



Stability conditions for a decentralised medium access algorithm: single- and multi-hop networks

Seva Shneer¹ · Alexander Stolyar²

Received: 19 October 2018 / Revised: 31 May 2019 / Published online: 11 October 2019
© The Author(s) 2019

Abstract

We consider a decentralised multi-access algorithm, motivated primarily by the control of transmissions in a wireless network. For a finite single-hop network with arbitrary interference constraints, we prove stochastic stability under the natural conditions. For infinite and finite single-hop networks, we obtain broad rate-stability conditions. We also consider symmetric finite multi-hop networks and show that the natural condition is sufficient for stochastic stability.

Keywords Stochastic stability · Rate stability · Infinite network · Queueing networks · Wireless systems · Carrier-sense multiple access · Decentralised medium access protocols · Single-hop networks · Multi-hop networks

Mathematics Subject Classification 60K25 · 68M12

1 Introduction

We consider a model motivated by wireless networks. A key feature of wireless transmissions is that they interfere with each other, especially if the receivers are in close proximity, and this interference may prevent some of the simultaneous transmissions from being received correctly. This creates the need for the design of algorithms regulating the behaviour of transmitters in wireless networks, so that simultaneous interfering transmissions do not occur at all, or occur rarely.

✉ Seva Shneer
v.shneer@hw.ac.uk

Alexander Stolyar
stolyar@illinois.edu

¹ School of MACS, Heriot-Watt University, Edinburgh EH14 4AS, UK

² ISE Department and Coordinated Science Lab, University of Illinois at Urbana-Champaign, Urbana IL 61801, USA

The transmitter–receiver pairs in a network are represented by nodes on a graph, and an edge between two nodes is present if the corresponding transmissions interfere with each other. Thus, the resulting interference graph represents the interference structure of the network.

We consider both single- and multi-hop networks. In a single-hop network, jobs (or messages to be transmitted) arrive at nodes in a network and, upon a successful transmission, leave the network. The dependence between states of different nodes exists because the interference graph imposes constraints on simultaneous transmissions. In a multi-hop network, a message, upon successful transmission at one node, may leave the network or may move to another node, where it needs to be transmitted again. Thus, multi-hop networks add a further layer of complexity, as the states of the nodes are dependent not only due to interference constraints, but in this case also due to message movement between the nodes.

We are interested in stability of a network. In finite networks, by stability we mean the stochastic stability of the nodes' queues. One usually calls a transmission scheduling algorithm maximally stable if it guarantees stability if such is feasible at all, under at least one algorithm. The celebrated BackPressure (sometimes referred to as MaxWeight) algorithm introduced in [21] is maximally stable. It is, however, in all but the simplest network structures, *centralised*, i.e., it requires the presence of a central entity that is aware of the state of the entire network. This is not practical in wireless networks that tend to be large and ever-changing.

There is thus the need for designing *decentralised* algorithms where each node regulates its behaviour on its own, without any global knowledge. In principle, decentralised behaviour may lead to conflicts, when several interfering transmissions will be attempted simultaneously, and the messages will not be received. We consider the so-called CSMA (carrier-sense multiple access, see [9]) networks where each transmitter can sense if a neighbouring node is transmitting and will never initiate an interfering transmission. Conflicts are thus avoided in CSMA networks.

In this paper, it is not our goal to design decentralised algorithms that are maximally stable. We study a different question: what is the stability performance of some simple specific decentralised protocols? We consider the following protocol. Assume that the network is finite. Assume also that time is slotted, i.e., arrivals happen at discrete-time instances denoted $1, 2, \dots$, all transmission times are equal to 1, and transmissions start at the beginning of a unit-long time slot and complete at its end. Assume that at the beginning of each time slot, each message is assigned a random number, drawn independently from some fixed absolutely continuous distribution; the lower this number, the higher the message transmission priority in the slot. Then, a given message is transmitted in a slot, if and only if its priority is the highest among all messages within its neighbourhood, which includes its own node and all adjacent (neighbouring) nodes.

In [17] a related protocol is considered, where nodes, rather than messages, compete for transmission, and where a node may transmit even if its neighbour has a higher priority, provided this neighbour does not transmit due to its other neighbour having an even higher priority. The protocol considered in this paper is significantly more conservative and thus, not surprisingly, leads to a smaller stability region.

