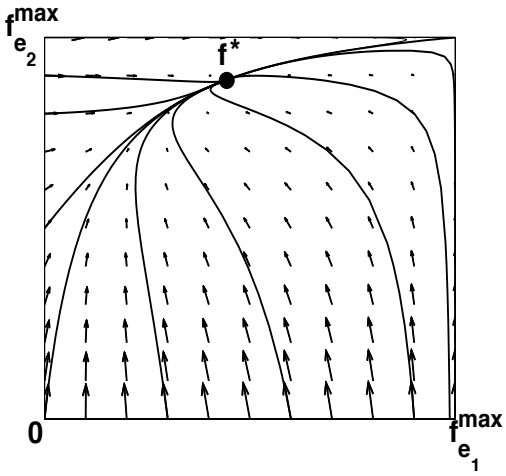
Fig. 5. Dependence of the limit flow f^* on the total outflow λ_0 at the origin for the dynamical network of Example 4. In (a), the two components of the limit flow, $f_{e_1}^*$ and $f_{e_2}^*$, are plotted as functions of λ_0. In (b), the curve of the limit flows is plotted in the $(f_{e_1}^*, f_{e_2}^*)$-plane. Observe that both the components of the limit flow are increasing from 0 to $f_e^{\max}$, as λ_0 is increasing from 0 to $\lambda_0^{\max}$, while they remain constant at $f_e^{\max}$, for all $\lambda_0 \geq \lambda_0^{\max}$.

Example 4: Consider a simple topology containing just the origin and the destination node, i.e., with $\mathcal{V} = \{0, 1\}$, and two parallel links $\mathcal{E} = \{e_1, e_2\}$. Assume that the flow functions on the two links are identical $\mu_{e_1}(\rho) = \mu_{e_2}(\rho) = 3(1 - e^{-\rho})/4$. Consider the routing policy

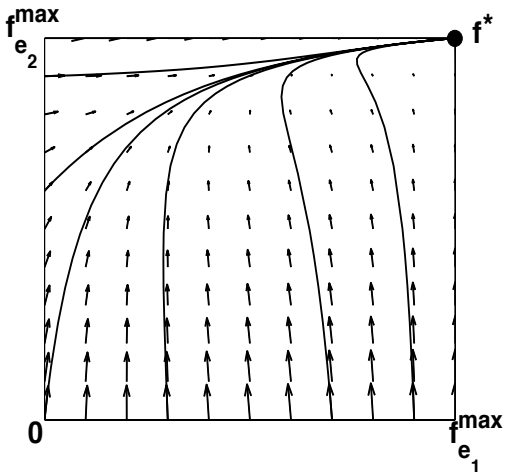
$$G_{e_1}^0(\rho) = \frac{3}{5} \frac{e^{-\rho_{e_1}}}{Z(\rho)}, \qquad G_{e_2}^0(\rho) = 6 \frac{e^{-\rho_{e_2}}}{Z(\rho)},$$



(a) $\lambda_0 = 0$



(b) $\lambda_0 = 1$



(c) $\lambda_0 = 2$

Fig. 6. Flow vector fields and flow trajectories for the dynamical network of Example 4, for three values of the inflow. In the first two cases $\lambda_0 < \lambda_0^{\max}$, and hence the limit flow f^* is an equilibrium flow. In contrast, in the last case, $\lambda_0 \geq \lambda_0^{\max}$, and consequently f^* is not an equilibrium flow and $f^*_{e_1} = f^{\max}_{e_1}$ and $f^{\max}_{e_2} = f^*_{e_2}$, as predicted by Theorem 1.

where $Z(\rho) := \frac{3}{5}e^{-\rho_{e_1}} + 6e^{-\rho_{e_2}}$. Then, the limit flow of the associated dynamical network can be explicitly computed for every constant total outflow $\lambda_0 \geq 0$ at the origin node, and is given by

$$f_1^*(\lambda_0) = \begin{cases} (12\lambda_0 - 11 + \omega(\lambda_0))/24 & \text{if} \quad 0 \leq \lambda_0 < 3/2 \\ 3/4 & \text{if} \quad \lambda_0 \geq 3/2\,, \end{cases}$$

$$f_2^*(\lambda_0) = \begin{cases} (12\lambda_0 + 11 - \omega(\lambda_0))/24 & \text{if} \quad 0 \leq \lambda_0 < 3/2 \\ 3/4 & \text{if} \quad \lambda_0 \geq 3/2\,, \end{cases}$$

where

$$\omega(\lambda_0) := \sqrt{(12\lambda_0 - 11)^2 + 48\lambda_0}\,.$$

Figure 5 shows the dependence of the limit flow f^* on the total outflow at the origin, λ_0. The two components $f^*_{e_1}$, and $f^*_{e_2}$, increase from 0 to $f^{\max}_{e_1}$, and, respectively, from 0 to $f^{\max}_{e_2}$, as λ_0 increases from 0 to $\lambda_0^{\max} := f^{\max}_{e_1} + f^{\max}_{e_2}$, and they remain constant for all $\lambda_0 \geq \lambda_0^{\max}$. Figure 6 reports the vector fields and flow trajectories associated to the dynamical network for three different values of the inflow, namely $\lambda_0 = 0$, $\lambda_0 = 1$, and $\lambda_0 = 2$. In the first two cases, $\lambda_0 < \lambda_0^{\max}$, and $f^* \in \mathcal{F}^*(\lambda_0)$ is an equilibrium flow; in the third case, $f^* \in \mathrm{cl}(\mathcal{F}(\lambda_0)) \setminus \mathcal{F}(\lambda_0)$ is not an equilibrium flow.

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

Theorem 2 (Weak resilience): Let $\mathcal{N}$ be a network satisfying Assumptions 1 and 2, $\lambda_0 > 0$ a constant inflow at the origin node, and $\mathcal{G}$ a locally responsive distributed routing policy such that $G^v_e(\rho^v) > 0$ 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 network is partially transferring with respect to f° and has weak resilience

$$\gamma_0(f^\circ, \mathcal{G}) = C(\mathcal{N})\,.$$

Proof: See Section V. ■

Theorem 2, combined with Proposition 1, shows that locally responsive distributed routing policies achieve the maximal weak resilience possible on a given network $\mathcal{N}$.

Remark 6: (i) A consequence of Theorem 2 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 networks presented in the companion paper [3]. There, it is shown that the strong resilience depends on the initial flow, and local information constraints reduce the maximal strong resilience achievable on a given network.

(ii) It is also interesting to note that the upper bound on the weak resilience, as given by Proposition 1, does not change even if we allow routing policies $\mathcal{G}$ that have knowledge of the perturbation. This, combined with Theorem 2, shows that the lack of knowledge of the

perturbations is not a hindrance in achieving maximal weak resilience.

