THE INSTITUTE FOR SYSTEMS RESEARCH

ISR TECHNICAL REPORT 2012-04

Flow Control in Time-Varying, Random Supply Chains

Ion Matei
Assane Gueye
John S. Baras

Flow Control in Time-Varying, Random Supply Chains

Ion Matei, Assane Gueye and John S. Baras

Abstract

Today's supply chains are more and more complex. They depend on a network of independent, yet interconnected moving parts. They rely on critical infrastructures and experience a lot of time variability and randomness. Designing strategies that deal with such constantly changing supply chains is necessary in this increasingly globalized economy where supply chain disruptions have impacts that propagate not only locally but also globally. In this paper we propose a randomized flow control algorithm for a time varying, random supply chain network. We formulate a constrained stochastic optimization problem that maximizes the profit function in terms of the long-run, time-average rates of the flows in the supply chain. We show that our algorithm, which is based on queueing theory and stochastic analysis concepts, can get arbitrarily close to the solution of the aforementioned optimization problem. In addition, we describe how the flow control algorithm can be extended to a multiple firms supply chain setup and present numerical simulations of our algorithm for different supply chain topologies.

I. Introduction

Among many possible definitions, the supply chain can be defined as the movement of materials as they flow from their source to the end customer. This includes activities such as purchasing, manufacturing, warehousing, transportation, customer service, demand planning, supply planning, and so on. As manufacturing is often outsourced around the world, with each component made in locations chosen for expertise and low costs [14], the nowadays supply chains are more and more complex, depending on a network of independent, yet interconnected moving parts and relying on critical infrastructures. In addition to the need for roads, railways, and airports to move goods, supply chains also need effective communications systems to transmit information between trading partners [13].

The global supply chain repeatedly demonstrates the co-existence of operational optimization with operational vulnerability [14]. This was most recently and dramatically demonstrated in the aftermath of several accidents and natural disasters. For example, a fire in the Phillips Semiconductor plant in Albuquerque, New Mexico, has caused its major customer, Ericsson, to lose 400 million in potential revenues. Another illustrative example concerns the impact of Hurricane Katrina, with the consequence that 10% - 15% of total U.S. gasoline production was halted, which not only raised the oil price in the U.S., but also overseas [3]. Moreover, the world price of coffee rose 22% after Hurricane Mitch struck the Central American republics of Nicaragua, Guatemala, and Honduras, which also affected supply chains worldwide [5]. More recently, the tragic earthquake of March 13, 2011 off the northeastern coast of Japan and the devastating tsunami which followed have shattered the nation, with immense loss of life, property, and uncertainty of the future, not the least of which is the expected decades long impact of the nuclear reactors in Fukushima [14].

Until recently, stakeholders have focused principally on operational optimization, but events such as the ones previously enumerated have commanded stakeholders to recognize operational vulnerabilities and underlined the *time-varying* and *random* nature of the supply chains.

Managing risks associated to critical infrastructure begins with understanding the types of failures that can disrupt commercial supply chains. Power outages occurs mainly due to storms or other natural events, but they can also occur due to equipment failure, surges, or demand exceeding capacity. Power outages can cause production downtimes, loss of raw materials or products, or potential damage to equipment or processes. Transportation failure, or failure to deliver a material or product on time, may occur due to equipment failure, departure delays, traffic, overcrowding or understaffing of ports, and numerous other reasons [13].

As economies around the world have become increasingly global, supply chain networks face many new types of risk, including natural disasters, political/social instability, cultural/communication inconsistency, exchange rate fluctuation, and local legislations [1]. Due to this increasingly globalized economy, supply chain disruptions have impacts that propagate not only locally but also globally. Hence, a holistic, system-wide approach to supply chain network modeling and analysis is essential in order to be able to capture the complex interactions among decision makers. Indeed, such rigorous modeling and analysis in the presence of possible disruptions is imperative since any incident may have lasting major financial consequences.

In this paper we focus on the logistics of the supply chain. Logistics is that part of supply chain management that plans, implements, and controls the efficient, effective forward and reverse flows as well as storage of goods and services from the point of origin to the point of consumption in order to meet customers requirements [2].

In the operation research literature a lot of attention was given to the management of flows in supply chains, with approaches mainly based on (stochastic) linear programming and game theory. The role of a supply chain, the key strategic drivers of its performance and the analytical methodologies for its analysis are extensively treated in [4]. In [12] a stochastic programming model and solution algorithm for solving supply chain network design problems are proposed, where the processing/transportation costs, demands, supplies, and capacities are stochastic parameters with known joint distribution. In [8] a survey of some applications of cooperative game theory to supply chain management is introduced. Special emphasis is placed on two important aspects of cooperative games: profit allocation and stability. More recently, in [10] the authors extend the supply chain research by capturing supply-side disruption risks, transportation and other cost risks, and demand-side uncertainty within an integrated modeling and robustness analysis framework.

In this paper we propose a randomized control algorithm for the flow of product in a time-varying, random supply chain aimed at maximizing the profit of a firm. The algorithm we propose is dynamic, it adapts to the current state of the supply chain and results as a solution of a stochastic optimization problem. The solution of the stochastic optimization problem is derived using a drift analysis technique. More importantly, the algorithm does not require knowledge of the probability distribution of the random process that drives the supply chain and deals with both supply changes and demand variability.

The paper is organized as following. In Section II we use a random graph to model the time-varying supply chain network. Section III introduces the notion of capacity region of a supply chain and formulates a constrained stochastic optimization problem, aimed at maximizing the profit function in terms of the long-run time-average of the flows' rates. In Section IV we describe a randomized, dynamic flow control algorithm for (approximately) solving the stochastic optimization problem, using queuing theory concepts to model the constraints. Section V presents a performance analysis of the flow control algorithm, which shows that the solution of our algorithm can get arbitrarily close to the solution of the optimization problem described in Section III. Section VI shows how our flow control algorithm can be extended to the case where multiple firms share the same market. We end the paper with numerical simulation of our algorithm for a variety of network topologies and some concluding remarks.

II. Supply chain model

We consider that a firm is involved in the production, storage and distribution of a homogeneous product. The firm is considering a set of manufacturing facilities, a set of warehouses and serves a set of retail outlets/demand markets.

Our supply chain model is similar to the one used in [9], with the additional assumption that the the network is *time-varying* and *random*. An example of a supply chain network is given in Figure 1, where node 1 represents the firm, nodes $\{2,3,4\}$ represent the set of manufacturing facilities, nodes $\{5,6\}$ are the warehouses and nodes $\{7,8,9\}$ designate the retail outlets/demand markets.

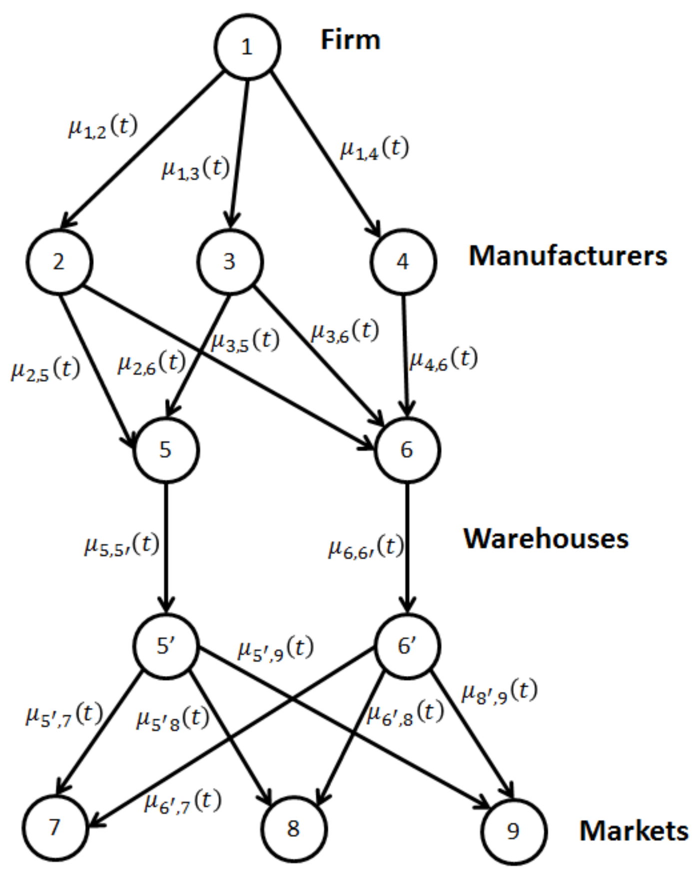

Fig. 1: Example of supply chain network

We first consider the supply chain with only one firm. Later, we will show how our approach can be extended to a multiple-firms scenario. The single-firm scenario is suitable for dominant-firm model, where a single firm controls a dominant share of the market [11]. Let us denote by $\mathcal{F}$, $\mathcal{M}$, $\mathcal{W}$ and $\mathcal{R}$ the sets of firms, manufacturers, warehouses and retailers, respectively, and let $\mathcal{N}$ be the set of all nodes in the network (with a typical node denoted by i), i.e., $\mathcal{N} = \{\mathcal{F} \cup \mathcal{M} \cup \mathcal{W} \cup \mathcal{R}\} \cup \{i' | i \in \mathcal{W}\}$, with cardinality $N = |\mathcal{N}|$. Note that similarly to [9], a warehouse i is represented by two nodes in the network (by using i' as well) in order to clearly emphasize the flow of product passing through the warehouse, i.e., through the link (i, i'). We let $\mathcal{L} = \{(i, j), i, j \in \mathcal{N}\}$ denote the set of links of the supply chain, where products "flow" from node i to node j for each $(i, j) \in \mathcal{L}$ and where the flow of product in the chain is driven by the demand at the retailers/markets (we assume that links of the form (i, i') are also included in $\mathcal{L}$).

We make the assumption that the supply chain operates in slotted time, with slots normalized to integral units so that slot times occur at times $t \in \{0, 1, 2, \ldots\}$. We denote by $S(t)$ the supply chain network state during slot t. The state process $S(t)$ reflects the uncontrollable conditions of the supply chain network such that possible disruptions in manufacturing and transportation due to natural disasters, power outages, technical malfunctions, etc. For example, the transport or manufacturing capacity can be at full capacity or at zero capacity in case of uncontrollable events. For simplicity, throughout the the rest of the paper, we assume that the links of the supply chain can be either active or inactive, as described by $S(t)$. This means that a transportation link may become unavailable at some time slot. We make the following assumption about the statistical properties of $S(t)$.

Assumption 2.1: The process $S(t)$ belongs to a finite set $\mathcal{S}$ and evolves according to an identically, independently distributed random process with stationary distribution given by $\pi = (\pi_s)$. Consequently

$$\lim_{t \to \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbf{1}_{\{S(\tau)=s\}} = \pi_s, \ \forall s \in \mathcal{S}, \tag{1}$$

where $\mathbf{1}_{\{S(\tau)=s\}} = \pi_s$ is the indicator function that takes value 1 whenever $S(t) = s$, and 0 otherwise.

We denote by $\mu_{i,j}(t)$ the amount of product flowing through the link (i, j) during time slot t. Without loss of generality we assume that the flows are measured in (final) product units; in order to recover other units (raw materials for example) the flows are multiplied by the process rate of the economic unit generating the flow). We denote by $d_i(t)$ the demand at market i at time slot t, assumed a random process. It is reasonable to assume that the quantity of product flowing between different entities is upper-bounded, and hence we make the following assumption.

Assumption 2.2: The flows $\mu_{i,j}(t)$ are positive for all time-slots t and there exist positive scalars μ_i^{max} such that

$$\sum_b \mu_{i,b}(t) \leq \mu_i^{max}, \ \forall i \in \mathcal{N}, \ \forall t, \tag{2}$$

where all pairs (i,b) belong to the set $\mathcal{L}$.

The above inequalities can be thought as production, transportation or storage capabilities limitations.

The following definitions introduce the time averages of the product flows in the supply chain.

Definition 2.1: The time average flows of product in the supply chain are given by

$$\bar{\mu}_{i,j}(t) = \frac{1}{t}\sum_{\tau=0}^{t-1} E\{\mu_{i,j}(\tau)\}, \tag{3}$$

and the long-run time averages of flow product are given by

$$\bar{\mu}_{i,j} = \lim_{t\to\infty} \bar{\mu}_{i,j}(t), \tag{4}$$

for all $(i,j) \in \mathcal{L}$.

Additionally, the market demands satisfy the following assumption.

Assumption 2.3: The random processes $d_i(t)$ are independent and identically distributed with mean given by

$$\bar{d}_i = E\{d_i(t)\}, \ \forall i \in \mathcal{R}. \tag{5}$$

Let us also define the aggregate vectors of product flows $\boldsymbol{\mu}(t) = (\mu_{i,j}(t), (i,j) \in \mathcal{L})$, and market demands $\boldsymbol{d}(t) = (d_i(t), \ i \in \mathcal{R})$.

III. Formulation of the stochastic optimization problem

We start this section by defining the capacity region of a supply chain network, which tells us how much demand the supply chain can support.

Definition 3.1: The capacity region Λ of a supply chain is the closure of all vector of demands $\boldsymbol{x} = (x_i)$ that can be supported by the supply chain network, considering all possible strategies for choosing the flows of product, under the limitations introduced by Assumption 2.2.

Let $C_{i,j}(s)$ be the set of flows on link (i,j) satisfying Assumption 2.2 when the supply network is in state s, and under all possible flow control policies. Let $\boldsymbol{C}(s)$ be the set of all link sets, i.e.,

$\boldsymbol{C}(s) = (C_{i,j}(s))_{(i,j)\in\mathcal{L}}$. To further characterize the capacity region of a supply chain network, we introduce the family of graphs Γ, given by

$$\Gamma \triangleq \sum_{s\in\mathcal{S}} \pi_s co\{\boldsymbol{C}(s)\},$$

where $co\{\boldsymbol{C}(s)\}$ represents the convex hull of the set $\boldsymbol{C}(s)$. We say that a matrix $\boldsymbol{G} = (G_{i,j})$ belongs to Γ if there exits a randomized flow control policy that depends on the state of the network, such that

$$\boldsymbol{G} = \sum_{s\in\mathcal{S}} \pi_s E\{\boldsymbol{\mu}(t)|S(t) = s\},$$

where $E\{\boldsymbol{\mu}(t)|S(t) = s\}$ is the expected flow matrix under the considered policy, given that the supply chain is in state s. The following Theorem inspired by [6] describes the capacity region of the supply chain.

Theorem 3.1: The capacity region of a supply chain is given by the set Λ of all demand vectors $\boldsymbol{x} = (x_i)$ such that there exits a flow matrix $\boldsymbol{G} = (G_{i,j})$ belonging to the closure of Γ, together with flow variables $f_{i,j}$ such that

$$f_{i,j} \geq 0,\ \forall (i,j) \in \mathcal{L},\ f_{i,j} = 0,\ \forall (i,j) \notin \mathcal{L}, \tag{6}$$

$$\sum_{a\in\mathcal{F}} f_{a,i} = \sum_{b\in\mathcal{W}} f_{i,b},\ \forall i \in \mathcal{M}, \tag{7}$$

$$\sum_{a\in\mathcal{M}} f_{a,i} = f_{i,i'},\ \forall i \in \mathcal{W}, \tag{8}$$

$$f_{i,i'} = \sum_{b\in\mathcal{R}} f_{i',b},\ \forall i \in \mathcal{W}, \tag{9}$$

$$\sum_{a\in\mathcal{W}} f_{a',i} = x_i,\ \forall i \in \mathcal{R}, \tag{10}$$

$$f_{i,j} \leq G_{i,j},\ \forall (i,j) \in \mathcal{L}. \tag{11}$$

In the particular case where the process $S(t)$ is i.i.d. (which is our assumption throughout this paper), the next Corollary presents a further characterization of the capacity region, where we use $Cl(A)$ to denote the closure of the set A.

Corollary 3.1 (adaptation of Corollary 3.9, [6]): If $\Gamma = Cl(\Gamma)$ and if the state process $S(t)$ is i.i.d. from slot to slot, the demand vector $\boldsymbol{x}$ is within the capacity region Λ if and only if there exists a stationary (randomized) policy that chooses $\boldsymbol{\mu}(t)$ based only on the current topology state $S(t)$, such that

$$E\left\{\sum_{a\in\mathcal{F}}\mu_{a,i}(t)\right\} = E\left\{\sum_{b\in\mathcal{W}}\mu_{i,b}(t)\right\}, \forall\ i\in\mathcal{M},$$

$$E\left\{\sum_{a\in\mathcal{M}}\mu_{a,i}(t)\right\} = E\left\{\mu_{i,i'}(t)\right\}, \forall\ i\in\mathcal{W},$$

$$E\left\{\mu_{i,i'}(t)\right\} = E\left\{\sum_{b\in\mathcal{R}}\mu_{i',b}(t)\right\}, \forall\ i\in\mathcal{W},$$

$$E\left\{\sum_{a\in\mathcal{W}}\mu_{a',i}(t)\right\} = x_i, \forall\ i\in\mathcal{R},$$

where the expectation is taken with respect to the random process $S(t)$ and the (potentially) random policy based on $S(t)$.

Note that if $\boldsymbol{x}\in\Lambda$, then any $\tilde{\boldsymbol{x}}$ such that $\tilde{\boldsymbol{x}}\le\boldsymbol{x}$ entrywise, also belongs to Λ. In addition, it can be shown that the set Λ is convex, closed and bounded and it contains the vector of all zeros, (i.e., $\mathbf{0}\in\Lambda$).

