Annual Reviews in Control 000 (2017) 1–11

Contents lists available at ScienceDirect

Annual Reviews in Control

journal homepage: www.elsevier.com/locate/arcontrol

On resilient control of dynamical flow networks☆

Giacomo Como[a,b,*]

[a] Department of Automatic Control, Lund University, BOX118, SE-22100 Lund, Sweden
[b] Lagrange Department of Mathematical Sciences, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129, Torino, Italy

ARTICLE INFO

Article history:
Received 30 September 2016
Revised 3 January 2017
Accepted 12 January 2017
Available online xxx

Keywords:
Dynamical flow networks
Transportation networks
Compartmental systems
Distributed control
Network resilience
Network flow control
Dynamical routing
Robust control
Monotone systems

ABSTRACT

Resilience has become a key aspect in the design of contemporary infrastructure networks. This comes as a result of ever-increasing loads, limited physical capacity, and fast-growing levels of interconnectedness and complexity due to the recent technological advancements. The problem has motivated a considerable amount of research within the last few years, particularly focused on the dynamical aspects of network flows, complementing more classical static network flow optimization approaches.

In this tutorial paper, a class of single-commodity first-order models of dynamical flow networks is considered. A few results recently appeared in the literature and dealing with stability and robustness of dynamical flow networks are gathered and originally presented in a unified framework. In particular, (differential) stability properties of monotone dynamical flow networks are treated in some detail, and the notion of margin of resilience is introduced as a quantitative measure of their robustness. While emphasizing methodological aspects —including structural properties, such as monotonicity, that enable tractability and scalability— over the specific applications, connections to well-established road traffic flow models are made.

1. Introduction

As critical infrastructure networks, such as transport and energy, are being utilized closer and closer to their capacity limits, the complex interaction between physical systems, cyber layers, and human decision makers has created new challenges in simultaneously achieving efficiency and reliability. The recent technological advancements in terms of smart sensors, high-speed communication, and real-time decision capabilities have exacerbated the large-scale interconnected nature of these systems, and increased both the potential gains associated to their optimization and their inherent systemic risks. In fact, while designed to perform well under normal operation conditions, such complex systems tend to exhibit critical fragilities in response to unforeseen disruptions. Even if simply started from small local perturbations, such disruptions have the potential to build up through cascading mechanisms driven by the interconnected dynamics of the infrastructure network, possibly leading to detrimental systemic effects. The term *resilience* refers to the ability of these systems "to plan and prepare for, absorb, respond to, and recover from disasters and adapt to new conditions" (definition by the US National Academy of Sciences 2012).

Whilst static network flow optimization has long been regarded as a fundamental design paradigm for infrastructure systems and represents a central area of mathematical programming (Ahuja, Magnanti, & Orlin, 1993; Bertsekas, 1998; Whittle, 2007), there is an increasing awareness that the full potential of the emerging technologies and the nature of the associated systemic risks can only be understood by developing systems modeling, robustness analysis, and control synthesis within a dynamical framework: this recognition is stimulating considerable interest in the control systems community. In this tutorial paper, based on a semi-plenary lecture given by the author at the 22nd International Symposium on the Mathematical Theory of Networks and Systems, a few recent results on stability and robustness of dynamical flow networks are presented. While emphasizing methodological aspects —in particular, structural properties enabling tractability and scalability of the considered models— over the specific applications, this paper also makes connections to well-established road traffic flow models.

Our focus is on first-order models of *dynamical flow networks*, describing the flow of mass among a finite set of interconnected cells. Such dynamical systems have sometimes been referred to as *compartmental systems* in some of the literature (Jacquez & Simon, 1993; Walter & Contreras, 1999). Our main interest is on *nonlinear* dynamical flow networks, with nonlinearities especially account-

☆ The author is a member of the excellence centres LCCC and ELLIIT. His research has been supported by the Swedish Research Council through a Project Research Grant.