IV. Proof of Theorem 1

Let $\mathcal{N}$ be a network satisfying Assumptions 1 and 2, $\mathcal{G}$ a locally responsive distributed routing policy, and $\lambda_0 \geq 0$ a constant outflow at the origin node. We shall prove that there exists a unique $f^* \in \mathrm{cl}(\mathcal{F})$ such that the flow $f(t)$ associated to the solution of the dynamical network (6) converges to f^* as t grows large, for every initial condition $\rho(0) \in \mathcal{R}$. Before proceeding, it is worth observing that, thanks to Property (a) of Definition 7 of locally responsive distributed routing policies, Assumption 2 on the monotonicity of the flow functions, and the structure of the dynamical network (6), one may rewrite (6) as

$$\frac{\mathrm{d}}{\mathrm{d}t}\rho_e = F_e(\rho)\,, \qquad \forall e \in \mathcal{E}\,,$$

where $F : \mathcal{R} \to \mathbb{R}^{\mathcal{E}}$ is differentiable and such that

$$\frac{\partial}{\partial \rho_e} F_e(\rho) \leq 0\,, \qquad \frac{\partial}{\partial \rho_j} F_e(\rho) \geq 0\,, \qquad \forall e \neq j \in \mathcal{E}\,.$$

The above shows that, the dynamical network (6) driven by a locally responsive distributed routing policy $\mathcal{G}$ is cooperative in the sense of Hirsch [11], [12]. Indeed, one may apply the standard theory of cooperative dynamical systems and monotone flows [32], [13] in order to prove some properties of (6), e.g., convergence from almost every initial condition.

However, we shall not rely on this general theory and rather use a direct approach leading us to much stronger results, i.e., *global* convergence to a *unique* limit flow. We shall proceed by proving a series of intermediate results some of which will prove useful also in the companion paper [3]. Our approach is based on the observation that, thanks to Assumption 1 on the acyclicity of the network topology, one can consider the dynamical network (6) as a cascade of monotone local systems (see [14]), each describing the flow dynamics on the set of outgoing links of a non-destination node. Specifically, for every $0 \leq v < n$, we shall focus on the input-output properties of the *local system*

$$\begin{aligned} &\frac{\mathrm{d}}{\mathrm{d}t}\rho_e(t) = \lambda(t) G_e^v(\rho^v(t)) - f_e(t)\,, \\ &f_e(t) = \mu_e(\rho_e(t))\,, \end{aligned} \qquad \forall e \in \mathcal{E}_v^+\,, \quad (18)$$

where $\lambda(t)$ is a nonnegative-real-valued, Lipschitz continuous input, and $f^v(t) := \{f_e(t) : e \in \mathcal{E}_v^+\}$ is interpreted as the output. We shall first prove existence and uniqueness of a globally attractive limit flow for the local system (18) under constant input (a property similar to *static input-output characteristic*, cf. [14, Def. V.I]). We shall then extend this result to show the existence and attractivity of a local equilibrium point under time-varying, convergent local input. Finally, we shall exploit this local input-output property, and the assumption of acyclicity of the network topology in order to establish the main result.

The following is a simple technical result, which will prove useful in order to apply Property (a) of Definition 7.

Lemma 1: Let $0 \leq v < n$ be a nondestination node, and $G^v : \mathcal{R}_v \to \mathcal{S}_v$ a continuously differentiable function satisfying Property (a) of Definition 7. Then, for any $\nu, \varsigma \in \mathcal{R}_v$,

$$\sum_{e \in \mathcal{E}_v^+} \mathrm{sgn}(\nu_e - \varsigma_e)\,(G_e^v(\nu) - G_e^v(\varsigma)) \;\leq\; 0. \qquad (19)$$

Proof: Define

$$\begin{aligned} &\mathcal{K} := \{e \in \mathcal{E}_v^+ : \nu_e > \varsigma_e\}\,, && G_{\mathcal{K}}(\zeta) := \textstyle\sum_k G_k^v(\zeta)\,, \\ &\mathcal{J} := \{e \in \mathcal{E}_v^+ : \nu_e \leq \varsigma_e\}\,, && G_{\mathcal{J}}(\zeta) := \textstyle\sum_j G_j^v(\zeta)\,, \\ &\mathcal{L} := \{e \in \mathcal{E}_v^+ : \nu_e < \varsigma_e\}\,, && G_{\mathcal{L}}(\zeta) := \textstyle\sum_l G_l^v(\zeta)\,, \end{aligned}$$

where $\zeta \in \mathcal{R}_v$, and the summation indices k, l, and j run over $\mathcal{K}$, $\mathcal{L}$, and $\mathcal{J}$, respectively. We shall show that, for any $\nu, \varsigma \in \mathcal{R}_v$,

$$G_{\mathcal{K}}(\nu) \leq G_{\mathcal{K}}(\varsigma), \qquad G_{\mathcal{L}}(\nu) \geq G_{\mathcal{L}}(\varsigma)\,. \qquad (20)$$

Let $\xi \in \mathcal{R}_v$ be defined by $\xi_k = \nu_k$ for all $k \in \mathcal{K}$, and $\xi_e = \varsigma_e$ for all $e \in \mathcal{E}_v^+ \setminus \mathcal{K}$. We shall prove that $G_{\mathcal{K}}(\nu) - G_{\mathcal{K}}(\varsigma) \leq 0$ by writing it as a path integral of $\nabla G_{\mathcal{K}}(\zeta)$ first along the segment $S_{\mathcal{K}}$ from ς to ξ, and then along the segment $S_{\mathcal{L}}$ from ξ to ν. Proceeding in this way, one gets

$$\begin{aligned} G_{\mathcal{K}}(\nu) - G_{\mathcal{K}}(\varsigma) &= \int_{S_{\mathcal{L}}} \nabla G_{\mathcal{K}}(\zeta) \cdot \mathrm{d}\zeta + \int_{S_{\mathcal{K}}} \nabla G_{\mathcal{K}}(\zeta) \cdot \mathrm{d}\zeta \\ &= \int_{S_{\mathcal{L}}} \nabla G_{\mathcal{K}}(\zeta) \cdot \mathrm{d}\zeta - \int_{S_{\mathcal{K}}} \nabla G_{\mathcal{J}}(\zeta) \cdot \mathrm{d}\zeta\,, \end{aligned} \qquad (21)$$