Next, we formulate a stochastic optimization problem that describes the objective of the firm under a set of constraints induced by the supply chain. The goal of the firm is to maximize its profit, that is the difference between the revenue and the cost functions. We consider the revenue function of the firm, that depends on the quantity of products that reach the retailers/markets in the long-run. We denote that by

$$\boldsymbol{f}(\bar{\boldsymbol{\mu}}) = \sum_{i\in\mathcal{W},j\in\mathcal{R}} \boldsymbol{f}_{i',j}(\bar{\mu}_{i',j}),$$

where (i',j) represent valid warehouse-retailers pairs, i.e., $i\in\mathcal{W}$, $j\in\mathcal{R}$ and $(i',j)\in\mathcal{L}$. We also consider cost functions associated with each link $(i,j)\in\mathcal{L}$ which we denote by $g_{i,j}(\bar{\mu}_{i,j})$. These cost functions depend on the flow of product on the links and are generated by activities such as acquiring raw materials, manufacturing, transportation or warehouse usage. The total cost function is given by

$$\boldsymbol{g}(\bar{\boldsymbol{\mu}}) = \sum_{i\in\mathcal{F},j\in\mathcal{M}} \boldsymbol{g}_{i,j}(\bar{\mu}_{i,j}) + \sum_{i\in\mathcal{M},j\in\mathcal{W}} \boldsymbol{g}_{i,j}(\bar{\mu}_{i,j}) + \sum_{i\in\mathcal{W}} \boldsymbol{g}_{i,i'}(\bar{\mu}_{i,i'}) +$$
$$+ \sum_{i\in\mathcal{W},j\in\mathcal{R}} \boldsymbol{g}_{i',j}(\bar{\mu}_{i',j}).$$

Assumption 3.1: The functions $\boldsymbol{f}_{i,j}$ are non-negative, continuously differentiable and concave, while the functions $\boldsymbol{g}_{i,j}$ are non-negative, continuously differentiable and convex.

We define the profit function $\boldsymbol{h}$ as the difference between the revenue and the cost functions, i.e.,

$$\boldsymbol{h}(\bar{\boldsymbol{\mu}}) = \boldsymbol{f}(\bar{\boldsymbol{\mu}}) - \boldsymbol{g}(\bar{\boldsymbol{\mu}}).$$

The firm's objective is to maximize the profit under the flow constraints induced by the (capacity region of the) supply chain network. Let x_i denote the long-run, average flow of product arriving at market (retailer) i, that is

$$x_i = \sum_{a \in \mathcal{W}} \bar{\mu}_{a,i}, \ \ \forall i \in \mathcal{R}.$$

We formulate the following stochastic optimization problem:

$$\begin{aligned} \max_{\bar{\boldsymbol{\mu}}, \boldsymbol{x}} \quad & \boldsymbol{h}(\bar{\boldsymbol{\mu}}) \\ \text{subject to:} \quad & \boldsymbol{x} \in \Lambda, \\ & \boldsymbol{x} \leq \bar{\boldsymbol{d}}. \end{aligned} \tag{12}$$

The first constraint introduced above ensures that the average product flows arriving at the markets (retailers) are within the capacity region of the supply chain network, i.e., can be supported by the network. The second inequality ensures that the long term flow of product arriving at the markets are no larger than the demands at the markets.

By Corollary 3.1, we can equivalently write the above stochastic optimization problem as,

$$\begin{aligned} \max_{\bar{\boldsymbol{\mu}}} \quad & \boldsymbol{h}(\bar{\boldsymbol{\mu}}) \\ \text{subject to:} \quad & \textstyle\sum_{a \in \mathcal{F}} \bar{\mu}_{a,i} = \sum_{b \in \mathcal{W}} \bar{\mu}_{i,b}, \forall i \in \mathcal{M}, \\ & \textstyle\sum_{a \in \mathcal{M}} \bar{\mu}_{a,i} = \bar{\mu}_{i,i'}, \forall i \in \mathcal{W}, \\ & \bar{\mu}_{i',i} = \textstyle\sum_{b \in \mathcal{R}} \bar{\mu}_{i',b}, \forall i \in \mathcal{W}, \\ & \textstyle\sum_{a \in \mathcal{W}} \bar{\mu}_{a',i} \leq \bar{d}_i, \forall i \in \mathcal{R}, \end{aligned} \tag{13}$$

where $\bar{\mu}_{i,j} = E\{\mu_{i,j}(t)\}$ for all $(i, j) \in \mathcal{L}$, with $\mu_{i,j}(t)$ being chosen by some stationary, randomized control algorithm, based only on the current state $S(t)$.

Assumption 3.2 (Interior point): There exist positive scalars ϵ_1 and ϵ_2 and two stationary randomized flow control policies based on the current state $S(t)$, corresponding to ϵ_1 and ϵ_2, respectively, such that

$$E\{\mu_{1,i}^{\epsilon_1}(t)\} + \epsilon_1 = E\left\{\sum_b \mu_{i,b}^{\epsilon_1}(t)\right\}, \forall\ i \in \mathcal{M},$$

$$\sum_a E\left\{\mu_{a,i}^{\epsilon_1}(t)\right\} + \epsilon_1 = E\left\{\mu_{i,i'}^{\epsilon_1}(t)\right\}, \forall\ i \in \mathcal{W},$$

$$E\left\{\mu_{i,i'}^{\epsilon_1}(t)\right\} + \epsilon_1 = E\left\{\sum_b \mu_{i',b}^{\epsilon_1}(t)\right\}, \forall\ i \in \mathcal{W},$$

$$E\left\{\sum_a \mu_{a',i}^{\epsilon_1}(t)\right\} + \epsilon_1 \le \bar{d}_i, \forall\ i \in \mathcal{R},$$

and

$$E\left\{\sum_b \mu_{i,b}^{\epsilon_2}(t)\right\} + \epsilon_2 = E\{\mu_{1,i}^{\epsilon_2}(t)\}, \forall\ i \in \mathcal{M},$$

$$E\left\{\mu_{i,i'}^{\epsilon_2}(t)\right\} + \epsilon_2 = \sum_a E\left\{\mu_{a,i}^{\epsilon_2}(t)\right\}, \forall\ i \in \mathcal{W},$$

$$E\left\{\sum_b \mu_{i',b}^{\epsilon_2}(t)\right\} + \epsilon_2 = E\left\{\mu_{i,i'}^{\epsilon_2}(t)\right\}, \forall\ i \in \mathcal{W},$$

$$E\left\{\sum_a \mu_{a',i}^{\epsilon_2}(t)\right\} \le \bar{d}_i, \forall\ i \in \mathcal{R},$$

In the above Assumption, ϵ_1 can be viewed as an additional flow on one of the links that arrives at a node and is produced by a source outside the supply chain, while ϵ_2 can be viewed as an additional flow leaving a node on one of the links but that fails to reach the destination node.

From the optimization point of view, it is most of the time more advantageous to work with inequality constraints rather than equality constraints. As a consequence, we replace each of the equality constraints in (13) by two inequality constraints, as shown in the following:

$$\begin{aligned} \max_{\bar{\mu}} \quad & \boldsymbol{h}(\bar{\boldsymbol{\mu}}) \\ \text{subject to:} \quad & \textstyle\sum_{a\in\mathcal{F}} \bar{\mu}_{a,i} \leq \sum_{b\in\mathcal{W}} \bar{\mu}_{i,b}, \forall i \in \mathcal{M}, \\ & \textstyle\sum_{a\in\mathcal{F}} \bar{\mu}_{a,i} \geq \sum_{b\in\mathcal{W}} \bar{\mu}_{i,b}, \forall i \in \mathcal{M}, \\ & \textstyle\sum_{a\in\mathcal{M}} \bar{\mu}_{a,i} \leq \bar{\mu}_{i,i'}, \forall i \in \mathcal{W}, \\ & \textstyle\sum_{a\in\mathcal{M}} \bar{\mu}_{a,i} \geq \bar{\mu}_{i,i'}, \forall i \in \mathcal{W}, \\ & \textstyle\bar{\mu}_{i',i} \leq \sum_{b\in\mathcal{R}} \bar{\mu}_{i',b}, \forall i \in \mathcal{W}, \\ & \textstyle\bar{\mu}_{i',i} \geq \sum_{b\in\mathcal{R}} \bar{\mu}_{i',b}, \forall i \in \mathcal{W}, \\ & \textstyle\sum_{a\in\mathcal{W}} \bar{\mu}_{a',i} \leq \bar{d}_i, \forall i \in \mathcal{R}. \end{aligned} \tag{14}$$

In the following sections we present a mathematical approach for solving the optimization problem (14), based on queueing theory and on drift analysis. We start with the case where the supply chain corresponds to only one firm, followed by a generalization for the multiple-firms case.

IV. Flow control algorithm

In this section we introduce a flow control algorithm which ensures that the long-run, time-average flows in the supply chain can get arbitrarily close to the optimal solution of (13). Our strategy is to associate to each of the inequality constraints a (virtual) queue. We show that the inequality constraints are satisfied if the queues associated to them are stable, in some sense that is about to be defined. By taking advantage of this property, we propose an algorithm that stabilizes the queues and gets arbitrarily close to the optimal solution of (14). The algorithm is derived as a result of a *drift analysis* approach on the (virtual) queues. This approach is closely related to the stochastic Lyapunov theory [7].

A. Modeling inequality constraints using queues

In this subsection we show why we can connect the feasibility of the inequality constraints defined in our optimization problem to the stability of a set of queues associated to them.

Consider a queue $U(t)$ (Figure 2) with (possibly random) input $\lambda(t)$ and output $\mu(t)$, whose dynamics is given by

$$U(t+1) = \max\{U(t) - \mu(t), 0\} + \lambda(t).$$

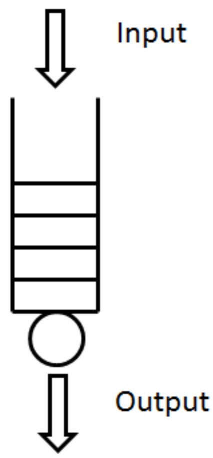

Fig. 2: Queue schematics

Definition 4.1: We say that the queue $U(t)$ is *strongly stable* if

$$\limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} E\{U(\tau)\} < \infty.$$

Let us now assume that there exists $\bar{\lambda}$ and $\bar{\mu}$ such that

$$\bar{\lambda} = \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} E[\lambda(\tau)],$$

and

$$\bar{\mu} = \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} E[\mu(\tau)].$$

Proposition 4.1 (Queue stability): A necessary condition for the strong stability of the queue $U(t)$ is

$$\bar{\lambda} \leq \bar{\mu}.$$

The necessary condition is quite intuitive. Indeed, if $\bar{\lambda} > \bar{\mu}$, the the expected queue backlog grows to infinity, leading to instability. Under additional assumptions on the processes $\lambda(t)$ and $\mu(t)$, it can be shown that $\bar{\lambda} < \bar{\mu}$ is also a sufficient condition (see [6] for more details).

Let us now consider a set of (virtual) queues associated with the constraints of our optimization problem, whose dynamics are given in the following.

In the case of a manufacturing unit, the dynamics of the queue levels are given by

$$U_i^1(t+1) = \max\left\{U_i^1(t) - \sum_b \mu_{i,b}(t), 0\right\} + \sum_a \mu_{a,i}(t), \forall i \in \mathcal{M}, \tag{15}$$

$$U_i^2(t+1) = \max\left\{U_i^2(t) - \sum_a \mu_{a,i}(t), 0\right\} + \sum_b \mu_{i,b}(t), \forall i \in \mathcal{M}. \tag{16}$$

In the case of the warehouses we have two sets of dynamic equations associated with a warehouse:

$$U_i^1(t+1) = \max\left\{U_i^1(t) - \mu_{i,i'}(t), 0\right\} + \sum_a \mu_{a,i}(t), \forall i \in \mathcal{W}, \tag{17}$$

$$U_i^2(t+1) = \max\left\{U_i^2(t) - \sum_a \mu_{a,i}(t), 0\right\} + \mu_{i,i'}(t), \forall i \in \mathcal{W}. \tag{18}$$

and

$$U_{i'}^1(t+1) = \max\left\{U_{i'}^1(t) - \sum_b \mu_{i',b}(t), 0\right\} + \mu_{i,i'}(t), \forall i \in \mathcal{W}, \tag{19}$$

$$U_{i'}^2(t+1) = \max\left\{U_{i'}^2(t) - \mu_{i,i'}(t), 0\right\} + \sum_b \mu_{i',b}(t), \forall i \in \mathcal{W}. \tag{20}$$

In the case of the retailers, we have

$$U_i(t+1) = \max\left\{U_i^1(t) - d_i(t), 0\right\} + \sum_a \mu_{a,i}(t), \forall i \in \mathcal{R}. \tag{21}$$

Remark 4.1: In the previous expressions, by $\sum_b \mu_{i,b}(t)$ we understand the summation over all active links carrying products from node i, at time slot t, as per the state of the supply chain state $S(t)$. A similar interpretation can be given to the term $\sum_a \mu_{a,i}(t)$.

From Proposition 4.1, we can infer that any flow control algorithm that *stabilizes* the queues will in fact *satisfy* the flow constraints defined in the optimization problem (13). Therefore, it make sense to look for an algorithm that stabilizes the queues defined above and in the same time maximizes the profit function.

B. Algorithm description

In this section we introduce a randomized flow control algorithm which can get arbitrarily close to the optimal solution of (13). The algorithm stabilizes the (virtual) queues and therefore ensures that the inequality constraints are satisfied, but most importantly shows how the economic entities in the supply chain dynamically adapt their flows based on the changes in the network.

The algorithm consists of actions taken by the entities involved in the economic activities of the firm, at each time slot t. Let δ be a positive scalar, that affects the performance of the algorithm. For simplicity, we assume that the set of firms $\mathcal{F}$ contains only one firm, say node 1 in the network. In the following we describe the flow control algorithm.

- *Control of the raw material flow:* At every time slot, the firm observes the current levels of the manufacturers' queues, $U_b^1(t)$ and $U_b^2(t)$. Then, at each time t it chooses the amount $\mu_{1,b}$ of raw material sent to manufacturer b, where $\mu_{1,b}$ is the solution of the following optimization problem:

$$\min_{\mu_{1,b}} \quad \textstyle\sum_{b\in\mathcal{M}}\Big(\delta \boldsymbol{g}_{1,b}(\mu_{1,b}) + \Big[U_b^1(t) - U_b^2(t)\Big]\mu_{1,b}\Big) \tag{22}$$

$$\text{subject to:} \qquad \textstyle\sum_{b\in\mathcal{M}}\mu_{1,b} \le \mu_1^{max}, \mu_{1,b} \ge 0, \forall b. \tag{23}$$

- *Control of the flow of product from manufacturers to warehouses:* At every time slot, each manufacturer i observes the current level of its queues $U_i^1(t)$ and $U_i^2(t)$ and the current levels of the queues of the warehouse b to which product is possible to be sent to (as per the state of $S(t)$), i.e., $U_b^1(t)$ and $U_b^2(t)$. The amount of product sent to each warehouse b at time slot t is given by $\mu_{i,b}$, obtained as solution of the following optimization problem:

$$\min_{\mu_{i,b}} \quad \textstyle\sum_b \delta \boldsymbol{g}_{i,b}(\mu_{i,b}) - \Big(\Big[U_i^1(t) - U_b^1(t)\Big] + \Big[U_b^2(t) - U_i^2(t)\Big]\Big)\mu_{i,b} \tag{24}$$

$$\text{subject to:} \qquad \textstyle\sum_{b\in\mathcal{W}}\mu_{i,b} \le \mu_i^{max},\ \mu_{i,b} \ge 0, \forall b, \tag{25}$$

  for all $i \in \mathcal{M}$, $b \in \mathcal{W}$ and $(i,b) \in \mathcal{L}$ which are active at time t, as per the state of the supply chain given by $S(t)$.

- *Control of the flow of product within the warehouses:* At every time slot, each warehouse i observes the current level of its queues $U_i^1(t)$, $U_{i'}^1(t)$, $U_i^2(t)$ and $U_{i'}^2(t)$. The amount of product allowed in the warehouse at time slot t is given by $\mu_{i,i'}$, obtained as solution of the following optimization problem:

$$\min_{\mu} \quad \delta \boldsymbol{g}_{i,i'}(\mu) - \Big(\Big[U_i^1(t) - U_{i'}^1(t)\Big] + \Big[U_{i'}^2(t) - U_i^2(t)\Big]\Big)\mu \tag{26}$$

$$\text{subject to:} \qquad 0 \le \mu \le \mu_i^{max} \tag{27}$$

for all $i \in \mathcal{W}$ and $(i,i') \in \mathcal{L}$ which are active at time t, as per the state of the supply chain given by $S(t)$.

- *Control of the flow of product from warehouses to retailers:* At every time slot, each warehouse i observes the current level of its queues backlog $U^1_{i'}(t)$ and $U^2_{i'}(t)$ and the current level of the queue of the retailer b to which the product is sent to, i.e., $U^1_b(t)$. The amount of product sent to retailer b at time slot t is given by $\mu_{i',b}$, where $\mu_{i',b}$ are obtained as solution of the following optimization problem:

$$\min_{\mu_{i',b}} \quad \sum_{b\in\mathcal{R}} \delta \boldsymbol{g}_{i',b}(\mu_{i',b}) - \delta \boldsymbol{f}_{i',b}(\mu_{i',b}) - \left[\left(U^1_{i'}(t) - U^1_b(t)\right) - U^2_{i'}(t)\right]\mu_{i',b} \tag{28}$$

$$\text{subject to:} \quad \sum_{b\in\mathcal{R}} \mu_{i',b} \le \mu^{max}_{i'}, \mu_{i',b} \ge 0, \ \forall b, \tag{29}$$

for all $i \in \mathcal{W}$, $b \in \mathcal{R}$ and $(i',b) \in \mathcal{L}$ which are active at time t, as per the state of the supply chain given by $S(t)$.