* Corresponding author.
E-mail addresses: giacomo.como@control.lth.se, giacomo.como@polito.it

http://dx.doi.org/10.1016/j.arcontrol.2017.01.001
1367-5788/

ing for *congestion* effects. Special attention is devoted to *demand* and *supply* constraints limiting, respectively, the maximum outflow from and the maximum inflow in the cells as a function of their current mass, as in Daganzo's cell transmission model for road traffic network flows (Daganzo, 1994, 1995). We introduce a class of *monotone* dynamical flow networks, characterized by structural properties of the dependence of the flow variables on the network state. We show that, while this class is large enough to encompass many examples of applicative interest, the system structure of monotone dynamical flow networks is such that their dynamic behaviors, and especially stability and robustness properties, are analyzable in a tractable and scalable way. In particular, in this paper, we: (i) present results relating the (differential) stability of (nonlinear) monotone dynamical flow networks to graph-theoretical properties; (ii) introduce the notion of *margin of resilience* as a measure of their robustness against exogenous perturbations; and (iii) study a class of *locally responsive feedback* routing and flow control policies that are able to achieve the maximum possible margin of resilience for a given network topology in spite of relying on local information only and requiring no global knowledge of the network.

The remainder of this paper is organized as follows. In Section 2, we first introduce dynamical flow networks. In Section 3, we focus on their simplest instance, affine dynamical flow networks, for which we gather some results relating their (global, exponential) stability to outflow-connectivity properties of the network topology. In Section 4, we define the notion of monotone dynamical flow networks, present stability results that can be deduced for this system structure, and show how they can be applied to dual ascent dynamics for static convex network flow optimization. In Section 5, we introduce nonlinear dynamical flow networks with demand and supply constraints and show how several examples of network flow dynamics from the literature that can be fit in this framework also belong to the class of monotone dynamical flow networks (either globally or locally), so that the stability results of Section 4 can be successfully applied. In Section 6, we study robustness of nonlinear dynamical flow networks with respect to perturbations of the demand functions (hence, of the flow capacities) as well as of the external inflows. We introduce the notion of margin of resilience as a quantitative measure of robustness and compute the margin of resilience of different classes of distributed routing and flow control policies.

We end this introductory section by gathering some notational conventions to be adopted throughout the paper. The sets of real numbers and of nonnegative real numbers are denoted by $\mathbb{R}$ and $\mathbb{R}_+$, respectively. The all-one vector is denoted by $\mathbb{1}$, the all-zero vector simply by 0, the identity matrix by $I = \mathrm{diag}\,(\mathbb{1})$, and the transpose of a matrix M by M^T. Inequalities between vectors are meant to hold true entrywise, i.e., if $a, b \in \mathbb{R}^n$, then $a \geq b$ means that $a_i \geq b_i$ for all $i = 1, \dots, n$. A square matrix M is called: *nonnegative* if all its entries are nonnegative; *Metzler* if all its nondiagonal entries are nonnegative, i.e., $M_{ij} \geq 0$ for all $i \neq j$ (Berman & Plemmons, 1994); (row) *diagonally dominant* if $|M_{ii}| \geq \Sigma_{j\neq i}|M_{ij}|$; *compartmental* if it is Metzler and diagonally dominant; *Hurwitz* if all its eigenvalues have negative real part; and *substochastic* if it is nonnegative and such that $M\mathbb{1} \leq \mathbb{1}$, i.e., its rows all sum up to less than or equal to 1. A directed graph, shortly *digraph*, is the pair $(\mathcal{V}, \mathcal{E})$ of a finite node set $\mathcal{V}$ and a link set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$, whereby links $(i, j) \in \mathcal{E}$ are interpreted as pointing from node i to node j. A length-l *path* from a node i to a node j in a digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a sequence of nodes $\{v_0, v_1, \dots, v_l\} \subseteq \mathcal{V}$ such that $v_0 = i$, $v_l = j$, $v_h \neq v_k$ for all $0 \leq h < k \leq l$, and $(v_{h-1}, v_h) \in \mathcal{E}$ for all $h = 1, \dots, l$. The gradient of a function $f : \mathbb{R}^n \to \mathbb{R}^m$ in a point x of its domain is the matrix $\nabla f(x) \in \mathbb{R}^{m\times n}$ whose entries are the partial derivatives $[\nabla f(x)]_{ij} = \partial f_i(x)/\partial x_j$: if $m = h + l$ and the variable is explicitly written as $x = (y, z)$ where $y \in \mathbb{R}^h$ and $z \in \mathbb{R}^l$, then $\nabla_y f(y, z) \in \mathbb{R}^{h\times n}$ and $\nabla_z f(y, z) \in \mathbb{R}^{l\times n}$ are the left and right blocks of $\nabla f(x) \in \mathbb{R}^{m\times n}$, so that $\nabla f(x) = [\nabla_y f(y, z), \nabla_z f(y, z)]$.