On the face of it, the protocol described in the previous paragraph is centralised, as priorities need to be assigned. It is however easily implemented in a decentralised fashion (with an arbitrarily small loss of efficiency), for example as follows. Assume that each time slot is split into two parts: the first part, of duration $0 < \varepsilon < 1$, is devoted to medium access competition, and the second part, of duration $1 - \varepsilon$, is devoted to actual message transmissions. Assume that each message transmission lasts exactly $1 - \varepsilon$. The protocol may be designed as follows: a random time uniformly distributed in $(0, \varepsilon)$ is chosen at the beginning of every time slot for every message currently present in the system, independently of all other messages, and these random variables are also independent over time slots. Once this time expires for a message, an access transmission is initiated. This access transmission does not start an actual transmission but is registered by all nodes in the neighbourhood. At time ε all access transmissions stop. If a node registered that one of its messages' access transmission was the earliest in its neighbourhood, the node will transmit that message. It is easy to see that the loss of efficiency versus the centralised protocol with transmission times of 1 is exactly ε .

In the rest of the paper, for simplicity we consider a “cleaner” version of the protocol, as described in the previous paragraph, that ignores the loss of an ε proportion of the throughput.

The following example illustrates the conservative nature of the protocol. Consider 4 nodes, with interference graph being a “circle”, so that either nodes 1 and 3, or nodes 2 and 4 can transmit simultaneously. Assume that, in a given time slot, the message with the highest priority is located at node 1, the message with the second-highest priority is located at node 2 and the message with the third-highest priority is located at node 3. Under the algorithm considered here, only the first message will be transmitted in this slot, when in fact nodes 1 and 3 could successfully transmit simultaneously. Of course, under other priority orderings, transmission of two messages will occur.

Denote by \mathcal{N}_i the neighbourhood of node i in the interference graph (all neighbours of the node and the node itself). Denote by X_i the number of messages at node i at the beginning of a time slot. It is easy to see that node i will transmit a message with probability

$$\varphi_i = \frac{X_i}{\sum_{j \in \mathcal{N}_i} X_j}. \quad (1)$$

A model closely related to ours has been considered in a recent paper [16] (see also [15]). It is a single-hop network with nodes located on a grid. The model is in *continuous time*, with each message having an exponentially distributed size with unit mean. All messages may transmit simultaneously, with the instantaneous transmission rate depending on the interference from the messages in the neighbourhood. A standard assumption that a message transmission rate is proportional to its signal-to-noise ratio is adopted in [16], which leads to a *node* i transmission rate given exactly by (1). The model in [16] is *symmetric* in that the message arrival rates at all nodes are equal. The authors focus on *infinite-grid* networks and are interested in their stability. The authors define this as the finiteness of the minimal stationary regime for the system starting with all queues being empty (see [16] for more details). They show that a network is

stable in this sense under the natural condition on the message arrival rates. The main tool in their analysis is *monotonicity*, i.e., the property that if one network starts with an initial condition dominating that of another, there exists a coupling preserving this dominance at all times.

In the single-hop scenario, we consider *arbitrary* networks: finite or infinite, arbitrary interference graph, arbitrary arrival intensities. For finite single-hop networks, we prove that the system is stochastically stable if the arrival rates belong to a certain set. For infinite single-hop networks, we obtain a rather broad sufficient condition for the *rate stability*, which is the property that, starting from any fixed initial state, the growth rates of the queues are sub-linear in time.

We also consider finite *multi-hop* networks; here we additionally assume that a network is symmetric: it is a regular graph (recall that a graph is regular if all nodes have the same number of neighbours), with equal exogenous arrival rates at the nodes and with a message path through the network being the standard random walk (until the message leaves the network). In the multi-hop setting, a further complication arises from the fact that monotonicity does not hold. We prove directly that the network is stochastically stable under a natural condition ensuring that, when states of all queues on the network are equal, the average rate at which work arrives is smaller than the average rate at which work is performed by the system. Our approach is *not* based on monotonicity, in either setting.