Note that the optimization problems (22)-(28) are convex constrained optimization problems, which can be solved efficiently at each time slot. Also, note that each of the entities involved in the economic activities does not need to know the entire state of the network, nor the probability distribution of $S(t)$. Indeed, in the case of the manufacturers, the raw material flow is determined only by the level of the queues' backlogs and the cost. When a manufacturer must decide the flow of product sent to warehouses, it looks at the current valid links, and it makes the decision based on the cost of utilizing the respective links and based on the difference between the queues' levels of the manufacturer and warehouses. In the case of the amount of product allowed in a warehouse, the decision is based on the cost of keeping the product in the warehouse and on the difference between the levels of the (virtual) queues. Finally, the amount of product sent to retailers from a warehouse is based on the current available links, on the (localized) profit obtained from sending products to a specific retailer and on the difference between the queues' levels of the warehouse and retailers. This limited need of information for implementing the algorithm makes it advantageous for controlling the flow of product in more and more complex and globalized supply chains. Another important observation is that the manufacturers, warehouses and retailers do not need to know the entire state of the network at a time slot, nor the statistics of the state process $S(t)$. They only need to observe the state of links which connect them to their neighbors. In addition, the virtual queues $U^1_i(t)$ can find an analogy in reality. Indeed, in the case of a manufacturer for example, the queue can be viewed as a deposit for the raw material waiting to be processed.

V. Derivation of the algorithm and performance analysis

In this section we show the considerations behind the development of the algorithm and analyze its performance. We show that the long-run time averages of the flows generated by this algorithm are feasible, and that they can get arbitrarily close to the optimal solution.

A. Derivation of the algorithm

Let $\boldsymbol{U}(t) = \left(U_i^j(t), i \in \mathcal{M}, U_i^j(t), U_{i'}^j(t), i \in \mathcal{W}, j \in \{1,2\}, U_i(t), i \in \mathcal{R}\right)$ be the vector of queues. We define the following quadratic Lyapunov function

$$V(\boldsymbol{U}(t)) \triangleq \frac{1}{2} \sum_{j\in\{1,2\}} \left[\sum_{i\in\mathcal{M}} U_i^j(t)^2 + \sum_{i\in\mathcal{W}} \left(U_i^j(t)^2 + U_{i'}^j(t)^2 \right) \right] + \frac{1}{2} \sum_{i\in\mathcal{R}} U_i(t)^2.$$

and introduce the queues' *drift*:

$$\Delta(\boldsymbol{U}(t)) \triangleq E[V(\boldsymbol{U}(t+1)) - V(\boldsymbol{U}(t)) | \boldsymbol{U}(t)].$$

The flow control algorithm for the supply chain results from minimizing an upper bound of the following quantity

$$\Delta(\boldsymbol{U}(t)) - \delta E\{\boldsymbol{h}(\boldsymbol{\mu}(t)) | \boldsymbol{U}(t)\}, \tag{30}$$

for each time slot t. Note that minimizing the previous expression means a trade-off between the stability of the queues through the Lyapunov drift $\Delta(\boldsymbol{U}(t))$ and the firm's profit through the profit function $\boldsymbol{h}$, where δ is a weighing factor. In fact, making δ large enough implies focusing on maximizing the profit (and getting arbitrarily close to the optimal solution), but at a cost in terms of an increased product congestion in the queues.

Let us consider the nonnegative reals Y, U, μ, A such that

$$Y \leq \max\{U - \mu, 0\} + A.$$

It is not difficult to show that the following inequality holds:

$$Y^2 \leq U^2 + \mu^2 + A^2 - 2U(\mu - A). \tag{31}$$

Using the previous inequality, an upper-bound for (30) is as follows:

$$\Delta(\boldsymbol{U}(t)) - \delta E\{\boldsymbol{h}(\boldsymbol{\mu}(t)) | \boldsymbol{U}(t)\} \leq B\bar{N} - E\left\{ \sum_{i\in\mathcal{M}} U_i^1(t) \left(\sum_b \mu_{i,b}(t) - \mu_{1,i}(t) \right) | \boldsymbol{U}(t) \right\} -$$

$$-E\left\{\sum_{i\in\mathcal{M}} U_i^2(t)\left(-\sum_b \mu_{i,b}(t)+\mu_{1,i}(t)\right)|\boldsymbol{U}(t)\right\}-E\left\{\sum_{i\in\mathcal{W}} U_i^1(t)\left(\mu_{i,i'}(t)-\sum_a \mu_{a,i}(t)\right)|\boldsymbol{U}(t)\right\}-$$

$$-E\left\{\sum_{i\in\mathcal{W}} U_i^2(t)\left(-\mu_{i,i'}(t)+\sum_a \mu_{a,i}(t)\right)|\boldsymbol{U}(t)\right\}-E\left\{\sum_{i\in\mathcal{W}} U_{i'}^1(t)\left(\sum_b \mu_{i,b}(t)-\mu_{i,i'}(t)\right)|\boldsymbol{U}(t)\right\}-$$

$$-E\left\{\sum_{i\in\mathcal{W}} U_{i'}^2(t)\left(-\sum_b \mu_{i,b}(t)+\mu_{i,i'}(t)\right)|\boldsymbol{U}(t)\right\}-E\left\{\sum_{i\in\mathcal{R}} U_i(t)\left(d_i(t)-\sum_a \mu_{a',i}(t)\right)|\boldsymbol{U}(t)\right\}-$$

$$-\delta E\left\{\sum_{(i',j)} \boldsymbol{f}_{i',j}(\mu_{i',j}(t))|\boldsymbol{U}(t)\right\}+\delta E\left\{\sum_{i\in\mathcal{M}} \boldsymbol{g}_i(r_i(t))|\boldsymbol{U}(t)\right\}+\delta E\left\{\sum_{(i,j)} \boldsymbol{g}_{i,j}(\mu_{i,j}(t))|\boldsymbol{U}(t)\right\}+$$

$$+\delta E\left\{\sum_{(i,i')} \boldsymbol{g}_{i,i'}(\mu_{i,i'}(t))|\boldsymbol{U}(t)\right\}+\delta E\left\{\sum_{(i',j)} \boldsymbol{g}_{i',j}(\mu_{i',j}(t))|\boldsymbol{U}(t)\right\},$$

where

$$B \triangleq \frac{1}{\bar{N}}\sum_{i\in\mathcal{N}} 2\left(\mu_i^{max}\right)^2,$$

and where $\bar{N}$ is the number of all queues.

Rearranging the sums in the previous inequality, we can further write

$$\Delta(\boldsymbol{U}(t))-\delta E\{\boldsymbol{h}(\boldsymbol{\mu}(t))|\boldsymbol{U}(t)\}\leq$$

$$\leq BN+E\left\{\sum_{i\in\mathcal{R}} U_i(t)d_i(t)|\boldsymbol{U}(t)\right\}+E\left\{\sum_{i\in\mathcal{M}} \delta\boldsymbol{g}_{1,i}(\mu_{1,i}(t))+\left[U_i^1(t)-U_i^2(t)\right]\mu_{1,i}(t)|\boldsymbol{U}(t)\right\}+$$

$$E\left\{\sum_{(i,b),i\in\mathcal{M},b\in\mathcal{W}} \delta\boldsymbol{g}_{i,b}(\mu_{i,b}(t))-\left(\left[U_i^1(t)-U_b^1(t)\right]+\left[U_b^2(t)-U_i^2(t)\right]\right)\mu_{i,b}(t)|\boldsymbol{U}(t)\right\}+$$

$$+E\left\{\sum_{(i\in\mathcal{W}} \delta\boldsymbol{g}_{i,i'}(\mu_{i,i'}(t))-\left(\left[U_i^1(t)-U_{i'}^1(t)\right]+\left[U_{i'}^2(t)-U_i^2(t)\right]\right)\mu_{i,i'}(t)|\boldsymbol{U}(t)\right\}+$$

$$+E\left\{\sum_{(i',b),i\in\mathcal{W},b\in\mathcal{R}} \delta\boldsymbol{g}_{i',b}(\mu_{i',b}(t))-\delta\boldsymbol{f}_{i',b}(\mu_{i',b}(t))-\left[\left(U_{i'}^1(t)-U_b(t)\right)-U_{i'}^2(t)\right]\mu_{i',b}(t)|\boldsymbol{U}(t)\right\}. \quad (32)$$

From the above inequality, the derivation of the algorithm is evident. Given queue levels $\boldsymbol{U}(t)$, the flow control algorithm follows from greedily minimizing the right-hand side of the inequality (32), in terms of the control variables $\boldsymbol{\mu}(t)$ over all possible flow options satisfying the constraints introduced in Assumption 2.2.

B. Performance analysis

In this section we show that the dynamic flow control algorithm introduced above gets arbitrarily close to the optimal solution of (14). We start our analysis by stating the following theorem.

Theorem 5.1: Let Assumptions 2.1 through 3.2 hold and assume that there exist positive constants δ, ϵ and B such that for all timeslots t and all backlog queue levels $\boldsymbol{U}(t)$, the Lyapunov drift satisfies:

$$\boldsymbol{\Delta}(\boldsymbol{U}(t)) - \delta E\{\boldsymbol{h}(\boldsymbol{\mu}(t))|\boldsymbol{U}(t)\} \leq B - \epsilon \sum_{i=1}^{\bar{N}} U_i(t) - \delta \boldsymbol{h}^*, \tag{33}$$

where $\boldsymbol{h}^*$ is the optimal cost function of the stochastic optimization problem (13). Then the follwing inequalities are satisfied

$$\limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \left[\sum_{j=1}^{2} \left(\sum_{i\in\mathcal{W}} E\{U_i^j(\tau)\} + \sum_{i\in\mathcal{M}} E\{U_i^j(\tau) + U_{i'}^j(\tau)\} + \sum_{i\in\mathcal{R}} E\{U_i(\tau)\} \right) \right] \leq \frac{B + \delta(\bar{\boldsymbol{h}} - \boldsymbol{h}^*)}{\epsilon} \tag{34}$$

$$\liminf_{t\to\infty} \boldsymbol{h}(\bar{\boldsymbol{\mu}}(t)) \geq \boldsymbol{h}^* - \frac{B}{\delta}, \tag{35}$$

where $\bar{\boldsymbol{\mu}}(t)$ was defined in (3) and $\bar{\boldsymbol{h}}$ is given by

$$\bar{\boldsymbol{h}} \triangleq \limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} E\{\boldsymbol{h}(\boldsymbol{\mu}(\tau))\}.$$

The previous Theorem is a slight modification of Theorem 5.4 in [6] and for brevity the proof is omitted.

Remark 5.1: Note that since the flows $\mu_{i,j}(t)$ are upper bounded by μ_i^{max} and the function $\boldsymbol{h}$ is continuous, there exists $\boldsymbol{h}_{max}$ such that $\bar{\boldsymbol{h}} - \boldsymbol{h}^* \leq \boldsymbol{h}_{max}$. In addition, let us define $\mu_{max} \triangleq \min\{\mu_i^{max}\}$.

The next Theorem describes the performance of the flow control algorithm.

Theorem 5.2: Let Assumptions 2.1 through 3.2 hold. For any positive parameter δ the flow control algorithm stabilizes the (virtual) queues associated with the constraints of the optimization problem (14) and gives the following upper bounds:

$$\limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \left(\sum_{j=1}^{2} \sum_{i\in\mathcal{W}} E\{U_i^j(\tau)\} + \sum_{i\in\mathcal{M}} E\{U_i^j(\tau) + U_{i'}^j(\tau)\} + \sum_{i\in\mathcal{R}} E\{U_i(\tau) \right) \leq \frac{NB + \delta \boldsymbol{h}_{max}}{\mu_{max}} \tag{36}$$

$$\liminf_{t\to\infty} \boldsymbol{h}(\bar{\boldsymbol{\mu}}(t)) \geq \boldsymbol{h}(\boldsymbol{\mu}^*) - \frac{BN}{\delta}, \tag{37}$$

where $\boldsymbol{\mu}^*$ is the solution of (13) and where $\bar{\boldsymbol{\mu}}(t)$ satisfies (3).

Proof: Let ϵ_1 be a small quantity of product flow added to the inputs of queues $U_i^1(t)$ for all $i \in \mathcal{M} \cup \mathcal{W}$ and queues $U_i(t)$, for $i \in \mathcal{R}$. It follows that the dynamics of the aforementioned queues become

$$\begin{aligned}
U_i^1(t+1) &= \max\left\{U_i^1(t) - \sum_b \mu_{i,b}(t), 0\right\} + \mu_{1,i}(t) + \epsilon_1, \forall i \in \mathcal{M}, \\
U_i^1(t+1) &= \max\left\{U_i^1(t) - \mu_{i,i'}(t), 0\right\} + \sum_a \mu_{a,i}(t) + \epsilon_1, \forall i \in \mathcal{W}, \\
U_{i'}^1(t+1) &= \max\left\{U_{i'}^1(t) - \sum_b \mu_{i',b}(t), 0\right\} + \mu_{i,i'}(t) + \epsilon_1, \forall i \in \mathcal{W}, \\
U_i(t+1) &= \max\left\{U_i(t) - d_i(t), 0\right\} + \sum_a \mu_{a',i}(t) + \epsilon_1, \forall i \in \mathcal{R},
\end{aligned}$$

and let Λ_{ϵ_1} denote the capacity region of the supply chain under the additional flow ϵ_1, and $\bar{\boldsymbol{\mu}}^*(\epsilon_1)$ denote the solution of (13), when Λ is replaced by Λ_{ϵ_1}. Then, by Corollary 3.1 applied to the capacity region Λ_{ϵ_1}, we have that there exists a stationary randomized flow control algorithm, that chooses the flows based on the current state of the supply chain, and gives

$$E\{\mu_{1,i}^*(\epsilon_1)\} + \epsilon_1 = E\left\{\sum_b \mu_{i,b}^*(\epsilon_1)\right\}, \forall\ i \in \mathcal{M},$$

$$\sum_a E\left\{\mu_{a,i}^*(\epsilon_1)\right\} + \epsilon_1 = E\left\{\mu_{i,i'}^*(\epsilon_1)\right\}, \forall\ i \in \mathcal{W},$$

$$E\left\{\mu_{i,i'}^*(\epsilon_1)\right\} + \epsilon_1 = E\left\{\sum_b \mu_{i',b}^*(\epsilon_1)\right\}, \forall\ i \in \mathcal{W},$$

$$E\left\{\sum_a \mu_{a',i}^*(\epsilon_1)\right\} + \epsilon_1 \leq \bar{d}_i, \forall\ i \in \mathcal{R},$$

where $\bar{\mu}_{i,j}^*(\epsilon_1) = E\left\{\mu_{i,j}^*(\epsilon_1)\right\}$.

Similarly, let us now assume that a small flow ϵ_2 is added to the inputs of queues $U_i(t)^2$, such that their dynamics become

$$\begin{aligned}
U_i^2(t+1) &= \max\left\{U_i^2(t) - \mu_{1,i}(t), 0\right\} + \sum_b \mu_{i,b}(t) + \epsilon_2, \forall i \in \mathcal{M}, \\
U_i^2(t+1) &= \max\left\{U_i^2(t) - \sum_a \mu_{a,i}(t), 0\right\} + \mu_{i,i'}(t) + \epsilon_2, \forall i \in \mathcal{W}, \\
U_{i'}^2(t+1) &= \max\left\{U_{i'}^2(t) - \mu_{i,i'}(t), 0\right\} + \sum_b \mu_{i',b}(t) + \epsilon_2, \forall i \in \mathcal{W}.
\end{aligned}$$

and let Λ_{ϵ_2} denote the capacity region under the additional flow ϵ_2, and $\bar{\boldsymbol{\mu}}^*(\epsilon_2)$ denote the solution of (13), when Λ is replaced by Λ_{ϵ_2}.

As before, by Corollary 3.1 applied to the capacity region Λ_{ϵ_2}, we have that there exists a stationary randomized flow control algorithm, that chooses the flows based on the current state of the supply chain, and gives

$$E\left\{\sum_b \mu^*_{i,b}(\epsilon_2)\right\} + \epsilon_2 = E\{\mu^*_{1,i}(\epsilon_2)\}, \forall\ i \in \mathcal{M},$$

$$E\left\{\mu^*_{i,i'}(\epsilon_2)\right\} + \epsilon_2 = \sum_a E\left\{\mu^*_{a,i}(\epsilon_2)\right\}, \forall\ i \in \mathcal{W},$$

$$E\left\{\sum_b \mu^*_{i',b}(\epsilon_2)\right\} + \epsilon_2 = E\left\{\mu^*_{i,i'}(\epsilon_2)\right\}, \forall\ i \in \mathcal{W},$$

where $\bar{\mu}^*_{i,j}(\epsilon_2) = E\left\{\mu^*_{i,j}(\epsilon_2)\right\}$.

Note that by Assumption 3.2, such ϵ_1 and ϵ_2 do exist.