where the second equality follows from the fact that $G_{\mathcal{K}}(\zeta) = 1 - G_{\mathcal{J}}(\zeta)$ since $G^v(\zeta) \in \mathcal{S}_v$. Now, Property (a) of Definition 7 implies that $\partial G_{\mathcal{K}}(\zeta)/\partial \zeta_l \geq 0$ for all $l \in \mathcal{L}$, and $\partial G_{\mathcal{J}}(\zeta)/\partial \zeta_k \geq 0$ for all $k \in \mathcal{K}$. It follows that $\nabla G_{\mathcal{J}}(\zeta) \cdot \mathrm{d}\zeta \geq 0$ along $S_{\mathcal{K}}$, and $\nabla G_{\mathcal{K}}(\zeta) \cdot \mathrm{d}\zeta \leq 0$ along $S_{\mathcal{L}}$. Substituting in (21), one gets the first inequality in (20). The second inequality in (20) follows by similar arguments. Then, one has

$$\begin{aligned} 0 \;&\geq\; G_{\mathcal{K}}(\nu) - G_{\mathcal{K}}(\varsigma) + G_{\mathcal{L}}(\varsigma) - G_{\mathcal{L}}(\nu) \\ &=\; \sum_{e \in \mathcal{E}_v^+} \mathrm{sgn}(\nu_e - \varsigma_e)\,(G_e^v(\nu) - G_e^v(\varsigma))\,, \end{aligned}$$

which proves the claim. ∎

We can now exploit Lemma 1 in order to prove the following key result guaranteeing that the solution of the local dynamical system (18) with constant input $\lambda(t) \equiv \lambda$ converges to a limit point which depends on the value of λ but not on the initial condition. (Cf. Example 4 and Figure 6.)

Lemma 2: (Existence of a globally attractive limit flow for the local dynamical system under constant input) Let $0 \leq v < n$ be a non-destination node, and λ a nonnegative-real constant. Assume that $G^v : \mathcal{R}_v \to \mathcal{S}_v$ is continuously differentiable and satisfies Property (a) of Definition 7. Then, there exists a unique $f^*(\lambda) \in \mathrm{cl}(\mathcal{F}_v)$ such that the solution of the dynamical system (18) with constant input $\lambda(t) \equiv \lambda$ satisfies

$$\lim_{t \to \infty} f_e(t) = f_e^*(\lambda)\,, \qquad \forall e \in \mathcal{E}_v^+\,,$$

for every initial condition $\rho^v(0) \in \mathcal{R}_v$.

Proof: Let us fix some $\lambda \in \mathbb{R}_+$. For $\nu \in \mathcal{R}_v$, and $t \geq 0$, let $\Phi^t(\nu) := \rho^v(t)$ be the value of the solution of (18) with constant input $\lambda(t) \equiv \lambda$ and initial condition $\rho^v(0) = \nu$, at time t. Also, let $\Psi^t(\nu) \in \mathcal{R}_v$ be defined by $\Psi^t_e(\nu) = \mu_e(\Phi^t_e(\nu))$, for every $e \in \mathcal{E}^+_v$. Now, fix two initial conditions $\nu, \varsigma \in \mathcal{R}_v$, and define $\chi(t) := \Phi^t(\nu) - \Phi^t(\varsigma)$, and $\xi(t) := \Psi^t(\nu) - \Psi^t(\varsigma)$. Since $\mu_e(\rho_e)$ is increasing by Assumption 2, for all $e \in \mathcal{E}^+_v$, one has that $\mathrm{sgn}(\chi_e(t)) = \mathrm{sgn}(\xi_e(t))$ for all $t \geq 0$. On the other hand, using Lemma 1, one gets

$$\sum_{e \in \mathcal{E}^+_v} \mathrm{sgn}(\chi_e(t)) \left(G^v_e(\Phi^t(\nu)) - G^v_e(\Phi^t(\varsigma)) \right) \leq 0 \,,$$

for all $t \geq 0$. It follows that, if

$$\varphi(t) := \lambda \sum_{e \in \mathcal{E}^+_v} \mathrm{sgn}(\chi_e(t)) \left(G^v_e(\Phi^t(\nu)) - G^v_e(\Phi^t(\varsigma)) \right) \,,$$

then, $\varphi(t) \leq 0$ for $t \geq 0$, and

$$\begin{aligned} ||\chi(t)||_1 &= ||\chi(0)||_1 + \int_0^t \left(\varphi(s) - ||\xi(s)||_1 \right) \mathrm{d}s \\ &\leq ||\chi(0)||_1 - \int_0^t ||\xi(s)||_1 \mathrm{d}s \,, \end{aligned}$$

for all $t \geq 0$. Rearranging the inequality above gives that

$$\int_0^t ||\xi(s)||_1 \mathrm{d}s \leq ||\chi(0)||_1 \,, \qquad \forall t \geq 0 \,. \tag{22}$$

Therefore, (each component of) $\xi(t)$ is absolutely integrable, and hence $\xi(t)$ is integrable for all $t \geq 0$. It is now standard to show how this implies global convergence to a unique limit flow. For completeness, we provide such argument below.

Let

$$h_t := \frac{\mathrm{d}}{\mathrm{d}t} \Psi^t(\nu) \,, \qquad t \geq 0 \,.$$

By applying the mean value theorem twice, one gets that, for all $z > 0$,

$$\Psi^t(\nu) = \frac{1}{z} \int_t^{t+z} \Psi^s(\nu) \mathrm{d}s - z_t h_{s^*_t} \,,$$

for some $z_t \in [0, z]$ and $s^*_t \in [t, t + z_t]$. On the other hand, observe that Assumption 2 implies that $\mathrm{d}\mu_e / \mathrm{d}\rho_e$ is bounded and (18) under the condition $\lambda(t) \equiv \lambda$ implies that

$$-f^{\max}_e \leq \frac{\mathrm{d}\rho_e}{\mathrm{d}s} \leq \lambda \quad \forall s \geq 0$$

Hence, there exists some positive constant M such that

$$||h_s||_1 = \sum_{e \in \mathcal{E}^+_v} \left| \frac{\mathrm{d}\mu_e}{\mathrm{d}\rho_e} \frac{\mathrm{d}\rho_e}{\mathrm{d}s} \right| \leq M \,, \qquad \forall s \geq 0 \,.$$

For a given $z > 0$, by choosing $\varsigma = \Phi^z(\nu)$, and defining the quantity $\kappa_z := z_t h_{s^*_t} - z_0 h_{s^*_0}$, one gets that, for all $t \geq 0$,

$$\begin{aligned} \Psi^t(\nu) + \kappa_z &= \Psi^0(\nu) + \frac{1}{z} \int_t^{t+z} \Psi^s(\nu) \mathrm{d}s - \frac{1}{z} \int_0^z \Psi^s(\nu) \mathrm{d}s \\ &= \Psi^0(\nu) + \frac{1}{z} \int_z^{t+z} \Psi^s(\nu) \mathrm{d}s - \frac{1}{z} \int_0^t \Psi^s(\nu) \mathrm{d}s \\ &= \Psi^0(\nu) + \frac{1}{z} \int_0^t \left(\Psi^{s+z}(\nu) - \Psi^s(\nu) \right) \mathrm{d}s \\ &= \Psi^0(\nu) + \frac{1}{z} \int_0^t \xi(s) \mathrm{d}s \,. \end{aligned}$$