Our stability proofs use the fluid-limit technique. The discrete-time setting motivating our work and the continuous-time network motivating [16] share the same fluid limits, and thus our results are valid in the continuous-time setting too. Note also that the random variables representing the number of successful transmissions from all nodes in the same time slot in our model are not independent. However, our stability results also apply to a different discrete-time model, where (perhaps rather unrealistically) each station transmits a message with a probability given by (1); again, this is due to the fact that this model has the same fluid-limit dynamics as ours.

Another important concept in wireless networks is utility maximisation. Utility-optimal algorithms are known to guarantee maximal stability for finite single-hop networks, under some assumptions on the utility functions. However, these algorithms are centralised as the *average* service rates φ_i assigned to nodes form a solution $\bar{\varphi} = \{\varphi_i\}$ to a global optimisation problem. An important example of such algorithms is presented by the well-known α -fair algorithms (see [8,12,13] for the introduction of the fair-allocation concepts and [1,6] for stability proofs). In α -fair algorithms, the average rates φ_i are such that

$$\bar{\varphi} \in \arg \max_{\bar{\mu} \in \mathcal{C}} \sum_i X_i \frac{1}{1-\alpha} \left(\frac{\mu_i}{X_i} \right)^{1-\alpha}, \quad \text{when } \alpha > 0, \quad \alpha \neq 1,$$

or

$$\bar{\varphi} \in \arg \max_{\bar{\mu} \in \mathcal{C}} \sum_i X_i \log(\mu_i/X_i), \quad \text{when } \alpha = 1,$$

where \mathcal{C} is some fixed set. In the finite single-hop setting, for a specific (natural) set \mathcal{C} , we provide two proofs of stochastic stability for the arrival rates within \mathcal{C} . The first proof follows from a much more general result for monotone 0-homogeneous service rates, which is of interest on its own and may have other applications. Our infinite-system rate stability proof is based on the same ideas. We provide a second proof for finite systems, as it is based on discovering an important property of the rates (1): they in fact happen to be α -fair in \mathcal{C} , with $\alpha = 2$. We believe this property to be interesting in its own right too as it presents an example of a decentralised protocol which happens to be centrally optimal in a certain sense, and it is known to imply stochastic stability in finite single-hop networks.

It is known that utility-maximising algorithms, with the exception of proportionally fair algorithms (i.e., α -fair algorithms with $\alpha = 1$; see [2,22] for a treatment of some of the multi-hop scenarios), in general do not guarantee stability in multi-hop networks. Therefore, our result on the fairness of the rates (1) does not imply stability in the multi-hop setting. In Sect. 3.2, we present our main result in the finite multi-hop setting, stating that a natural stability property does hold under our algorithm for a class of symmetric networks. Specifically, these networks are such that the interference graph is regular, the exogenous arrival rates to all nodes are equal, and upon a successful transmission a message may either leave the network or move to a neighbour node chosen uniformly at random. We show stochastic stability under a natural condition on the per-node exogenous arrival rate.

The paper is organised as follows. We define our model in Sect. 2 and then present our main results in Sect. 3 (for single-hop networks in Sect. 3.1 and for multi-hop networks in Sect. 3.2). Section 4 contains the necessary fluid-limit constructions, definitions and results. The proofs of our main results are presented in Sect. 5.2 for the finite single-hop case, in Sect. 5.3 for the infinite single-hop case, and in Sect. 6 for the multi-hop case. We discuss some open problems in Sect. 7.

Basic notation. We will use the following notation throughout: \mathbb{R} and \mathbb{R}_+ are the sets of real and real non-negative numbers, respectively; \bar{y} means a (finite- or infinite-dimensional) vector (y_i) ; for a finite-dimensional vector \bar{y} , $\|\bar{y}\| = \sum_i |y_i|$; for a set of functions (f_i) and a vector (y_i) , $\bar{f}(\bar{y})$ denotes the vector $(f_i(\bar{y}))$; vector inequalities are understood component-wise; $\frac{d^+}{dt}$ is the right derivative; $\frac{d^+}{dt} y(t_0) = \liminf_{t \downarrow t_0} \frac{y(t) - y(t_0)}{t - t_0}$ – the lower right Dini derivative; $y(\cdot) = (y(t), t \geq 0)$; we also use the convention that $0/0 = 0$. The indicator function of an event or condition A is denoted by $I(A)$; $\lfloor a \rfloor$ is the largest integer not exceeding a . The abbreviation *r.v.* means *random variable*; *w.p.1* means *with probability 1*; *i.i.d.* means *independent identically distributed*; *u.o.c.* means *uniformly on compact sets*.