The flow control algorithm described in the previous section, minimizes the right-hand side of inequality (32) for all possible policies based on the current state of the supply chain, and in particular against the previously mentioned stationary policies, generated by adding the additional flows ϵ_1 and ϵ_2. Consequently, under the flow control algorithm, we have that

$$\Delta(\boldsymbol{U}(t)) - \delta E\{\boldsymbol{h}(\boldsymbol{\mu}(t))|\boldsymbol{U}(t)\} \le B\bar{N} - \sum_{i\in\mathcal{M}} U^1_i(t)\left(\sum_b \bar{\mu}^*_{i,b}(\epsilon_1) - \bar{\mu}^*_{1,i}(\epsilon_1)\right) -$$
$$-\sum_{i\in\mathcal{M}} U^2_i(t)\left(-\sum_b \bar{\mu}^*_{i,b}(\epsilon_2) + \bar{\mu}^*_{1,i}(\epsilon_2)\right) - \sum_{i\in\mathcal{W}} U^1_i(t)\left(\bar{\mu}^*_{i,i'}(\epsilon_1) - \sum_a \bar{\mu}^*_{a,i}(\epsilon_1)\right) -$$
$$-\sum_{i\in\mathcal{W}} U^2_i(t)\left(-\bar{\mu}^*_{i,i'}(\epsilon_2) + \sum_a \bar{\mu}^*_{a,i}(\epsilon_2)\right) - \sum_{i\in\mathcal{W}} U^1_{i'}(t)\left(\sum_b \bar{\mu}^*_{i,b}(\epsilon_1) - \bar{\mu}^*_{i,i'}(\epsilon_1)\right) -$$
$$-\sum_{i\in\mathcal{W}} U^2_{i'}(t)\left(-\sum_b \bar{\mu}^*_{i,b}(\epsilon_2) + \bar{\mu}^*_{i,i'}(\epsilon_2)\right) - \sum_{i\in\mathcal{R}} U_i(t)\left(E\{d_i(t)\} - \sum_a \bar{\mu}^*_{a',i}(\epsilon_1)\right) -$$
$$-\delta \sum_{(i',j)} \boldsymbol{f}_{i',j}(\bar{\mu}^*_{i',j}(\epsilon_1)) + \delta \sum_{i\in\mathcal{M}} \boldsymbol{g}_{1,i}(\bar{\mu}^*_{1,i}(\epsilon_1)) + \delta \sum_{(i,j)} \boldsymbol{g}_{i,j}(\bar{\mu}^*_{i,j}(\epsilon_1)) +$$
$$+\delta \sum_{(i,i')} \boldsymbol{g}_{i,i'}(\bar{\mu}^*_{i,i'}(\epsilon_1)) + \delta \sum_{(i',j)} \boldsymbol{g}_{i',j}(\bar{\mu}^*_{i',j}(\epsilon_1)),$$

Denoting $\epsilon = \min\{\epsilon_1, \epsilon_2\}$, we further have

$$\Delta(\boldsymbol{U}(t)) - \delta E\{\boldsymbol{h}(\boldsymbol{\mu}(t))|\boldsymbol{U}(t)\} \le$$

$$\leq B\bar{N} - \epsilon \sum_{j=1}^{2} \sum_{i \in \mathcal{M}} U_i^j(t) - \epsilon \sum_{j=1}^{2} \sum_{i \in \mathcal{W}} \left[U_i^j(t) + U_{i'}^j(t) \right] - \epsilon \sum_{i \in \mathcal{R}} U_i(t) - \delta \boldsymbol{h}(\bar{\boldsymbol{\mu}}^*(\epsilon_1)).$$

By Theorem 5.1, it follows that

$$\limsup_{t\to\infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \left[\sum_{j=1}^{2} \left(\sum_{i \in \mathcal{W}} E\{U_i^j(\tau)\} + \sum_{i \in \mathcal{M}} E\{U_i^j(\tau) + U_{i'}^j(\tau)\} + \sum_{i \in \mathcal{R}} E\{U_i(\tau)\} \right) \right] \leq$$

$$\leq \frac{BN + \delta(\bar{\boldsymbol{h}} - \boldsymbol{h}(\boldsymbol{\mu}^*(\epsilon_1)))}{\epsilon} \leq \frac{BN + \delta \boldsymbol{h}_{max}}{\epsilon} \tag{38}$$

and

$$\liminf_{t\to\infty} \boldsymbol{h}(\bar{\boldsymbol{\mu}}(t)) \geq \boldsymbol{h}(\boldsymbol{\mu}^*(\epsilon_1))) - \frac{BN}{\delta}, \tag{39}$$

The performance bounds in (38) and (39) hold for any values of ϵ_i such that $0 < \epsilon_i \leq \mu_{max}$, for $i = 1,2$. However the particular values of ϵ_i only affect the values of the bounds and not the control algorithm. Therefore, we can optimize the bounds separately over all possible values of ϵ_i, $i = 1,2$. Obviously the bound (38) is minimized when ϵ approaches μ_{max}. It can be shown that the optimal solution of (13) when the capacity region is replaced by Λ_{ϵ_1}, is continuous in ϵ_1. Consequently, as ϵ_1 approaches zero, the capacity region Λ_{ϵ_1} approaches Λ and $\boldsymbol{\mu}^*(\epsilon_1)$ approaches $\boldsymbol{\mu}^*$. Therefore the bound (39) is minimized when ϵ_1 goes to zero, and the result follows. ■

Remark 5.2: Note that inequality (36) shows that under the flow control algorithm, the queues remain stable, i.e., the long-run flows are feasible. In addition, inequality (37) shows that under the flow control algorithm we can get arbitrarily close to the optimal solution, by making δ arbitrarily large.

VI. Extensions to multiple firms and products

In the previous sections we introduced a dynamic algorithm for controlling the flow of product in a time-varying supply chain, in the case of a single firm. In this section we show how the algorithm can be extended to multiple firms. We assume that several firms produce a homogeneous product and that they share the same markets. Figure 3 shows an example of a supply chain corresponding to the multiple-firms scenario, which is similar to the supply chain network introduced in [9], where multiple firms share the same market.

Let $\boldsymbol{h}^l(\bar{\boldsymbol{\mu}}) = \boldsymbol{f}^l(\bar{\boldsymbol{\mu}}) - \boldsymbol{g}^l(\bar{\boldsymbol{\mu}})$ be the profit functions corresponding to each of the firms, where the notations used are similar to the ones used in the previous sections. Note that the main

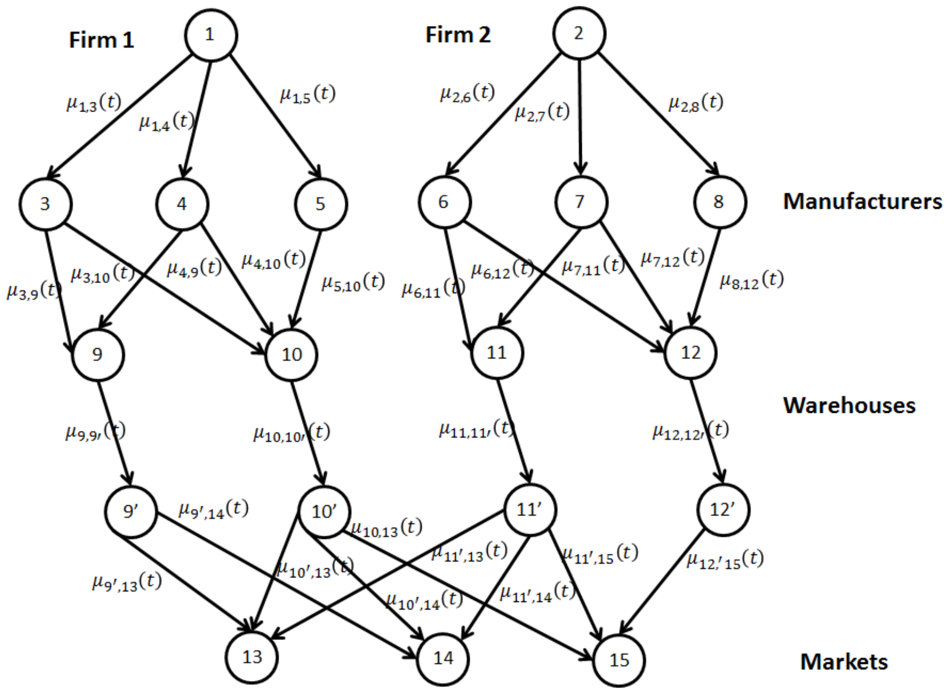

Fig. 3: Multiple firms supply chain network

difference between the single-firm scenario and the multiple-firms scenario consists of the fact that the markets receive flows of products from multiple firms rather than a single one.

We define the global cost function

$$\boldsymbol{H}(\bar{\boldsymbol{\mu}}) = \sum_{l} w_l \boldsymbol{h}^l(\bar{\boldsymbol{\mu}}),$$

where w_l are positive scalars. In the multiple firms scenario, the weighting factors w^l could represent bargaining power coefficients for the market sharing. Let $\mathcal{F}$ be the set of firms, $\mathcal{R}$ be the set of retailers, and $\mathcal{M}_l$, $\mathcal{W}_l$ denote the manufacturing units and the warehouses corresponding to the firm l, respectively.

Let us now define the following stochastic optimization problem

$$\max_{\bar{\mu}} \quad \boldsymbol{H}(\bar{\boldsymbol{\mu}}) \tag{40}$$

$$\text{subject to:} \quad \bar{\mu}_{l,i} = \textstyle\sum_{b\in\mathcal{W}_l} \bar{\mu}_{i,b}, \forall i \in \mathcal{M}_l, \forall l \in \mathcal{F}, \tag{41}$$

$$\textstyle\sum_{a\in\mathcal{M}_l} \bar{\mu}_{a,i} = \bar{\mu}_{i,i'}, \forall i \in \mathcal{W}_l, \forall l \in \mathcal{F},$$

$$\bar{\mu}_{i',i} = \textstyle\sum_{b\in\mathcal{R}} \bar{\mu}_{i',b}, \forall i \in \mathcal{W}_l, \forall l \in \mathcal{F},$$

$$\textstyle\sum_{l\in\mathcal{F}} \sum_{a\in\mathcal{W}_l} \bar{\mu}_{a',i} \leq \bar{d}_i, \forall i \in \mathcal{R}.$$

Using a similar approach as in the case of a single firm, we can derive a flow control algorithm that gets arbitrary close to the maximum of (40), while satisfying the equality constraints (41).

VII. Numerical example

We implement the flow control algorithm described above for several examples of supply chain networks. For each link (i, j) of the network, we assume that the cost function has the form $g_{i,j}(\mu_{i,j}) = a_{i,j}\mu_{i,j}^2 + b_{i,j}\mu_{i,j}$. The revenue function is assumed to have the form $f(\mu) = c\mu_{i,j}^{\frac{1}{p}} + d$. In all examples we have taken $a_{i,j} = 0.1$, $a_{i,j} = 0.3$, $c = 3$, $d = 2$, $p = 1.8$. The maximum output rate at node i is assumed to be equal to $\mu_{max} = 6 \times L_i$, where L_i is the number of links going out of i. This sets an "average" maximum rate of 6 for each link. The links ON-OFF processes are assumed to be i.i.d. with an 'ON' probability of 0.9. The demand processes are taken to be independent and uniformly distributed between 0 and 3 at each time, with an average of 1.5. In addition, we vary the value of the parameter δ within the set $\{0.1, 0.5, 0.9\}$ to study its influence on the queues' backlog.

We consider six different supply chain network topologies shown in Figure 4. Example 4a is a 1-branch network with one firm and one retailer. In our examples, a branch is a sequence Firm-Manufacturer-Warehouses(1,2)-Retailer. Network 4b has one firm and two retailers. The firm is connected to each retailer by a separate branch. The topology in 4c has one source and one destination connected by two disjoint branches. In the last three examples (Figures 4d,4e,4f) we consider the network in Figure 4b with cross links between the two branches.

For each network, we plot the queues' backlog over time as well as the running averages of the queues (Figures 5). We show both the plots for the forward $U_i^{(1)}$ and backward $U_i^{(2)}$ queues. Recall that these queues are virtual queues introduced as a consequence of modeling the

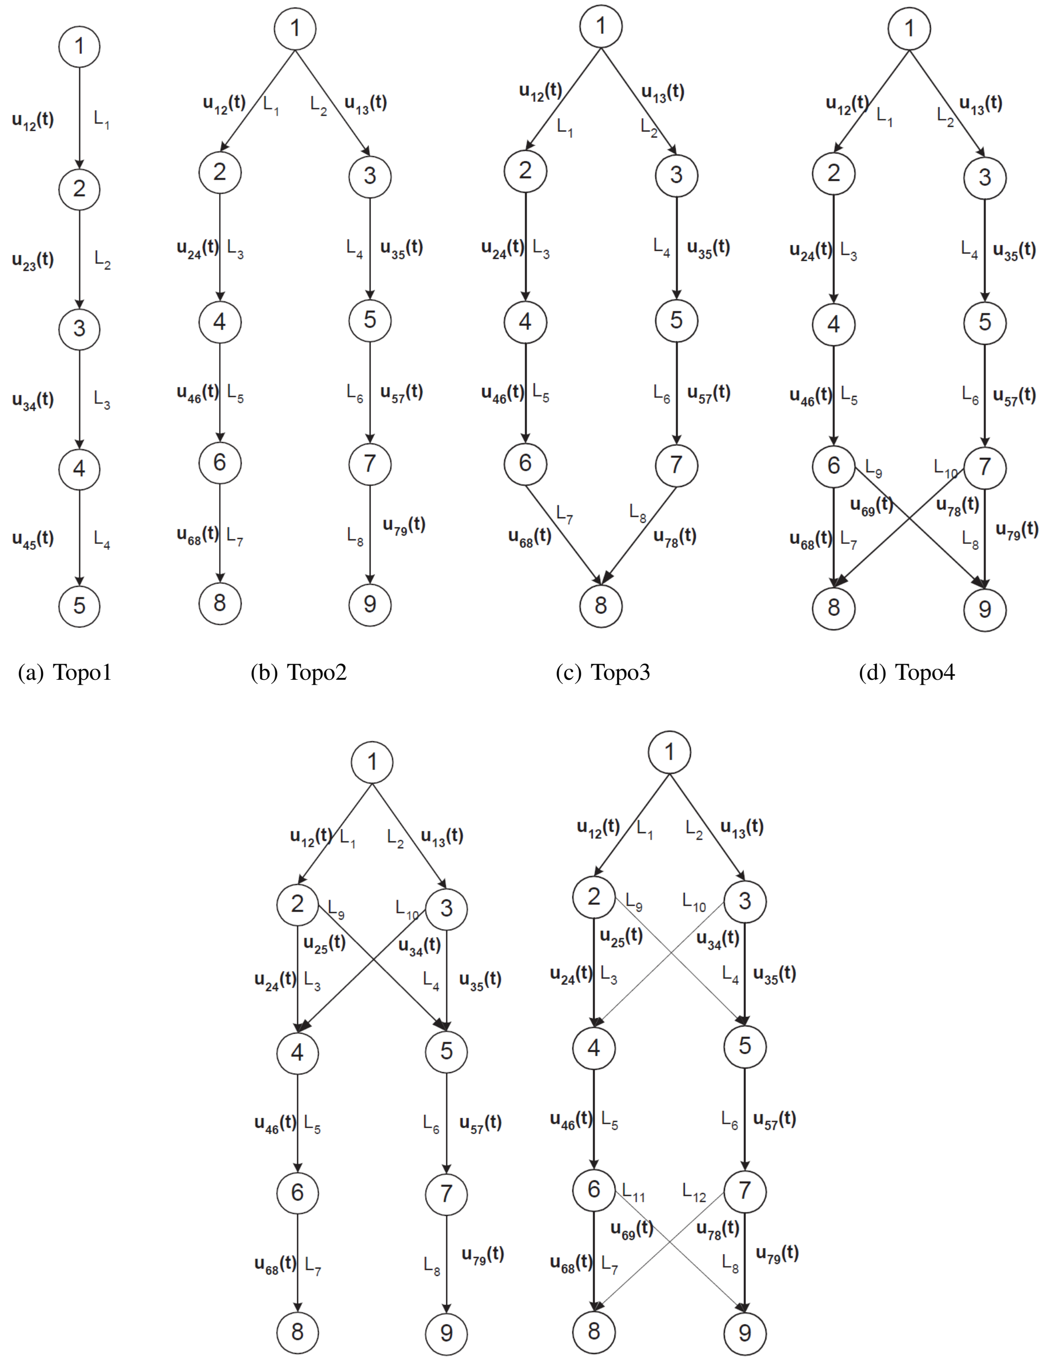

(a) Topo1 (b) Topo2 (c) Topo3 (d) Topo4

(e) Topo5 (f) Topo6

Fig. 4: Examples of network topologies: (a) single branch network; (b) 2 branches, 2 retailers; (c) 2 branches, 1 retailer; (d) 2 branches, 2 retailers, downstream crossing; (e) 2 branches, 2 retailers, upstream crossing; (f) 2 branches, 2 retailers, up and downstream crossings.

inequality constraints as queues. However, the forward queues can be interpreted as real queues at nodes of the network. Also, notice that there is no backward queue defined for the queue at a retailer. In all examples, the queues' backlog of each branch are shown in a 4-by-2 panel where the left column corresponds to the froward queues (from top to bottom nodes) and the right column corresponds to the backward queues. For instance in Figure 5b, the top left plot shows the forward queue backlog at node 2 and the top right plot is its backward queue. The second top row shows the forward and backward queues at node 4. The next row shows the queues at node 6 and the bottom row shows the queue at node 8 (notice that there is no backward queue at node 8). The queues at nodes 3,5,7,9 are shown in a similar manner in Figure 5c.

We also plot the flow rates on the different links and their running averages (Figures 6). In all examples, the link flow rates are shown in a 2-column panel, except for example 1 (Figure 6a) where the link rates are shown in a 1-column panel. The left column shows links (originating) in the left branch and the right column displays the rates of links (originating) in the right branch. For instance, in Figure 6b, the left column shows (from top to bottom) the rates in links L_1, L_3, L_5, L_7 of the topology on Figure 4b; the right column shows the rate at links L_2, L_4, L_6, L_8. When there are cross-branch links as in examples of Figures 6d-6h, we show them in the bottom subplots. For instance, in Figure 6d, the rate at link L_9 is shown in the bottom of the left column and link L_{10}'s rate in the bottom of the right column. Finally, to emphasize on the convergence of the average rates, we zoom-in the rate plots to focus only on the $[0, 1.8]$-range of the y-axis (i.e. the rates). This is shown in Figure 7.

A. Discussion: Queues' backlog

The queues' backlog are shown in Figures 5. It can be observed that in all examples, the queues are oscillating but are not growing unbounded. This is exactly the stability of queues predicted by the theory. We will see in the discussion of the rates (in Figures 6) that the average rate also satisfies the stability condition of Proposition 4.1 for all examples. The theory however, does not predict anything about the convergence of the average queues' backlog. Yet, we have observed in all our examples that the average backlog seems to converge for all (forward and backward) queues with a faster convergence in networks 4d-4f. An interesting follow up of this study is to prove/disprove convergence of average queue and to determine under which conditions convergence is guaranteed.