Since $\xi(t)$ is integrable, the above shows that, for every $z > 0$, $\Psi^t(\nu) + \kappa_z$ is convergent as t grows large. Then, for every nondecreasing sequence of time instants $t_1 \leq t_2 \leq \ldots$, such that $\lim_k t_k = \infty$, arbitrariness of z and the bound

$$||\kappa_z||_1 \leq z_t ||h_{s^*_t}||_1 + z_0 ||h_{s^*_0}|| \leq 2zM$$

imply that the sequence $\{\Psi^{t_k}(\nu) : k \geq 1\}$ is Cauchy, and hence convergent in $\mathrm{cl}(\mathcal{F}_v)$. It follows that $\Psi^t(\nu)$ converges to some limit flow $f^*(\lambda, \nu) \in \mathrm{cl}(\mathcal{F}_v)$ depending both on the inflow λ, as well as possibly on the initial density ν. Moreover, using (22) again, one gets that

$$\begin{aligned} 0 &= \lim_{t \to \infty} \frac{1}{t} \int_0^t ||\Psi^s(\nu) - \Psi^s(\varsigma)||_1 \mathrm{d}s \\ &= ||f^*(\lambda, \nu) - f^*(\lambda, \varsigma)||_1 \,, \end{aligned}$$

for every $\nu, \varsigma \in \mathcal{R}_v$, which shows that in fact the limit flow does not depend on the initial condition. ∎

Now, let us define

$$\lambda^{\max}_v := \sum_{e \in \mathcal{E}^+_v} f^{\max}_e \,.$$

The following result characterizes the way the local limit flow $f^*(\lambda)$ depends on the local input λ. (Cf. Example 4 and Figure 5.)

Lemma 3 (Dependence of the limit flow on the input): Let $0 \leq v < n$ be a non-destination node, and λ a nonnegative-real constant. Assume that $G^v : \mathcal{R}_v \to \mathcal{S}_v$ is continuously differentiable and satisfies Properties (a) and (b) of Definition 7. Let $f^*(\lambda) \in \mathrm{cl}(\mathcal{F}_v)$ be the limit flow of the local system (18) with constant input $\lambda(t) \equiv \lambda$. Then, the following statements hold true:

(i) if $\lambda < \lambda^{\max}_v$, then

$$f^*_e(\lambda) < f^{\max}_e \,, \ \lambda G^v_e(\mu^{-1}(f^*(\lambda))) = f^*_e(\lambda) \ \forall e \in \mathcal{E}^+_v$$

(ii) if $\lambda \geq \lambda^{\max}_v$, then $f^*_e(\lambda) = f^{\max}_e \quad \forall e \in \mathcal{E}^+_v$.

Moreover, $f^*(\lambda)$ is continuous as a map from $\mathbb{R}_+$ to $\mathrm{cl}(\mathcal{F}_v)$, and each component $f_e(\lambda)$ is nondecreasing in λ.

Proof: Let $\rho^* \in \mathcal{R}_v$ be such that

$$\rho^*_e := \begin{cases} \mu^{-1}_e(f^*_e(\lambda)) & \text{if} \quad f^*_e(\lambda) < f^{\max}_e, \\ +\infty & \text{if} \quad f^*_e(\lambda) = f^{\max}_e \,, \end{cases}$$

for every $e \in \mathcal{E}^+_v$. Now, by contradiction, assume that there exists a nonempty proper subset $\mathcal{J} \subset \mathcal{E}^+_v$ such that ρ^*_j is finite for every $j \in \mathcal{J}$, and ρ^*_k is infinite for every $k \in \mathcal{K} := \mathcal{E}^+_v \setminus \mathcal{J}$. Thanks to Property (b) of Definition 7, one would have that, for any initial condition $\rho^v(0) \in \mathcal{R}_v$, the solution of (18) satisfies

$$\lim_{t \to \infty} \sum_{k \in \mathcal{K}} \left(\lambda G^v_k(\rho^v(t)) - f_k(t) \right) = - \sum_{k \in \mathcal{K}} f^{\max}_k < 0 \,,$$

so that there would exist some $z \geq 0$ such that $\sum_{k \in \mathcal{K}} (\lambda G^v_k(\rho^v(t)) - f_k(t)) \leq 0$ for every $t \geq z$. Hence, if $\rho_{\mathcal{K}}(t) := \sum_{k \in \mathcal{K}} \rho_k(t)$, then for every $t \geq z$ one would have

$$\begin{aligned} \rho_{\mathcal{K}}(t) &= \rho_{\mathcal{K}}(z) + \int_z^t \sum_{k \in \mathcal{K}} \left(\lambda G^v_k(\rho^v(s)) - f_k(s) \right) \right) \mathrm{d}s \\ &\leq \rho_{\mathcal{K}}(z) \\ &< +\infty \,, \end{aligned}$$

which would contradict the assumption that $\rho_k^* = +\infty$ for every $k \in \mathcal{K}$. Therefore, either ρ_e^* is finite for every $e \in \mathcal{E}_v^+$, or ρ_e^* is infinite for every $e \in \mathcal{E}_v^+$.

If ρ_e^* is finite for every $e \in \mathcal{E}_v^+$, then ρ^* is necessarily an equilibrium, being a finite limit point of the autonomous dynamical system (18) with continuous right-hand side, and so $f^*(\lambda)$ is an equilibrium flow for the local dynamical system (18). The contrapositive of this proves (ii) of the Lemma. On the other hand, when ρ_e^* is infinite for every $e \in \mathcal{E}_v^+$, so that $\sum_{e\in\mathcal{E}_v^+} f_e^* = \sum_{e\in\mathcal{E}_v^+} f_e^{\max} = \lambda_v^{\max}$. Then, necessarily $\lambda_v^{\max} \leq \lambda$ for otherwise, if $\lambda_v^{\max} - \lambda = \varepsilon > 0$, then $\frac{\mathrm{d}}{\mathrm{d}t}\sum_{e\in\mathcal{E}_v^+} \rho_e(t) \leq -\varepsilon/2 < 0$ for all t large enough, thus contradicting the fact that $\sum_{e\in\mathcal{E}_v^+} \rho_e(t)$ diverges as t grows large. The contrapositive of this proves (i) of the Lemma.