2. Dynamical flow networks

This paper focuses on single-commodity, first-order models of dynamical flow networks, also referred to as compartmental systems in some of the literature (Jacquez & Simon, 1993; Walter & Contreras, 1999). These are dynamical systems with n-dimensional state vector $x = x(t)$ that belongs to the nonnegative orthant $\mathbb{R}^n_+$ at any time $t \geq 0$. The entries $x_i = x_i(t)$ of the state vector represent the mass in each *cell* $i \in \mathcal{I}$, where $\mathcal{I} := \{1, \dots, n\}$ is a finite set of interconnected cells. The network flow dynamics can be compactly expressed as

$$\dot{x} = u + F^T\mathbb{1} - F\mathbb{1} - w \tag{1}$$

where:

(i) $u \in \mathbb{R}^n_+$ is a nonnegative vector supported on a subset $\mathcal{R} \subseteq \mathcal{I}$ whose entries u_i model the *external inflows* in the cells $i \in \mathcal{R}$;
(ii) $F \in \mathbb{R}^{n\times n}_+$ is a nonnegative matrix supported on a subset $\mathcal{A} \subseteq \mathcal{I} \times \mathcal{I}$ of ordered pairs of adjacent cells whose nonzero entries F_{ij} represent the *flow* from cell i to cell j for all pairs $(i, j) \in \mathcal{A}$[1];
(iii) And $w \in \mathbb{R}^n_+$ is a nonnegative vector supported on a subset $\mathcal{S} \subseteq \mathcal{I}$ whose entries w_i model the *outflows* from the cells $i \in \mathcal{R}$ towards the external environment.

Hence, in particular, we have

$$u_i = 0\,, \quad i \notin \mathcal{R}\,, \qquad F_{ij} = 0\,, \quad (i, j) \notin \mathcal{A}\,, \qquad w_i = 0\,, \quad i \notin \mathcal{S}\,. \tag{2}$$

While in typical applications the external inflows are — either constant or time-varying — exogenous inputs, the flow variables F and w in general may depend both on the state x (thus allowing for feedback) and directly on the time t (thus allowing for exogenous time variability). Invariance of the nonnegative orthant $\mathbb{R}^n_+$ for the state vector x is guaranteed by the additional constraint

$$x_i = 0 \qquad \Longrightarrow \qquad w_i = 0\,, \quad F_{ij} = 0\,, \quad i, j \in \mathcal{I}\,, \tag{3}$$

i.e., the outflow from an empty cell is always 0. Entrywise rewriting of the dynamics (1) reads

$$\dot{x}_i = u_i + \sum_{j\in\mathcal{I}} F_{ji} - \sum_{j\in\mathcal{I}} F_{ij} - w_i\,, \qquad i \in \mathcal{I}\,,$$

which is physically interpreted as a *mass conservation law*: the rate of change of the mass in cell i equals the imbalance between the total inflow in it and the total outflow from it, the former coinciding with the sum of the external inflow u_i and the aggregate inflow from the other cells $\sum_{j\in\mathcal{I}} F_{ji}$, and the latter being given by the aggregate outflow towards other cells $\sum_{j\in\mathcal{I}} F_{ij}$ and the outflow towards the external environment w_i. An equivalent form of (1) that will often prove convenient is

$$\dot{x} = u - (I - R^T)z\,, \tag{4}$$

where

$$z = F\mathbb{1} + w$$

is a nonegative n-dimensional vector whose entries

$$z_i = \sum_{j\in\mathcal{I}} F_{ij} + w_i$$

represent the *total outflows* from the cells $i \in \mathcal{I}$, and $R \in \mathbb{R}^{n\times n}$ is a *routing matrix* whose entries R_{ij}, sometimes referred to as *split*

[1] Throughout, we will always assume that $(i, i) \notin \mathcal{A}$, so that $F_{ii} = 0$, for all $i \in \mathcal{I}$.