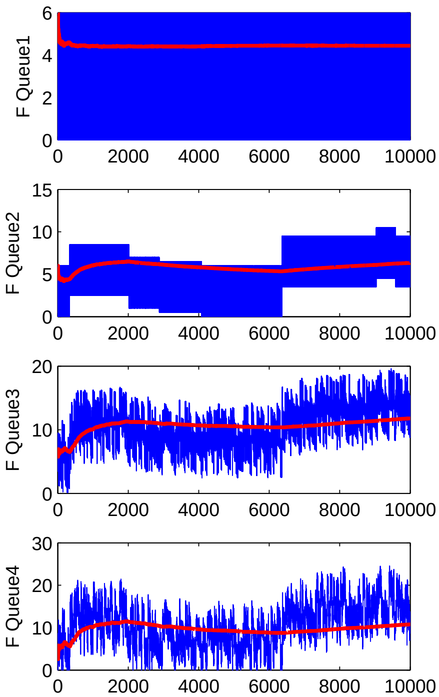

(a) Topo1: Queues

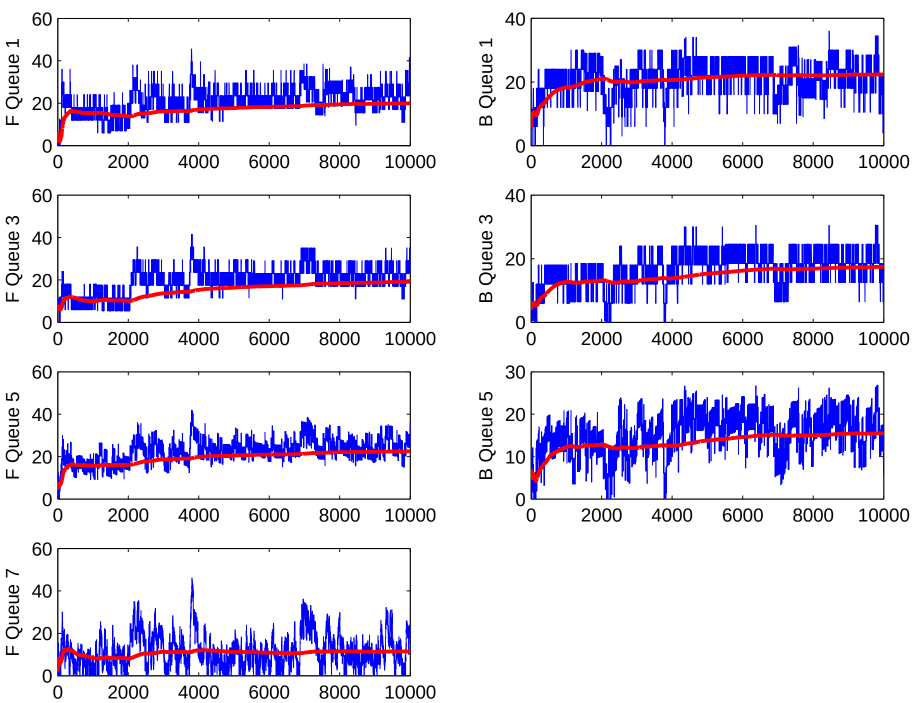

(b) Topo2: Queues branch 1

Fig. 5: Queue levels

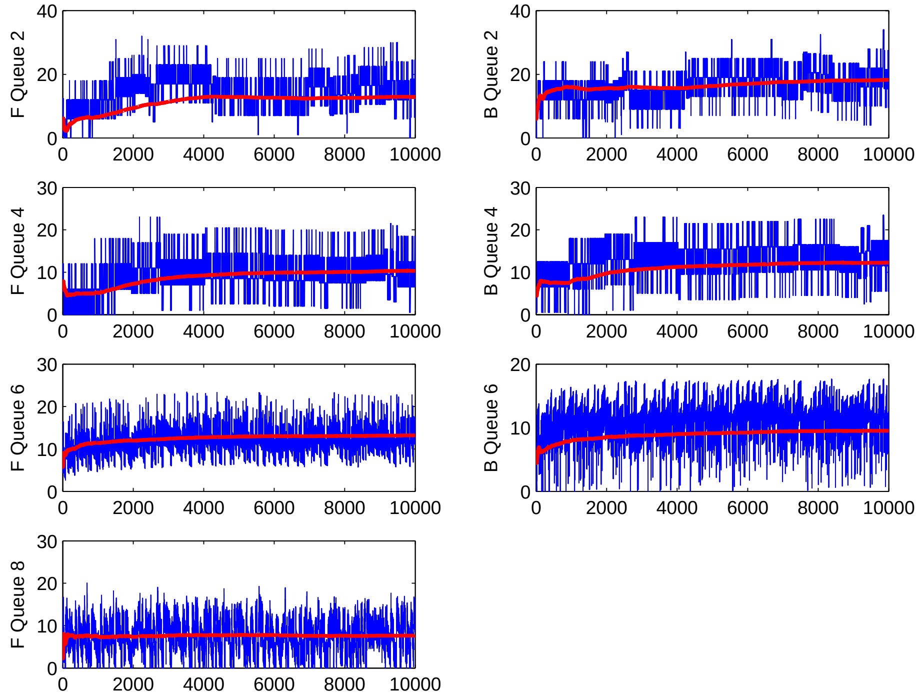

(c) Topo2: Queues branch 2

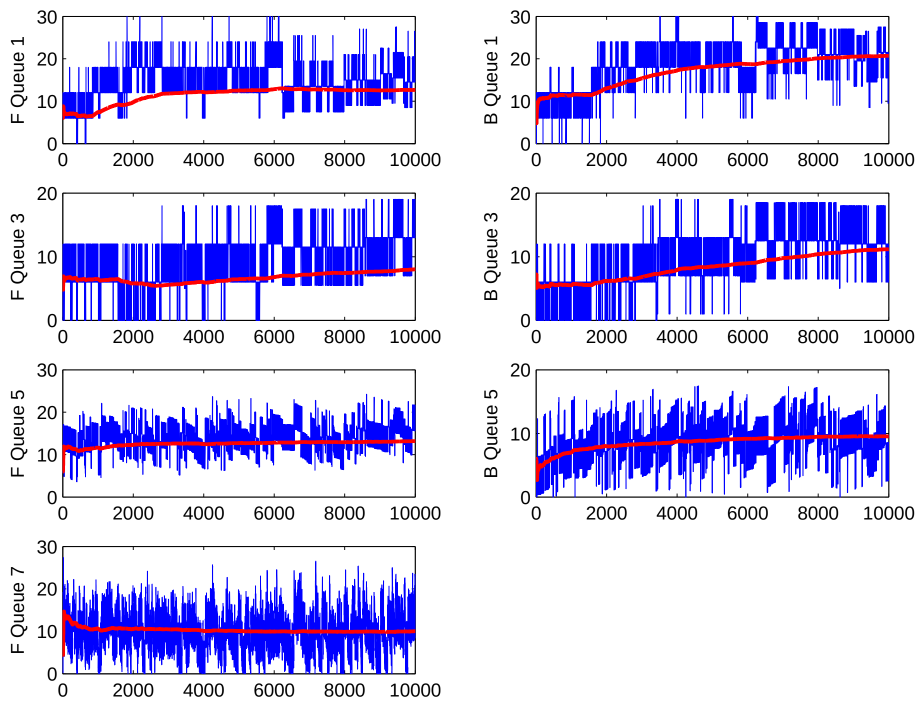

(d) Topo3: Queues branch 1

Fig. 5: Queue levels (cont.).

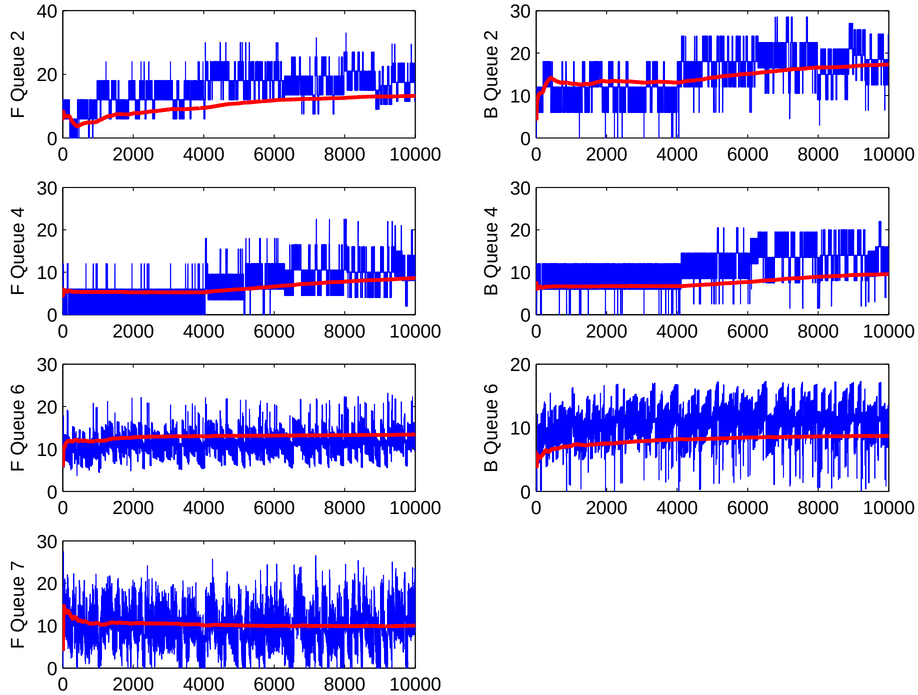

(e) Topo3: Queues branch 2

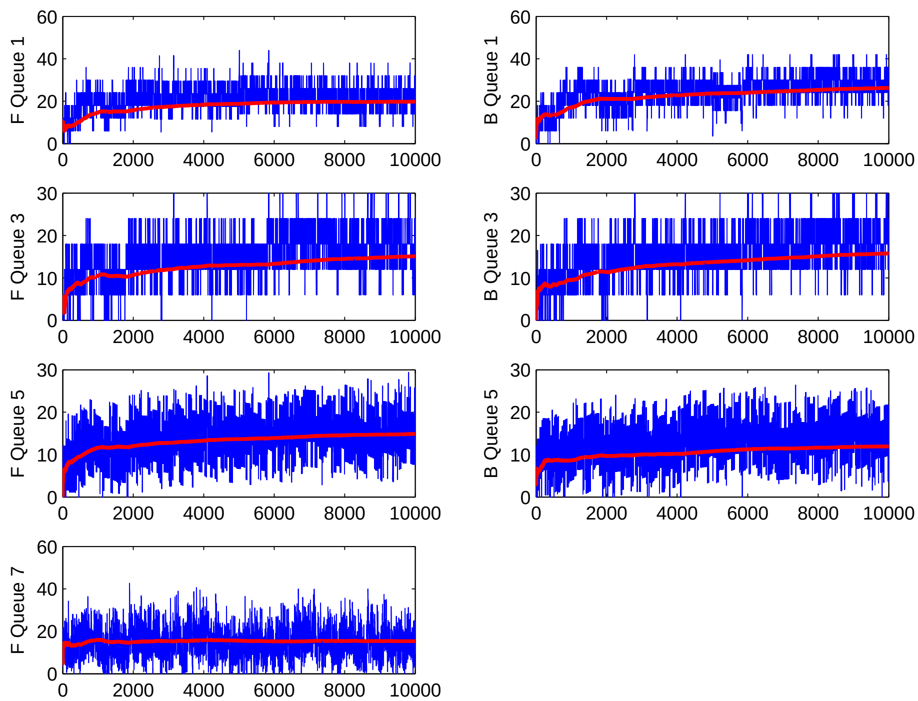

(f) Topo4: Queues branch 1

Fig. 5: Queue levels (cont.).

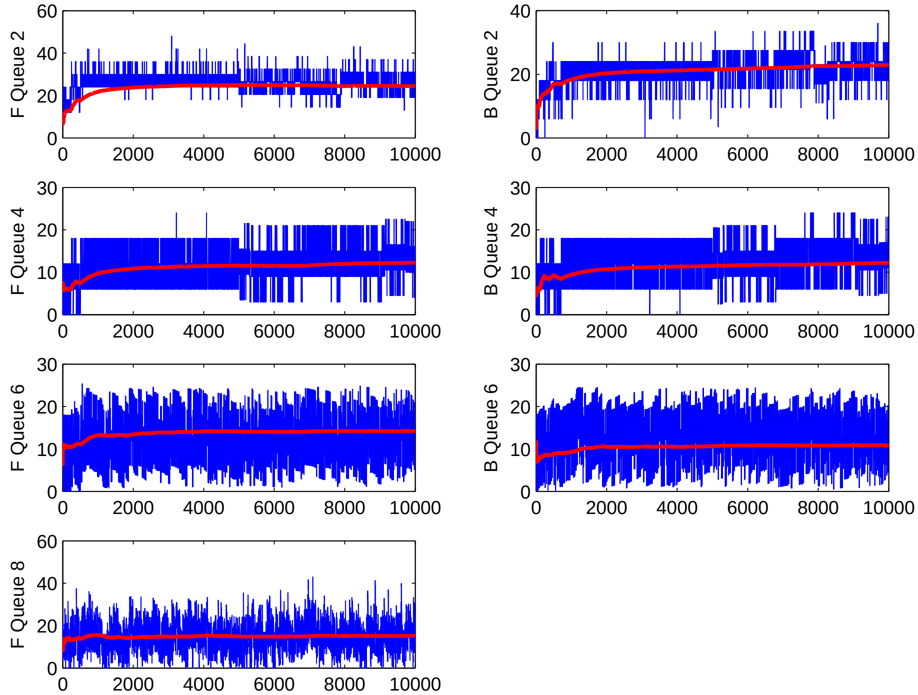

(g) Topo4: Queues branch 2

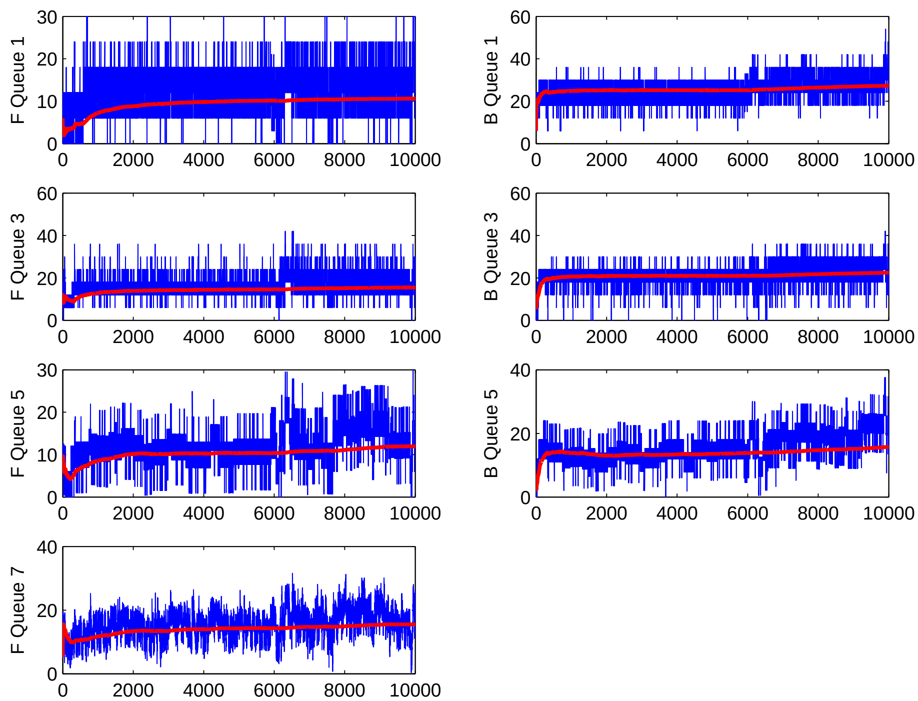

(h) Topo5: Queues branch 1

Fig. 5: Queue levels (cont.).