Finally, it remains to prove continuity of $f^*(\lambda)$ as a function of λ. For this, consider the function $H : (0,+\infty)^{\mathcal{E}_v^+} \times (0,\lambda_v^{\max}) \to \mathbb{R}^{\mathcal{E}_v^+}$ defined by $H_e(\rho^v,\lambda) := \lambda G_e^v(\rho^v) - \mu_e(\rho_e)$. Clearly, H is differentiable and such that

$$\begin{aligned} \frac{\partial}{\partial_{\rho_e}} H_e(\rho^v,\lambda) &= \lambda \frac{\partial}{\partial_{\rho_e}} G_e^v(\rho^v) - \mu_e'(\rho_e) \\ &= -\sum_{j\neq e} \lambda \frac{\partial}{\partial_{\rho_e}} G_j^v(\rho^v) - \mu_e'(\rho_e) \\ &< -\sum_{j\neq e} \frac{\partial}{\partial_{\rho_e}} H_j(\rho^v,\lambda)\,, \end{aligned} \tag{23}$$

where the inequality follows from the strict monotonicity of the flow function (see Assumption 2). Property (a) in Definition 7 implies that $\partial H_j(\rho^v,\lambda)/\partial\rho_e \geq 0$ for all $j \neq e \in \mathcal{E}_v^+$. Hence, from (23), we also have that $\partial H_e(\rho^v,\lambda)/\partial\rho_e < 0$ for all $e \in \mathcal{E}_v^+$. Therefore, for all $\rho^v \in (0,+\infty)^{\mathcal{E}_v^+}$, and $\lambda \in (0,\lambda_v^{\max})$, one has that

$$\begin{aligned} \left|\frac{\partial}{\partial_{\rho_e}} H_e(\rho^v,\lambda)\right| &= -\frac{\partial}{\partial_{\rho_e}} H_e(\rho^v,\lambda) \\ &> \sum_{j\neq e} \frac{\partial}{\partial_{\rho_e}} H_j(\rho^v,\lambda) \\ &= \sum_{j\neq e} \left|\frac{\partial}{\partial_{\rho_e}} H_j(\rho^v,\lambda)\right|, \end{aligned}$$

i.e., (the transpose of) the Jacobian matrix $\nabla_{\rho^v} H(\rho^v,\lambda)$ is strictly diagonally dominant, and hence invertible by a standard application of the Gershgorin Circle Theorem, e.g., see [33, Theorem 6.1.10]. It then follows from the implicit function theorem that $\rho^*(\lambda)$, which is the unique zero of $H(\,\cdot\,,\lambda)$, is continuous on the interval $(0,\lambda_v^{\max})$. Hence, also $f^*(\lambda) = \mu(\rho^*(\lambda))$ is continuous on $(0,\lambda_v^{\max})$, since it is the composition of two continuous functions. Moreover, since $\sum_{e\in\mathcal{E}_v^+} f_e^*(\lambda) = \lambda$ for every $\lambda \in (0,\lambda_v^{\max})$, and $0 \leq f_e^*(\lambda) \leq f_e^{\max}$, one gets that $\lim_{\lambda\downarrow 0} f_e^*(\lambda) = 0$ and $\lim_{\lambda\uparrow\lambda_v^{\max}} f_e^*(\lambda) = f_e^{\max}$. Now, one has that $\sum_e f_e^*(0) = 0$, so that $\lim_{\lambda\downarrow 0} f_e^*(\lambda) = f_e^*(0) = 0$ for all $e \in \mathcal{E}_v^+$. Moreover, as previously shown, $f_e^*(\lambda) = f_e^{\max} = \lim_{\lambda\uparrow\lambda_v^{\max}} f_e^*(\lambda)$ for $\lambda \geq \lambda_v^{\max}$. This completes the proof of continuity of $f^*(\lambda)$ on $[0,+\infty)$. Monotonicity of each component $f_e(\lambda)$ follows in turn from standard arguments for monotone dynamical systems, see, e.g. [14, Remark V.2]. ∎

While Lemma 2 ensures existence of a unique limit point for the local system (18) with constant input $\lambda(t) \equiv \lambda$, the following lemma establishes that the output of the local system (18) is convergent, provided that the input is convergent.

Lemma 4 (Attractivity of the local dynamical system): Let $0 \leq v < n$ be a nondestination node, $G^v : \mathcal{R}_v \to \mathcal{S}_v$ a continuously differentiable map, satisfying Properties (a) and (b) of Definition 7, and $\lambda(t)$ a nonnegative-real-valued Lipschitz continuous function such that

$$\lim_{t\to\infty} \lambda(t) = \lambda\,. \tag{24}$$

Then, for every initial condition $\rho^v(0) \in \mathcal{R}_v$, the solution of the local dynamical system (18) satisfies

$$\lim_{t\to\infty} f_e(t) = f_e^*(\lambda)\,, \qquad \forall e \in \mathcal{E}_v^+\,, \tag{25}$$

where $f^*(\lambda)$ is the limit flow of the local system (18) with constant input $\lambda(t) \equiv \lambda$.

Proof: Lemma 2 guarantees that the local systems (18) is endowed with the static input-output characteristic $f^*(\lambda)$. Then, the result follows immediately from [14, Proposition V.8]. ∎

We are now ready to prove Theorem 1 by showing that, for any initial condition $\rho(0) \in \mathcal{R}$, the solution of the dynamical network (6) satisfies

$$\lim_{t\to\infty} f_e(t) = f_e^*\,, \tag{26}$$

for all $e \in \mathcal{E}$. We shall prove this by showing via induction on $v = 0, 1, \ldots, n-1$ that, for all $e \in \mathcal{E}_v^+$, there exists $f_e^* \in [0, f_e^{\max}]$ such that (26) holds true. First, observe that, thanks to Lemma 2, this statement is true for $v = 0$, since the total outflow at the origin node is constant at λ_0. Now, assume that the statement is true for all $0 \leq v < w$, where $w \in \{1,\ldots,n-2\}$ is some intermediate node. Then, since $\mathcal{E}_w^- \subseteq \cup_{v=0}^{w-1}\mathcal{E}_v^+$, one has that

$$\lim_{t\to\infty} \lambda_w^-(t) = \lim_{t\to\infty} \sum_{e\in\mathcal{E}_w^-} f_e(t) = \sum_{e\in\mathcal{E}_w^-} f_e^* = \lambda_w^*\,.$$

Then, Lemma 4 implies that, for all $e \in \mathcal{E}_w^+$, (26) holds true with $f_e^* = f_e^*(\lambda_w^*)$, thus proving the statement for $v = w$. This proves the existence of a globally attractive limit flow f^*. The proof of Theorem 1 is completed by Lemma 3.

V. Proof of Theorem 2

This section is devoted to the proof of Theorem 2 on the weak resilience of dynamical networks with locally responsive distributed routing policies $\mathcal{G}$.

To start with, let us recall that, in this case Theorem 1 implies the existence of a globally attractive limit flow $\tilde{f}^* \in \mathrm{cl}(\tilde{\mathcal{F}})$ for the perturbed dynamical 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 5: Consider a dynamical 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 $\tilde{f}^* \in \mathrm{cl}(\tilde{\mathcal{F}})$ is the limit flow vector of the corresponding perturbed dynamical network (10), 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 $\tilde{f}_e^* \leq \tilde{f}_e^{\max}/2$.

Proof: First, observe that the claim is trivially true when $\tilde{f}_e^* > \tilde{f}_e^{\max}/2$ for all $e \in \mathcal{E}$. Therefore, let us assume that there exists a node $v \in \{0, \dots, n-1\}$ such that $\tilde{f}_e^* \leq \tilde{f}_e^{\max}/2$ for some link $e \in \mathcal{E}_v^+$. Define $\rho^\theta \in \mathcal{R}_v$ by $\rho_j^\theta = 0$ for all $j \in \mathcal{E}_v^+$, $j \neq e$, and $\rho_e^\theta = \theta \rho_e^\mu$, where recall that ρ_e^μ is the median density of the flow function μ_e. Since the stretching coefficient of $\tilde{\mathcal{N}}$ is less than or equal to θ, one has that the median densities of the perturbed and the unperturbed flow functions satisfy $\tilde{\rho}_e^\mu \leq \theta \rho_e^\mu$. This and the fact that $\tilde{f}_e^* \leq \tilde{f}_e^{\max}/2$ imply that $\tilde{\rho}_e^* \leq \tilde{\rho}_e^\mu \leq \rho_e^\theta$, while clearly $\tilde{\rho}_j^* \geq 0 = \rho_j^\theta$ for all $j \in \mathcal{E}_v^+$, $j \neq e$. Now, let $\beta_\theta := G_e^v(\rho^\theta)$, and observe that, thanks to the assumption on the strict positivity of $G_e^v(\rho^v)$, one has $\beta_\theta > 0$. Then, from Lemma 1 one gets that

$$\begin{aligned} G_e^v(\tilde{\rho}^*) &= \frac{1}{2}\left(G_e^v(\tilde{\rho}^*) + 1 - \sum\nolimits_{j \neq e} G_j^v(\tilde{\rho}^*)\right) \\ &\geq \frac{1}{2}\left(G_e^v(\rho^\theta) + 1 - \sum\nolimits_{j \neq e} G_j^v(\rho^\theta)\right) \\ &= G_e^v(\rho^\theta) \\ &= \beta_\theta \,. \end{aligned} \tag{27}$$

On the other hand, since $\tilde{f}_e^* \leq \tilde{f}_e^{\max}/2 < \tilde{f}_e^{\max}$, Lemma 2 implies that necessarily

$$\tilde{\lambda}_v^* G_e^v(\tilde{\rho}^*) = \tilde{f}_e^* \,.$$

The claim now follows by combining this and (27). ∎

As a consequence of Lemma 5, we now prove the following result showing that the dynamical network is partially transferring and providing a lower bound on its weak resilience.

Lemma 6: Let $\mathcal{N}$ be a network satisfying Assumptions 1 and 2, $\lambda_0 \geq 0$ a constant outflow at the origin node, 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 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 5.

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

$$\delta \leq C(\mathcal{N}) - 2|\mathcal{E}|\lambda_0 \beta_\theta^{1-n} \alpha \,, \tag{28}$$

and stretching coefficient less than or equal to $\theta \geq 1$. We shall iteratively select a sequence of nodes $v_0, v_1, \dots, v_k$ such that $v_0 = 0$, $v_k = n$, and, for every $1 \leq j \leq k$, there exists $i \in \{0, \dots, j-1\}$ such that

$$(v_i, v_j) \in \mathcal{E} \,, \qquad \tilde{f}_{(v_i,v_j)}^* \geq \lambda_0 \alpha \beta_\theta^{j-n} \,. \tag{29}$$

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

$$\lim_{t \to \infty} \tilde{\lambda}_n(t) = \tilde{\lambda}_n^* = \sum_{e \in \mathcal{E}_n^-} \tilde{f}_e^* \geq \alpha \lambda_0 \beta_\theta^{k-n} \geq \alpha \lambda_0 \,, \tag{30}$$

so that the perturbed dynamical network is α-transferring. Moreover, observe that the trivial perturbation $\tilde{\mathcal{N}} = \mathcal{N}$ has magnitude $\delta = 0$, hence it satisfies (28) for all $\alpha \in (0, C(\mathcal{N})\beta_\theta^{n-1}/(2|\mathcal{E}|\lambda_0)]$. Therefore, (30) will imply the partial transferring property of the original dynamical 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 5, that $\tilde{f}_e^{\max} < 2\tilde{f}_e^*$ for all $e \in \mathcal{E}_0^+$, so that

$$\sum_e \tilde{f}_e^{\max} < 2 \sum_e \tilde{f}_e^* < 2\alpha |\mathcal{E}_0^+| \beta_\theta^{1-n} \lambda_0 \leq 2\alpha |\mathcal{E}| \beta_\theta^{1-n} \lambda_0 \,,$$

where the summation index e runs over $\mathcal{E}_0^+$. Combining the above with the inequality $C(\mathcal{N}) \leq \sum_{e \in \mathcal{E}_0^+} f_e^{\max}$, one would get

$$\delta \geq \sum_{e \in \mathcal{E}_0^+} \left(f_e^{\max} - \tilde{f}_e^{\max}\right) > C(\mathcal{N}) - 2\alpha |\mathcal{E}| \beta_\theta^{1-n} \lambda_0 \,,$$

thus contradicting the assumption (28). Hence, necessarily there exists $e \in \mathcal{E}_0^+$ such that $\tilde{f}_e^* \geq \lambda_0 \alpha \beta_\theta^{1-n}$, and choosing $v_1 \in \mathcal{V}$ to be the head node of such e, one sees that (29) holds true with $j = 1$.