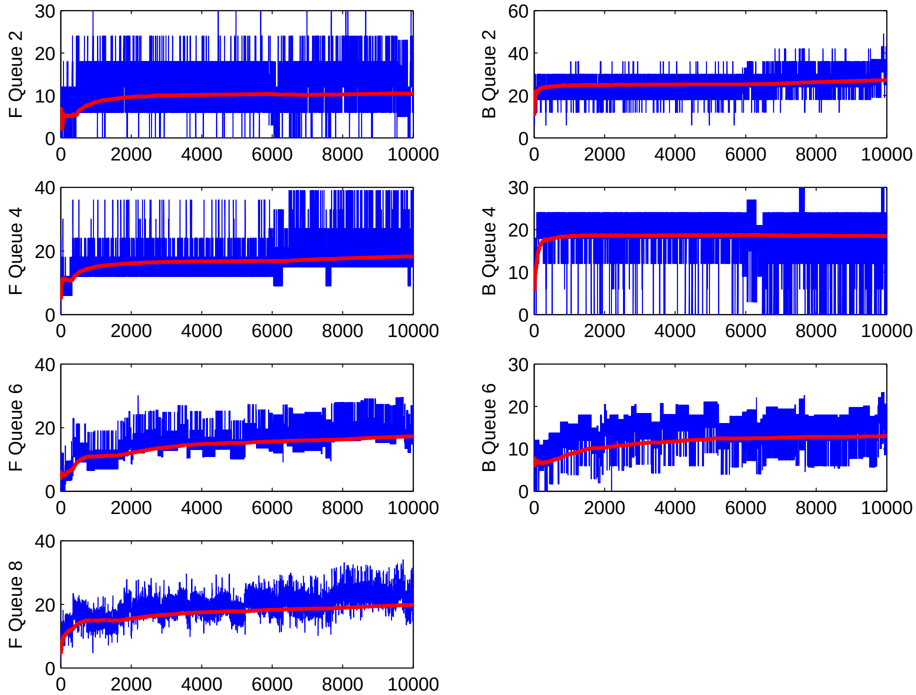

(i) Topo5: Queues branch 2

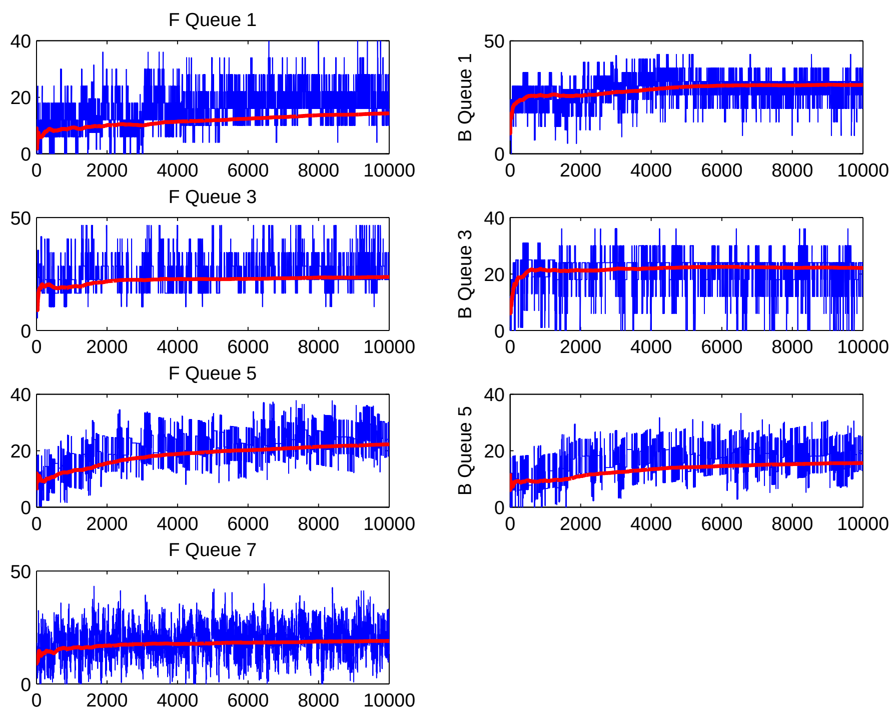

(j) Topo6 ($\delta = 0.1$): Queues branch 1

Fig. 5: Queue levels (cont.).

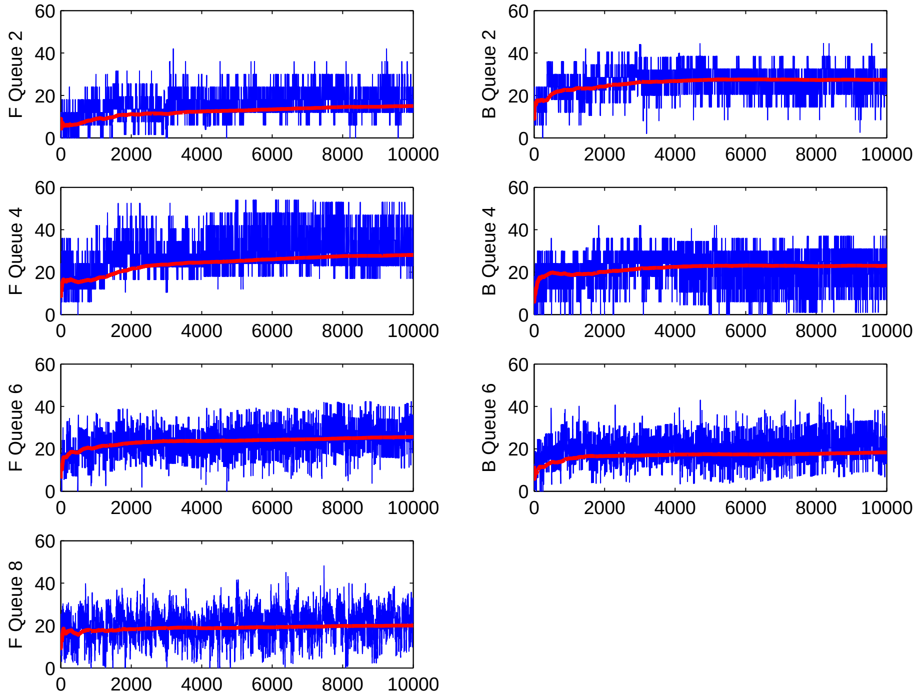

(k) Topo6 ($\delta = 0.1$): Queues branch 2

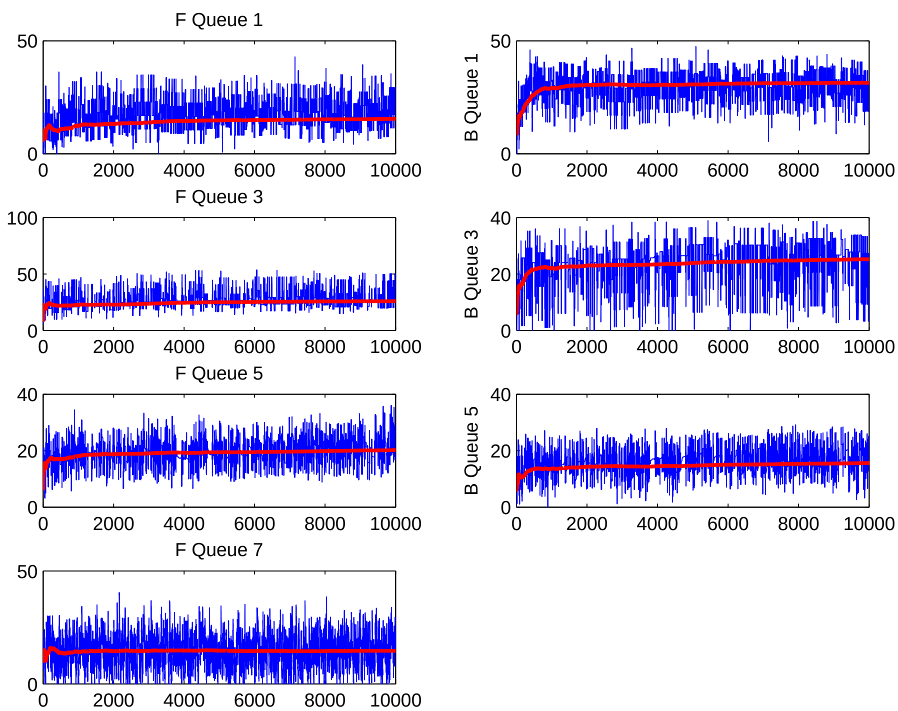

(l) Topo6 ($\delta = 0.5$): Queues branch 1

Fig. 5: Queue levels (cont.).

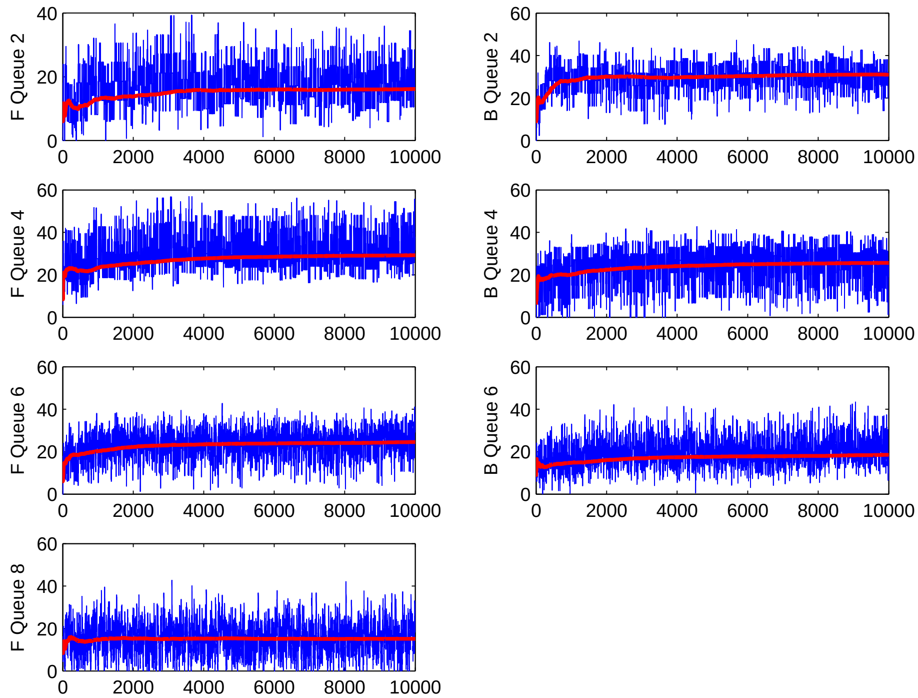

(m) Topo6 ($\delta = 0.5$): Queues branch 2

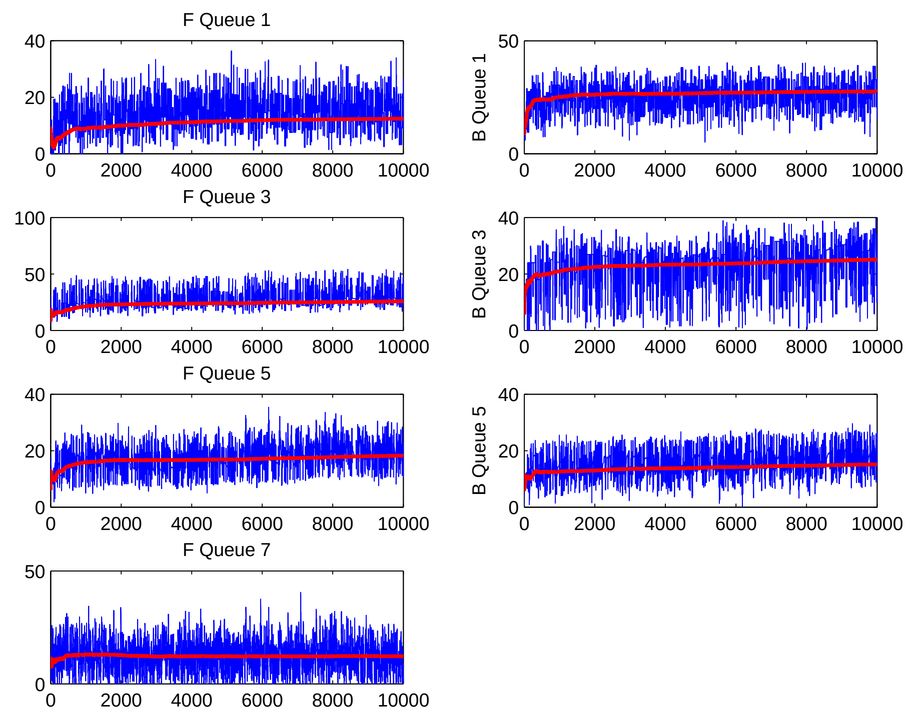

(n) Topo6 ($\delta = 0.9$): Queues branch 1

Fig. 5: Queue levels (cont.).

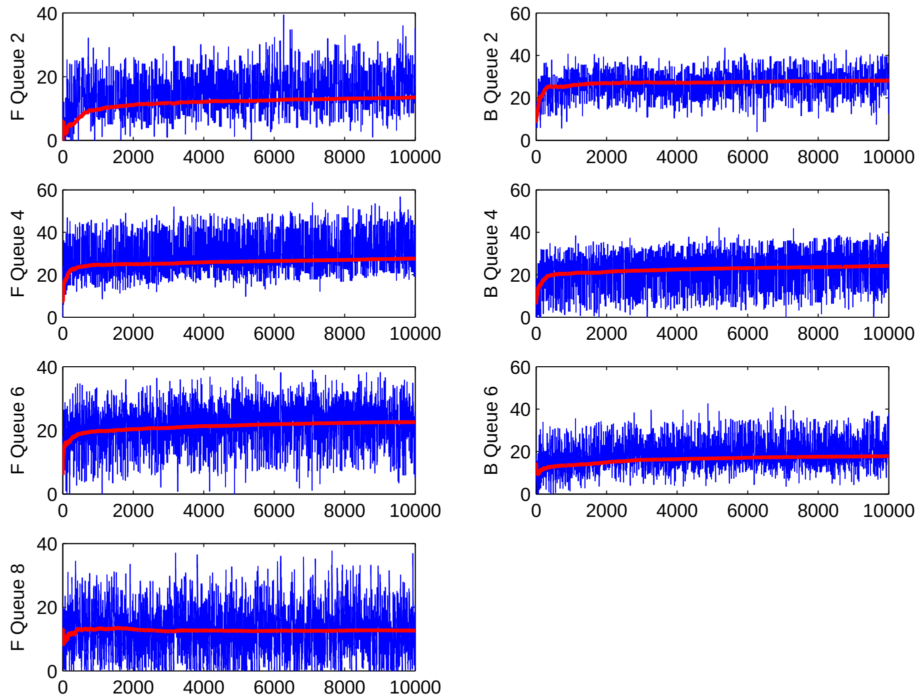

(o) Topo6 ($\delta = 0.9$): Queues branch 2

Fig. 5: Queues levels (cont.).

From the figures, it can also be noticed that in general, the forward queues at the manufacturers and at the first (upper) warehouses are in average more loaded that the queues at the second (lower) warehouses and at that retailers. This is a consequence of the back-pressure algorithm which forces upstream nodes to reduce their rate and consequently build up their queues when downstream nodes are congested. Hence, in general, queues close to the destination tend to have a smaller backlog.

We also observe that adding cross links decreases the rate of change (fluctuation) of the queues. This can be seen by comparing the fluctuations of Figures 5a-5e to the rest of the plots. This is expected because additional links give more option for load balancing and help regularize the queues. However, the average queues' backlogs are comparable across all examples. This is also not surprising. Indeed, due to the maximum output rate constraint and the way it is implemented in the simulations, from the perspective of any given node, the topologies in examples 4b-4f are identical (same "aggregate" input/output process). As a consequence, the behavior of each queue is expected to be the same in average. Hence, the additional links (only) serve to balance the load to reduce the variations in each queue. This explains the faster convergence observed for networks 4d-4f.

Finally, we have noticed that the queue fluctuations increase for higher values of the parameter δ (Figures 5j-5k, 5l-5m, and 5n-5o). Recall that setting δ large implies focusing on maximizing the profit (and getting arbitrarily close to the optimal solution), but at the cost of increased product congestion in the queues.

B. Discussion: Link Rates

The rates at the different links are shown in Figures 6 and 7. A certain number of observations can be made from the figures.

First, we see that the rates are random due to our randomized control algorithm. However, for all runs of the simulation, the average rate converges for each link. Furthermore, at each retailer, the value to which the average aggregate rate converges is less than 1.5, the average demand at each market. This is a necessary condition for the stability of the queues as was stated in Proposition 4.1. The average rate at the other links are such that the conservation of flow principle is satisfied at each node (which is what we expected).

In the case of topology 4a, the average rates at all links converge to 1.5 (see Figure 7a); similarly for each branch of example 4b. For example 4c, the average rate on each branch is equal to 0.75, giving an aggregate average rate of 1.5 at the retailer. The average rate in the first three links of each branch of network 4d converges to 1.5. At the branching nodes 6 and 7, the traffic evenly splits on the two links, as is shown in Figure 7d. The next examples show a similar behavior with even splitting at branching nodes. Notice that the introduction of additional links (in topologies 4d-4f) slows down the convergence of the average rate.

Recall that in our simulations, the maximum output rate at node i is equal to $\mu_{max} = 6 \times L_i$, where L_i is the number of links going out of i. This implies that when there is only one link going out of a node, the maximum rate at the link is equal to 6. This can be observed especially in Figure 6a, where the maximum rate at all links is equal to 6. When there are 2 links departing from a node, traffic can either be split (when both links are up), or entirely sent over one link (especially when the other link is down). To see which choice will be made at a given node, one can analyze the cost function $g_{i,j}(\mu_{i,j}) = a\mu_{i,j}^2 + b\mu_{i,j}$. Assume that at a branching node, traffic is split such that a rate of μ is sent over one link and $12-\mu$ over the other link $0 \leq \mu \leq 12$. The (local) total cost of such routing is $a\mu^2 + b\mu + a(12-\mu)^2 + b(12-\mu) = a(2\mu^2 - 24\mu) + 12b + 144a$. Analyzing this cost as a function of μ, we observe that it is minimized when $\mu = 12$, implying that at a branching node, when both links are up, the entire traffic should be sent over one of the links. This is what we observe at nodes 2 and 3 for network 4d, nodes 6 and 7 for network 4e, and nodes 2,3,6,7 for network 4f shown respectively in Figures 6d, 6e, 6f, where the traffic on the links departing from such nodes is (almost all the time) either 0 or equal to the maximum rate of 12. This general rule is however not verified at branching node 1. Here (top plots of Figures 6b-6f) the maximum rates at links L_1 and L_2 are most of the time equal to 6 each. In fact, a closer look at the simulations shows that the whole traffic is sent over one link mostly because the other is down.

VIII. Conclusions

In this paper we addressed the flow control in a time-varying, random supply chain network. For a single-firm supply chain network, we proposed a randomized flow control algorithm for maximizing the profit function in terms of the time averages of the flows, in the long-run. We have also shown how the algorithm can be used for multiple-firms networks. The algorithm can get arbitrarily close to the optimal solution, can be implemented in the distributed manner and does not require knowledge of the probability distribution of the random process that drives the supply chain network. The design of the algorithm was based on concepts from queuing theory and stochastic optimization. We presented numerical simulations of our algorithm for different supply chain topologies; simulations that confirmed the intuition induced by our theoretical results.