Now, fix some $1 < j^* \leq k$, and assume that (29) holds true for every $1 \leq j < j^*$. Then, by choosing i as in (29), one gets that

$$\tilde{\lambda}_{v_j}^* = \sum_{e \in \mathcal{E}_{v_j}^-} \tilde{f}_e^* \geq \tilde{f}_{(v_i,v_j)}^* \geq \lambda_0 \alpha \beta_\theta^{j-n} \geq \lambda_0 \alpha \beta_\theta^{j^*-1-n} \,, \tag{31}$$

for every $1 \leq j < j^*$. Moreover,

$$\tilde{\lambda}_{v_0}^* = \lambda_0 \geq \lambda_0 \alpha \beta_\theta^{-n} \geq \lambda_0 \alpha \beta_\theta^{j^*-1-n} \,. \tag{32}$$

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 every $e \in \mathcal{E}_\mathcal{U}^+$. Thanks to (31) and (32), this would imply that $\tilde{f}_e^* < \beta_\theta \tilde{\lambda}_{v_j}^*$, for every $e \in \mathcal{E}_{v_j}^+ \cap \mathcal{E}_\mathcal{U}^+$ with $0 \leq j < j^*$. Then, since $\mathcal{E}_\mathcal{U}^+ = \cup_{j=0}^{j^*-1} (\mathcal{E}_{v_j}^+ \cap \mathcal{E}_\mathcal{U}^+)$, Lemma 5 would imply that

$$\tilde{f}_e^{\max} < 2\tilde{f}_e^* \,, \qquad \forall e \in \mathcal{E}_\mathcal{U}^+ \,.$$

This would yield

$$\sum_e \tilde{f}_e^{\max} < \sum_e 2\tilde{f}_e^* < 2 \sum_e \lambda_0 \alpha \beta_\theta^{j^*-n} \leq 2|\mathcal{E}| \lambda_0 \alpha \beta_\theta^{1-n} \,,$$

where the summation index e runs over $\mathcal{E}_\mathcal{U}^+$. From the above and the inequality $C(\mathcal{N}) \leq \sum_{e \in \mathcal{E}_\mathcal{U}^+} f_e^{\max}$, one would get

$$\delta \geq \sum_{e \in \mathcal{E}_\mathcal{U}^+} \left(f_e^{\max} - \tilde{f}_e^{\max}\right) > C(\mathcal{N}) - 2\alpha |\mathcal{E}| \beta_\theta^{1-n} \lambda_0 \,,$$

thus contradicting the assumption (28). Hence, necessarily there exists $e \in \mathcal{E}_{\mathcal{U}}^{+}$ such that $\tilde{f}_e^* \geq \lambda_0 \alpha \beta_\theta^{j^*-n}$, and choosing $v_{j^*} \in \mathcal{V} \setminus \mathcal{U}$ to be the head node of e one sees that (29) holds true with $j = j^*$. Iterating the argument above until $v_{j^*} = n$ yields the claim. ∎

It is now easy to see that Lemma 6 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.

VI. Conclusion

In this paper, we studied robustness properties of dynamical networks, where the dynamics on every link is driven by the difference between the inflow, which depends on the upstream routing decisions, and the outflow, which depends on the particle density, on that link. 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 the min-cut capacity of the network, and that it is independent of the locality constraint on the information available, as well as of the initial flow. Strong resilience of dynamical networks is studied in the companion paper [3]. The main technical assumptions of this paper are acyclicity of the network topology, monotonicity of the flow function, and the absence of a bound on the density on the links. Relaxing such assumptions is part of ongoing work.

Acknowledgements

The authors wish to thank Prof. Devavrat Shah of MIT for helpful discussions. They are also grateful to the Associate Editor and three anonymous Referees for several comments and additional references which have allowed to improve the presentation and contextualize the paper.

References

[1] G. Como, K. Savla, D. Acemoglu, M. A. Dahleh, and E. Frazzoli, "Robust distributed routing in dynamical flow networks," in *IEEE Conference on Decision and Control and European Control Conference*, (Orlando, FL), 2011.

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

[3] 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*, vol. 58, no. 2, 2013. To appear, Available at `http://arxiv.org/abs/1103.4893`.

[4] M. Garavello and B. Piccoli, *Traffic Flow on Networks*. American Institute of Mathematical Sciences, 2006.

[5] D. Bertsekas and R. Gallager, *Data Networks*. Prentice Hall, 2nd ed., 1992.

[6] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *Journal of the Operation Research Society*, vol. 49, pp. 237–252, March 1998.

[7] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: a mathematical theory of network architectures," *Proceedings of the IEEE*, vol. 95, pp. 255–312, January 2007.

[8] U. S. Karmarkar, "Capacity loading and release planning with work-in-progress (WIP) and leadtimes," *Journal of Manufacturing and Operations Management*, vol. 2, pp. 105–123, 1989.

[9] C. D'Aspice, S. Gottlich, M. Herty, and B. Piccoli, *Modeling, Simulation, and Optimization of Supply Chains: A Continuous Approach*. SIAM, 2010.

[10] S. Gottlich, M. Herty, and A. Klar, "Network models for supply chains," *Communications in Mathematical Sciences*, vol. 3, no. 4, pp. 545–559, 2005.

[11] 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.

[12] 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.

[13] M. W. Hirsch and H. L. Smith, "Monotone dynamical systems," in *Ordinary Differential Equations, vol. II* (A. Canada, P. Drabek, and A. Fonda, eds.), Handbook of Differential Equations, Elsevier, 2005.

[14] D. Angeli and E. D. Sontag, "Monotone control systems," *IEEE Transactions on Automatic Control*, vol. 48, no. 10, pp. 1684–1698, 2003.

[15] G. Vinnicombe, "Robust congestion control for the internet," in *SIGCOMM*, 2002.

[16] F. Paganini, "A global stability result in network flow control," *Systems and Control Letters*, vol. 46, no. 3, pp. 165–172, 2002.

[17] S. H. Low, F. Paganini, and J. C. Doyle, "Internet congestion control," *IEEE Control Systems Magazine*, vol. 22, no. 1, pp. 28–43, 2002.

[18] 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.

[19] R. Albert, H. Jeong, and A. Barabasi, "Error and attack tolerance of complex networks," *Nature*, vol. 406, pp. 378–382, 2000.

[20] J. C. Doyle, D. A. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger, "The "robust yet fragile" nature of the internet," *Proceedings of the National Academy of Sciences*, vol. 102, no. 41, pp. 14497–14502, 2005.

[21] Y. Vardi and C. Zhang, "Measures of network vulnerability," *IEEE Signal Processing Letters*, vol. 14, pp. 313–316, May 2007.

[22] S. Bulteau and G. Rubino, "Evaluating network vulnerability with the mincuts frequency vector," Tech. Rep. 1089, IRISA, France, March 1997.

[23] 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.

[24] R. K. Wood, "Deterministic network interdiction," *Mathematical and Computer Modelling*, vol. 17, no. 2, pp. 1–18, 1993.

[25] 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.

[26] A. S. Acampora and S. I. A. Shah, "Multihop lightwave networks: a comparison of store-and-forward and hot-potato routing," *IEEE Transactions on Communications*, vol. 40, no. 6, pp. 1082–1090, 1992.