References

[1] B. Behdani. Risk management and paradoxical situations in modern supply chains. *The CIP Report on Supply Chain*, July 2011.

[2] R. Blasgen. Supply chain management: The driving force behind the global economy. *The CIP Report on Global Supply Chain*, July 2011.

[3] Canadian Competition Bureau. Competition bureau concludes examination into gasoline price spike following hurricane Katrina (http://www.competitionbureau.gc.ca), 2006.

[4] S. Chopra and P. Meindl. *Supply Chain Management, 3rd ed.* Prentice Hall, 2006.

[5] Fairtrade Foundation. Spilling the bean (http://www.fairtrade.org.uk), 2002.

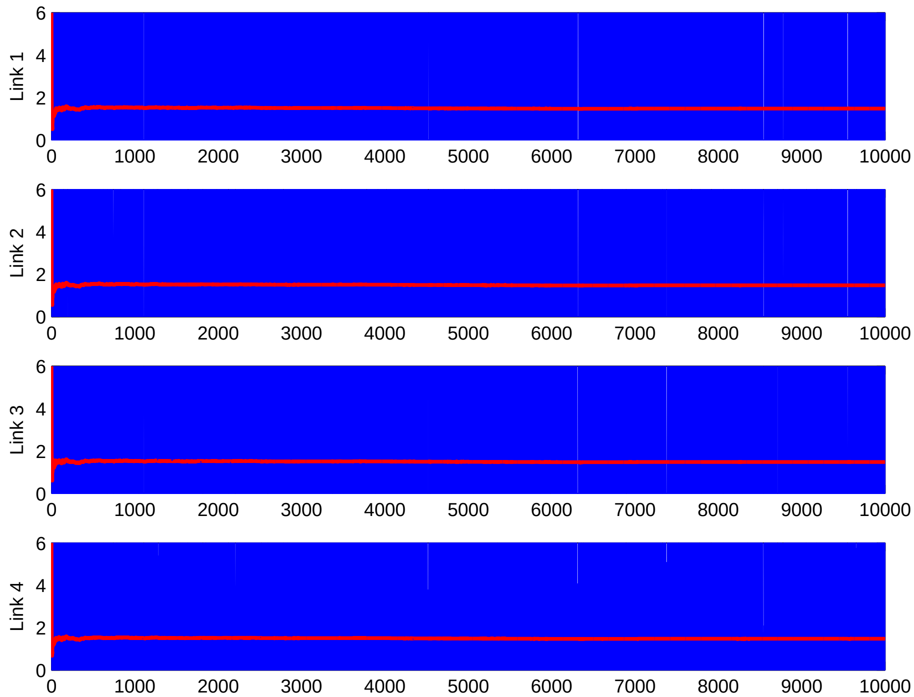

(a) Topo1: Rates

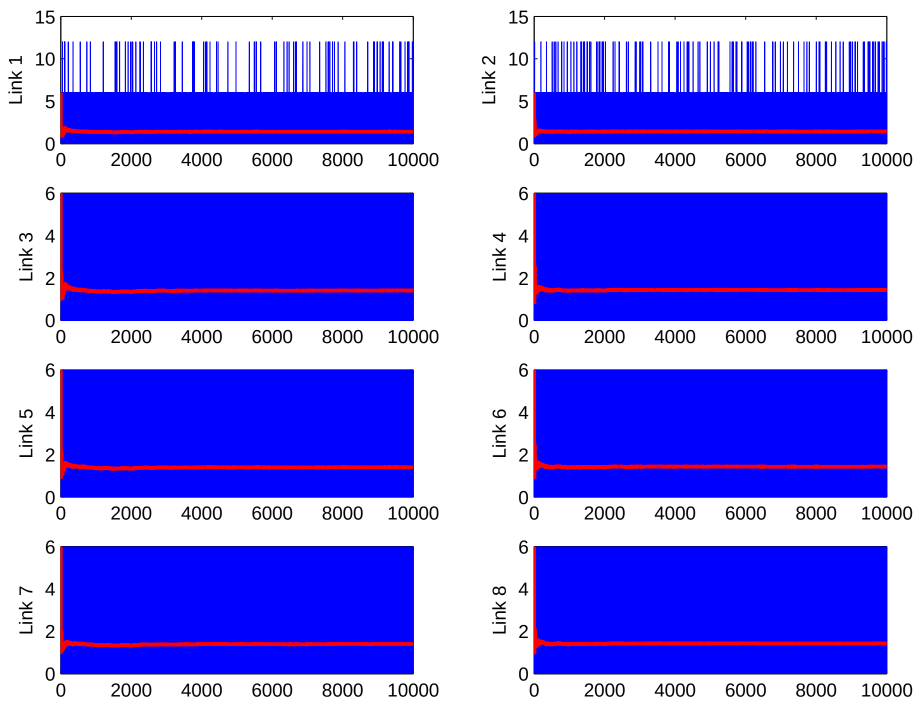

(b) Topo2: Rates

Fig. 6: Link Rates.

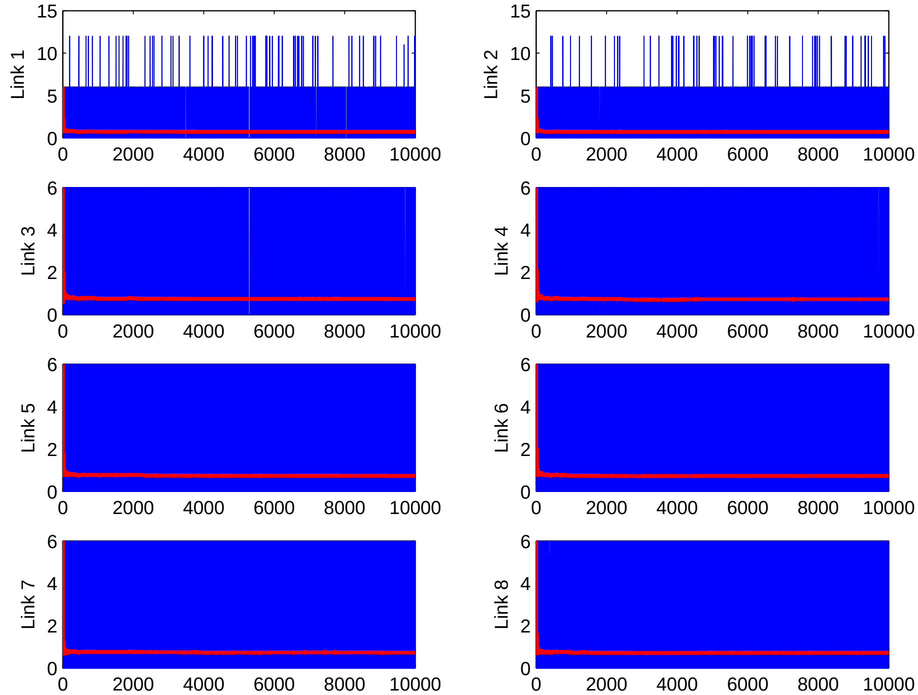

(c) Topo3: Rates

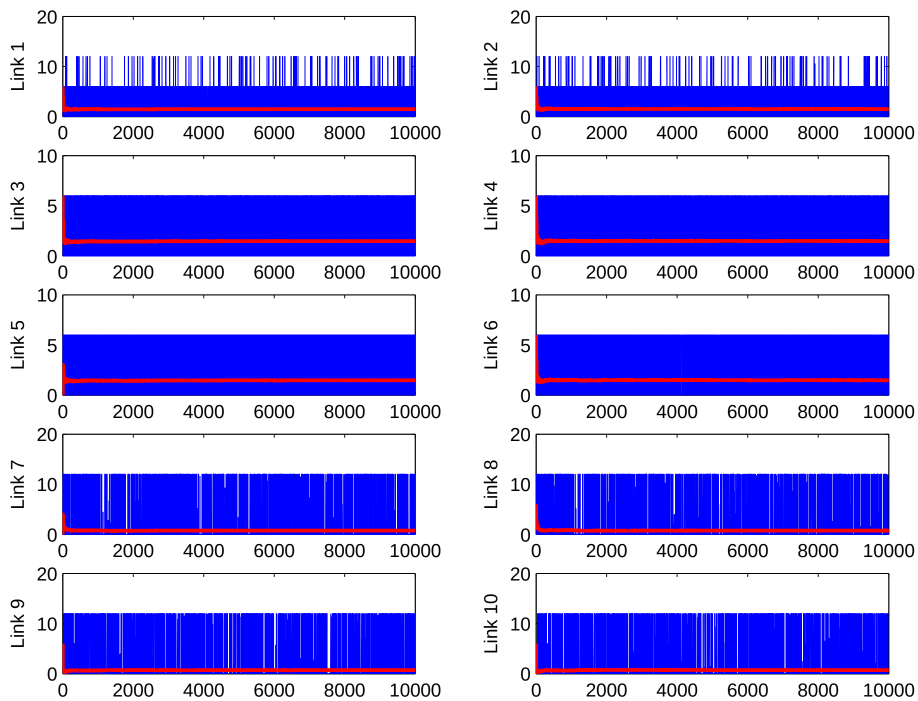

(d) Topo4: Rates

Fig. 6: Link Rates (cont.).

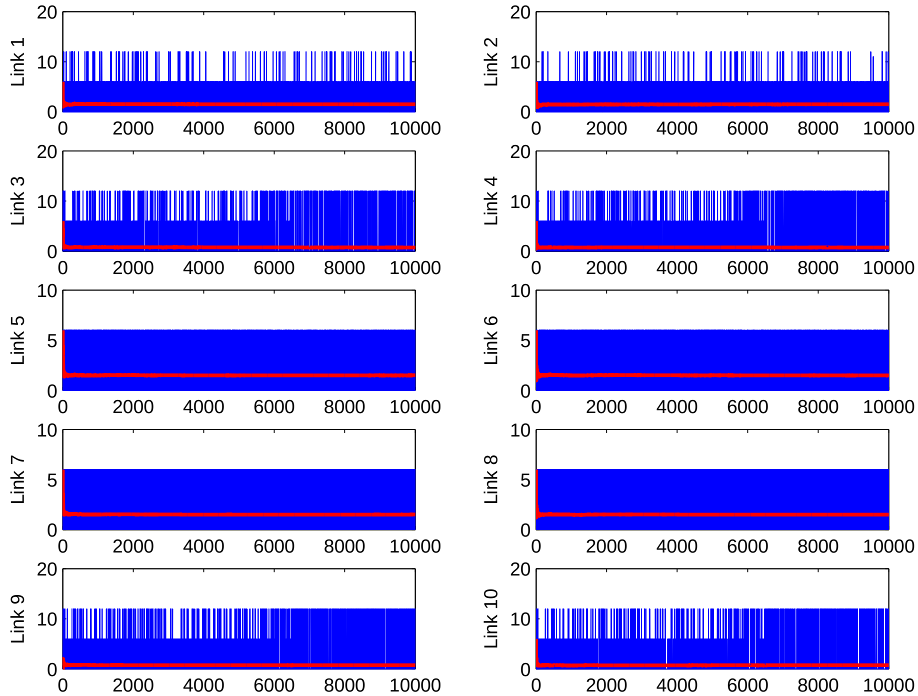

(e) Topo5: Rates

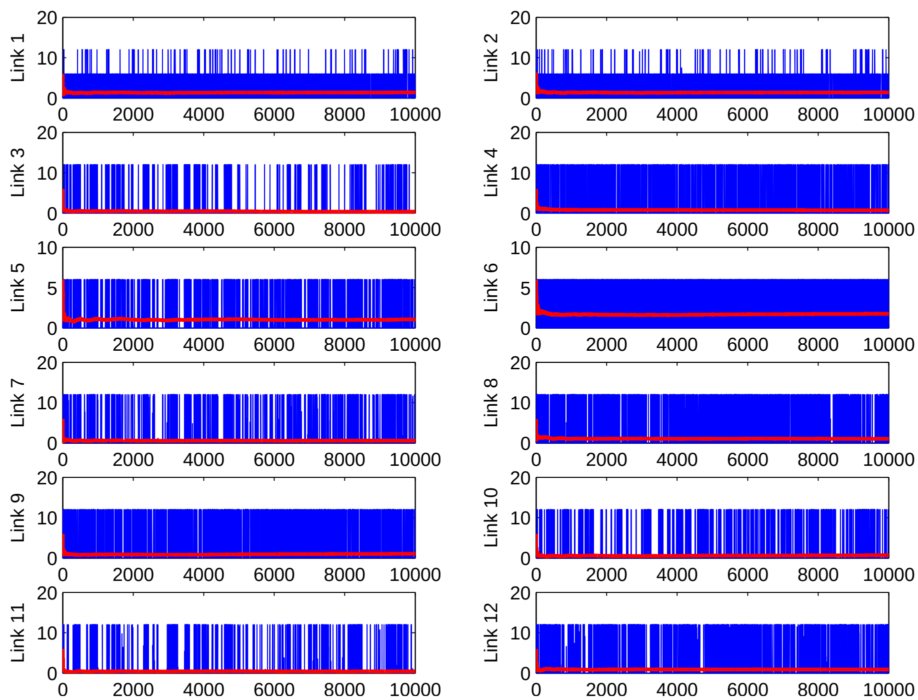

(f) Topo6 ($\delta = 0.1$): Rates

Fig. 6: Link Rates (cont.).

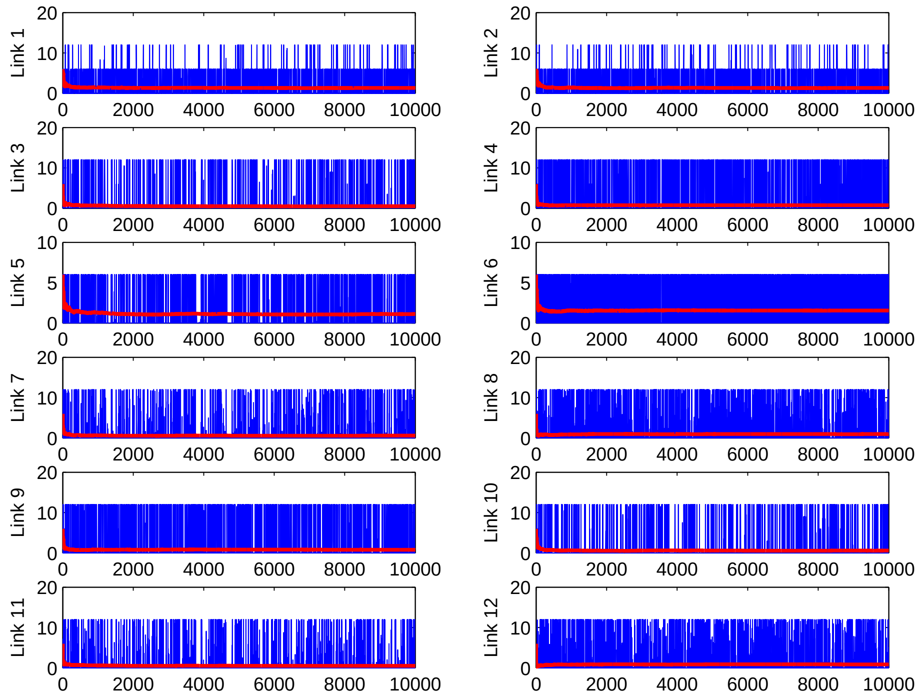

(g) Topo6 ($\delta = 0.5$): Rates

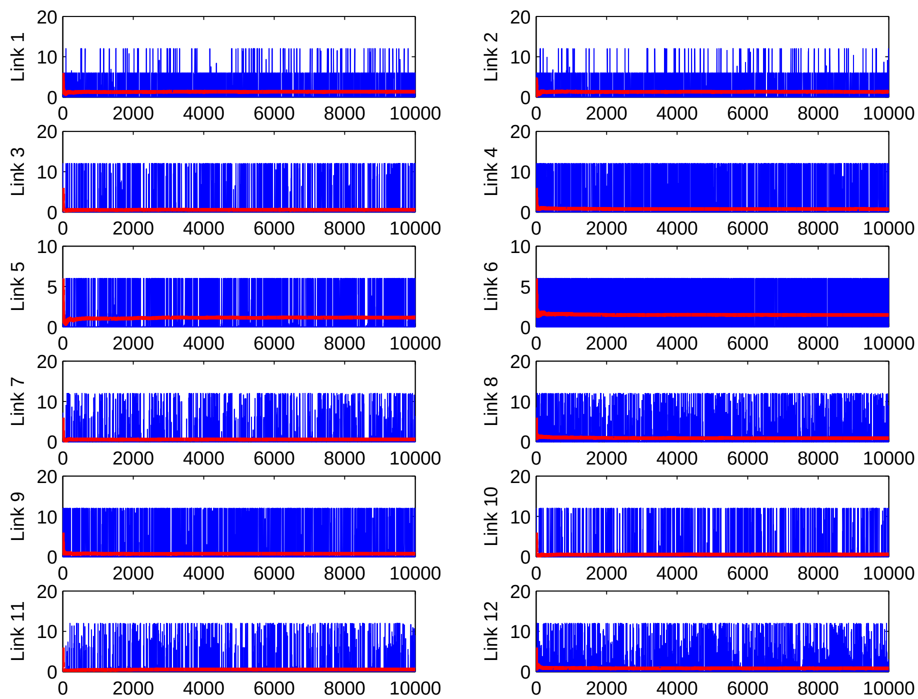

(h) Topo6 ($\delta = 0.9$): Rates

Fig. 6: Link Rates (cont.).

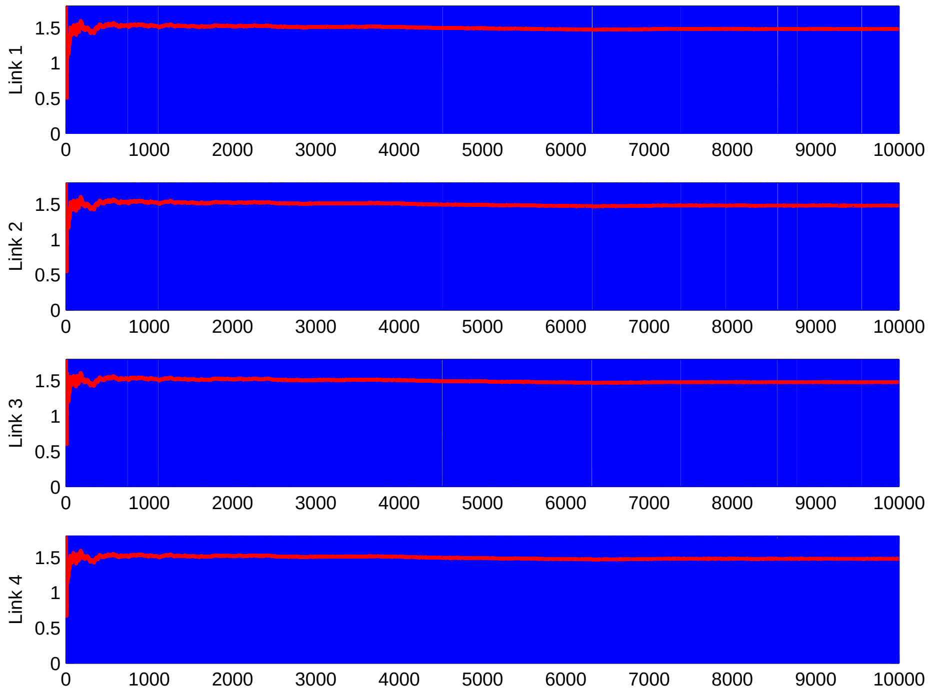

(a) Topo1: Rates

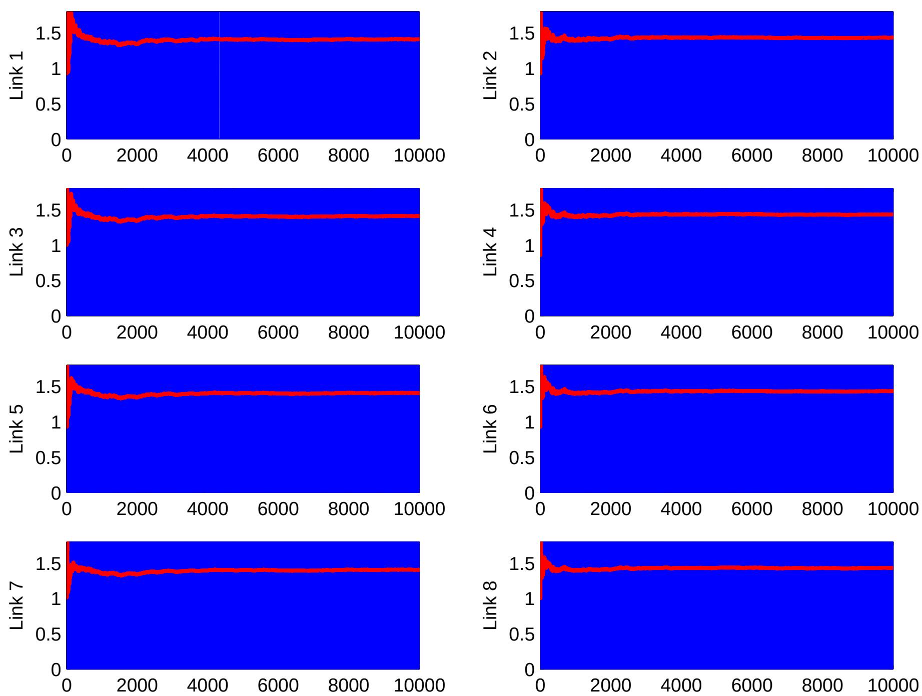

(b) Topo2: Rates

Fig. 7: Link Rates - zooming.

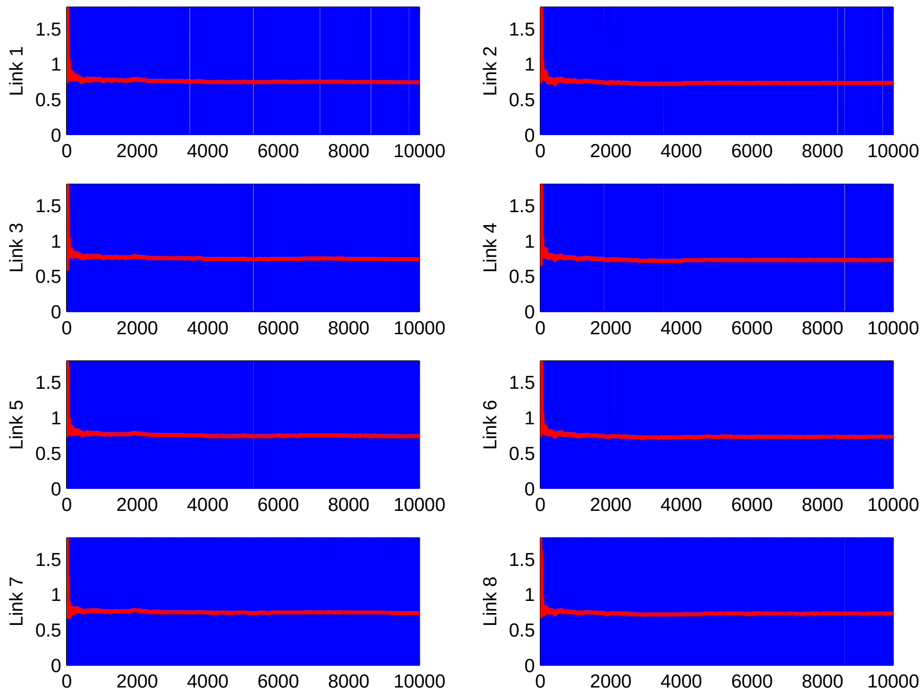

(c) Topo3: Rates

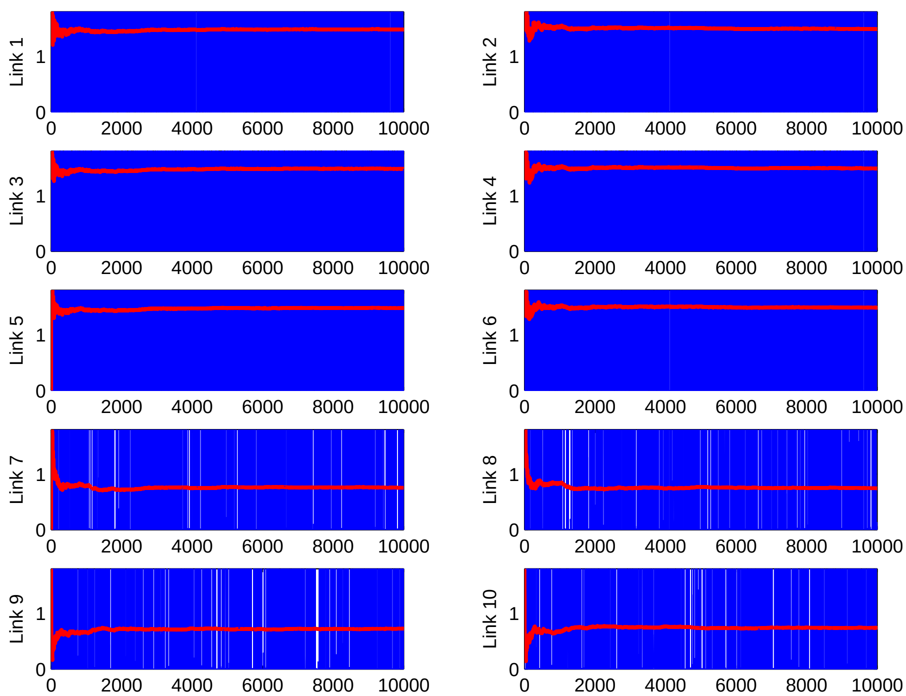

(d) Topo4: Rates

Fig. 7: Link Rates - zooming (cont.).

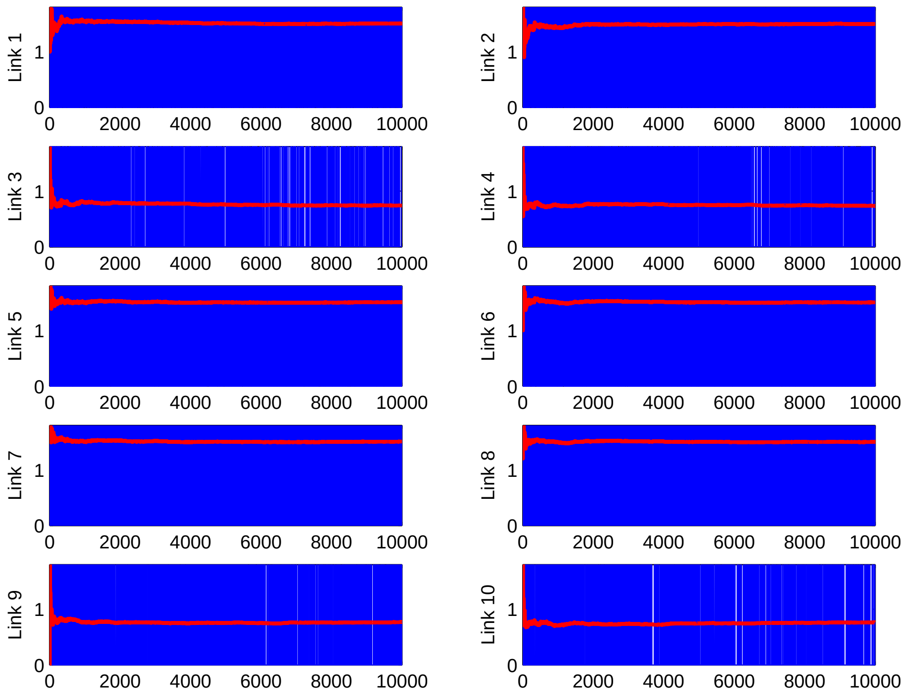

(e) Topo5: Rates

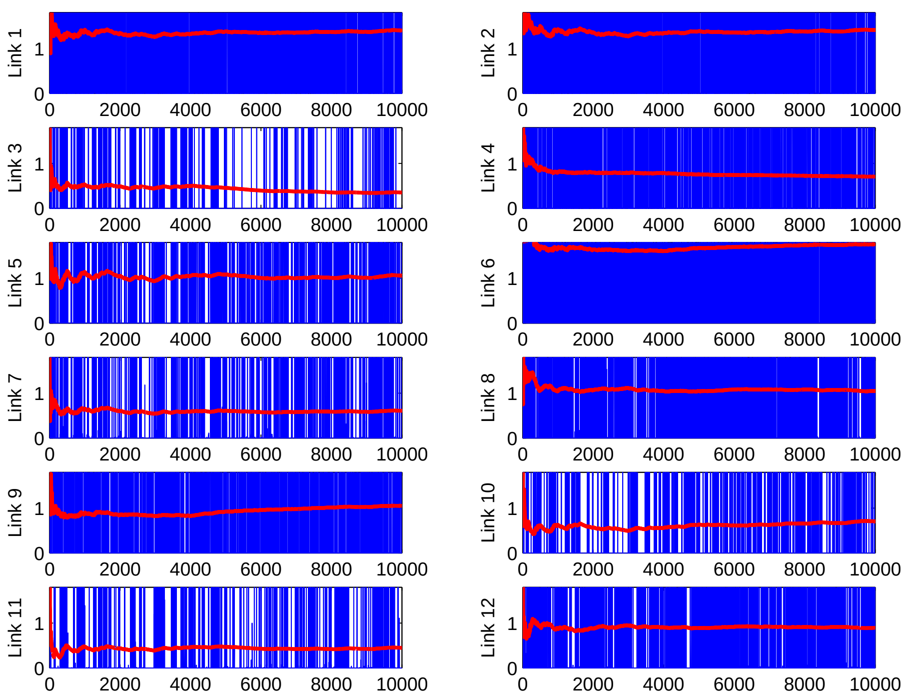

(f) Topo6 ($\delta = 0.1$): Rates

Fig. 7: Link Rates - zooming (cont.).

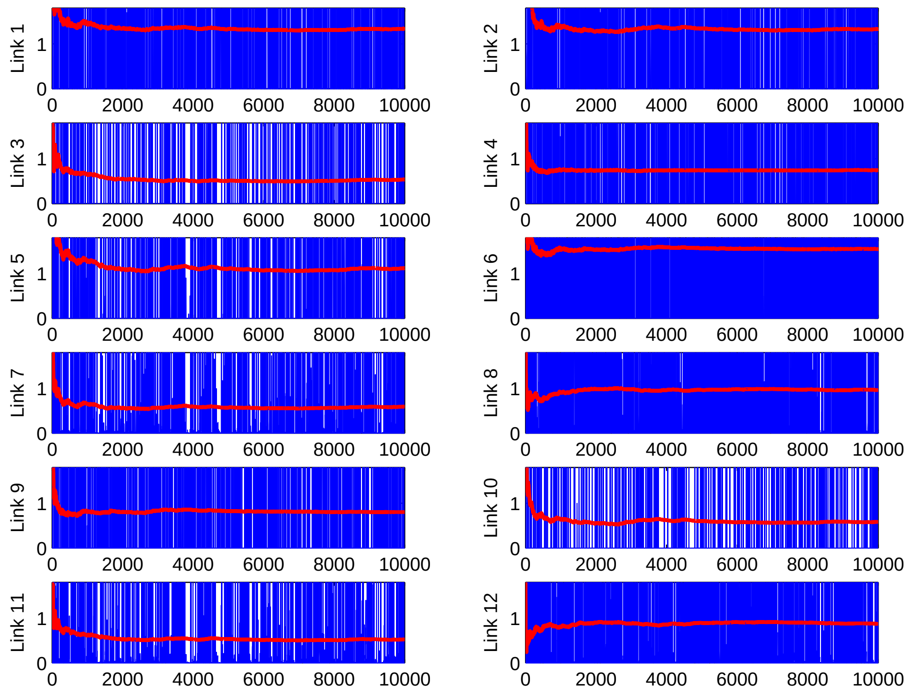

(g) Topo6 ($\delta = 0.5$): Rates

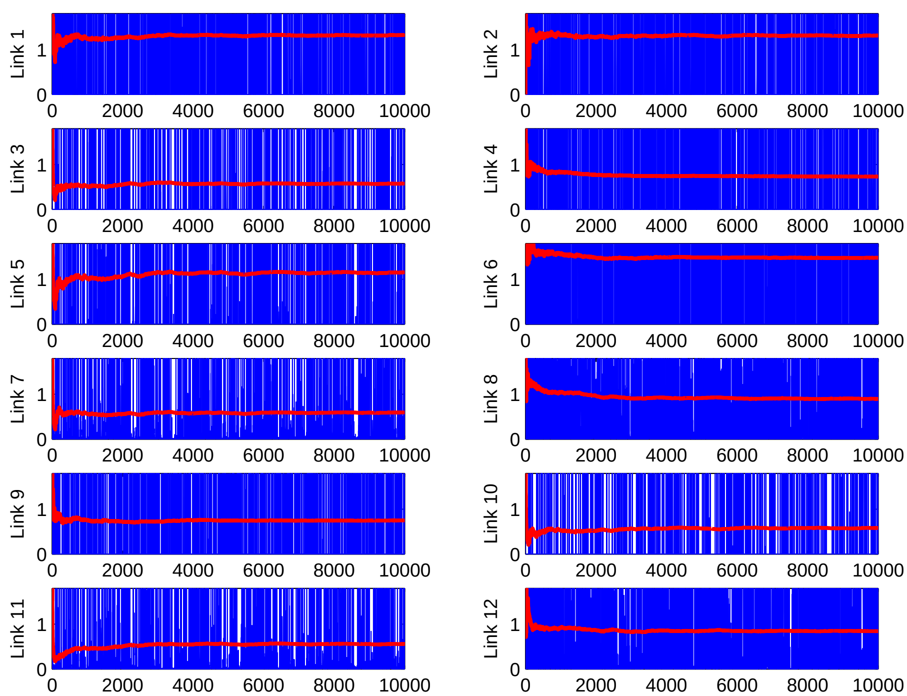

(h) Topo6 ($\delta = 0.9$): Rates

Fig. 7: Link Rates - zooming (cont.).

[6] L. Georgiadis, M.J. Neely, and L. Tassiulas. Resource allocation and cross layer control in wireless networks. *Foundations and Trends in Networking*, 1(1):1–144, 2006.

[7] H.J. Kushner. *Stochastic stability and control*. Acad. Press, 1967.

[8] M. Nagarajan and G. Sosic. Game-theoretic analysis of cooperation among supply chain agents: Review and extensions. *European Journal of Operational Research*, May 2006.

[9] A. Nagurney. Supply chain network design under profit maximization and oligopolistic competition. *Transportation Research E*, (46):281–294, 2010.

[10] Q. Qiang and A. Nagurney. Modeling of supply chain risk under disruptions with performance measurement and robustness analysis. *in Managing Supply Chain Risk and Vulnerability: Tools and Methods for Supply Chain Decision Makers, Editors T. Wu and J. Blackhurst, Springer, Berlin, Germany, 2009*, pages 91–111, 2009.

[11] W.F. Samuelson and S.G. Marks. *Managerial Economics, 4th ed*. Wiley, 2003.

[12] T. Santoso, S. Ahmed, M. Goetschalckx, and A. Shapiro. A stochastic programming approach for supply chain network design under uncertainty. *European Journal of Operational Research*, (167):96–115, 2005.

[13] R. Skulte and T. Wikerson. Critical infrastructure and the supply chain. *The CIP Report on Global Supply Chain*, July 2011.

[14] I. Varkony. Critical infrastructure protection: Volatility and risks in the global chain. *The CIP Report on Global Supply Chain*, July 2011.