[27] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*. MIT Press, 2nd ed., 2001.

[28] C. Daganzo, *A theory of supply chains*. Springer, 2003.

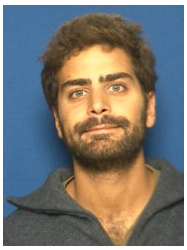
[29] G. Como, K. Savla, D. Acemoglu, M. A. Dahleh, and E. Frazzoli, "Stability analysis of transportation networks with multiscale driver decisions," *SIAM Journal on Control and Optimization*, 2011. Under revision, Available at `http://arxiv.org/abs/1101.2220`.

[30] 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.

[31] M. Ben-Akiva and S. Lerman, *Discrete Choice Analysis: Theory and application to travel demand*. MIT press, 1985.

[32] H. L. Smith, *Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems*. American Mathematical Society, 1995.

[33] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1990.



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



Ketan Savla received the B.Tech. degree in mechanical engineering from the Indian Institute of Technology Bombay, Mumbai, India, in 2003, the M.S. degree in mechanical engineering from the University of Illinois at Urbana-Champaign, Urbana, in 2004, and the M.A. degree in applied mathematics and the Ph.D. degree in electrical engineering from the University of California at Santa Barbara (UCSB), Santa Barbara, both in 2007.

He is currently an Assistant Professor in the Sonny Astani Department of Civil and Environmental Engineering at the University of Southern California, Los Angeles, CA. Prior to that, he was a Postdoctoral Associate and then a Research Scientist in the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA. His current research interests include control and optimization techniques for mobile autonomous systems, human-in-loop systems and complex service networks.

Dr. Savla was awarded the Center for Control, Dynamical Systems, and Computation best thesis from UCSB in 2010.



Daron Acemoglu is the Elizabeth and James Killian Professor of Economics in the Department of Economics at the Massachusetts Institute of Technology, Cambridge, MA and a member of the Economic Growth program of the Canadian Institute of Advanced Research. He is also affiliated with the National Bureau Economic Research, the Center for Economic Performance, the Center for Economic Policy Research, and Microsoft Research Center. He is an elected fellow of the American Academy of Arts and Sciences, the Econometric Society, the European Economic Association, and the Society of Labor Economists.

Daron Acemoglu has received a BA in economics at the University of York, 1989, M.Sc. in mathematical economics and econometrics at the London School of Economics, 1990, and Ph.D. in economics at the London School of Economics in 1992. Since 1993, he has held the academic positions of Lecturer at the London School of Economics, and Assistant Professor, Pentti Kouri Associate Professor and Professor of Economics at MIT. He has received numerous awards and fellowships, including the award for best paper published in the Economic Journal in 1996 for his paper "Consumer Confidence and Rational Expectations: Are Agents' Beliefs Consistent With the Theory?", the inaugural T. W. Shultz Prize from the University of Chicago in 2004, and the inaugural Sherwin Rosen Award for outstanding contribution to labor economics in 2004, Distinguished Science Award from the Turkish Sciences Association in 2006, the John von Neumann Award, Rajk College, Budapest in 2007. He was also awarded the John Bates Clark Medal in 2005, given every two years to the best economist in the United States under the age of 40 by the American Economic Association, and holds an Honorary Doctorate from the University of Utrecht.

Daron Acemoglu's research covers a wide range of areas within economics, including political economy, economic development and growth, human capital theory, growth theory, innovation, search theory, network economics and learning. Daron Acemoglu is also the co-editor of Econometrica and of the National Bureau of Economic Research Macroeconomic Annual. Daron has recently completed a new book, along with Professor James Robinson (Harvard University), entitled "Why Nations Fail: The Origins of Power, Prosperity and Poverty".



Munther A. Dahleh (S'84–M'87–SM'97–F'01) was born in 1962. He received the B.S. degree in electrical engineering from Texas A&M University, College Station, and the Ph.D. degree in electrical engineering from Rice University, Houston, TX, in 1983 and 1987, respectively.

Since then, he has been with the Department of Electrical Engineering and Computer Science, MIT, Cambridge, where he is currently a Full Professor and the Associate Head. He is the coauthor (with I. Diaz-Bobillo) of the book "Control of Uncertain Systems: A Linear Programming Approach" (Englewood Cliffs, NJ: Prentice-Hall, 1995) and the coauthor (with N. Elia) of the book "Computational Methods for Controller Design" (New York: Springer, 1998). His research interests include problems at the interface of robust control, filtering, information theory, and computation. In addition to methodology development, he has been interested in the application of distributed control in the future electric grid and the future transportation system. He is also interested in various problems in network science. He is also interested in the interface between systems theory and neurobiology, and, in particular, in providing an anatomically consistent model of the motor control.

Dr. Dahleh served as an Associate Editor of the IEEE Transactions on Automatic Control and of Systems and Control Letters. He has been the recipient of the Ralph Budd Award in 1987 for the best thesis at Rice University, the George Axelby Outstanding Paper Award (paper coauthored with J. B. Pearson in 1987), the National Science Foundation Presidential Young Investigator Award in 1991, the Finmeccanica Career Development Chair in 1992, the Donald P. Eckman Award from the American Control Council in 1993, the Graduate Students Council Teaching Award in 1995, the George Axelby Outstanding Paper Award (paper coauthored with Bamieh and Paganini in 2004), the Hugo Schuck Award for Theory (for the paper coauthored with Martins), and the George Axelby Outstanding Paper Award (paper coauthored with Martins in 2010). He was a Plenary Speaker at the 1994 American Control Conference, at the Mediterranean Conference on Control and Automation in 2003, at the Mathematical Theory for Networks and Systems in 2006, at the Symposium on System Identification in 2009, at Asian Control Conference in 2009, and at the Sixth Spanish, Italian, and Netherlands Meeting on Game Theory in 2010.



Emilio Frazzoli (Senior Member, IEEE) received the Laurea degree in aerospace engineering from the University of Rome "La Sapienza", Rome, Italy, in 1994 and the Ph.D. degree in navigation and control systems from the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, in 2001.

He is an Associate Professor of Aeronautics and Astronautics with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology. His main research interests lie in the general area of planning and control for mobile cyber-physical systems, with a particular emphasis on autonomous vehicles, mobile robotics, and transportation networks.

Dr. Frazzoli was the recipient of a National Science Foundation (NSF) CAREER award in 2002. He is an Associate Fellow of the American Institute of Aeronautics and Astronautics. He is currently serving as an Associate Editor for the IEEE Transactions on Automatic Control, the AIAA Journal of Guidance, Control, and Dynamics, and the Journal of Optimization Theory and Applications.