2 Model and notation

Denote by \mathcal{V} the set of nodes of a graph, and by \mathcal{E} the set of its edges. The set \mathcal{V} may be finite (in which case we will refer to the network as finite) or countable (in which case we will refer to the network as infinite). Denote by N the (finite or infinite) cardinality of \mathcal{V} .

For a node i , denote by $\mathcal{N}_i = \{j \in \mathcal{V} : (i, j) \in \mathcal{E}\} \cup \{i\}$ its neighbourhood. We assume throughout that \mathcal{N}_i is finite for each i , the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is connected, and that the neighbourhood relationship is symmetric (or that the graph is undirected), i.e., if $i \in \mathcal{N}_j$, then $j \in \mathcal{N}_i$.

Each node has an infinite buffer for storing messages but there is no queue. Time is slotted, and at the beginning of each time slot first transmissions are initiated, and then arrivals happen. Each transmission time is equal to 1.

At the beginning of each time slot, every message in the system is assigned a random number which is drawn, independently of everything else, from a certain fixed absolutely continuous distribution; the smaller this number, the higher the message transmission priority. A message is transmitted if and only if it has the highest priority in its neighbourhood, i.e., node i transmits a message if that message's priority is the maximal over all the messages in \mathcal{N}_i . We refer the reader to the introduction for an explanation of how this may be implemented in a decentralised way, by using a uniform distribution on a small interval, with an arbitrarily small loss of throughput.

At each time slot k , a random number $\xi_i(k)$ new messages arrive at node i . We assume that $\xi_i(k)$ are i.i.d. with $\mathbb{E}(\xi_i(k)) = \lambda_i > 0$.

Throughout the paper, we use the notation

$$\varphi_i(\bar{p}) = \frac{p_i}{\sum_{j \in \mathcal{N}_i} p_j}, \quad \bar{\varphi}(\bar{p}) = (\varphi_i(\bar{p})), \quad (2)$$

where $\bar{p} = (p_i)$ is a vector with finite non-negative components. By convention, $\varphi_i(\bar{p}) = 0$ when $\sum_{j \in \mathcal{N}_i} p_j = 0$.

3 Main results

The results are split into two subsections covering single- and multi-hop networks. The subsection on single-hop networks contains results for both finite and infinite networks.

3.1 Single-hop network

For the single-hop network, we consider any undirected graph \mathcal{G} where each node has a finite neighbourhood. We assume that, upon a successful transmission, a message leaves the system. The evolution of the state of the queue of node i may then be written as

$$X_i(k+1) = X_i(k) + \xi_i(k) - \eta_i(k),$$

where by $X_i(k)$ we denote the state of the queue of node i at time k , and by $\eta_i(k)$ the number of messages leaving node i during the k th time slot. The random variable $\eta_i(k)$ can only take the values 0 and 1, and it is easy to see that, as priorities are chosen independently from the same fixed distribution,

$$\mathbb{P}(\eta_i(k) = 1 | \bar{X}(k) = \bar{X}) = \varphi_i(\bar{X}) = \frac{X_i}{\sum_{j \in \mathcal{N}_i} X_j}.$$

Denote

$$\mathcal{C} = \left\{ \bar{\lambda} : \bar{\lambda} \leq \bar{\varphi}(\bar{p}) \text{ for some } \bar{p} \in \mathbb{R}_+^N \right\}. \tag{3}$$

We will call a finite network *stochastically stable* if the countable Markov chain $\bar{X}(\cdot)$ is positive recurrent.

Theorem 1 *Consider a finite network. If $\bar{\lambda} < \bar{v}$ for some $\bar{v} \in \mathcal{C}$, then the system is stochastically stable.*

The fact that $\bar{\lambda}$ is component-wise strictly smaller than a vector from \mathcal{C} means that there exists at least one configuration of queue states, represented by the corresponding vector \bar{p} from (3), such that for that configuration the amount of work arriving at any node is strictly smaller than the average amount of work this node performs. This is a natural sufficient stability condition. We conjecture that this condition is essentially necessary as well, in that if $\bar{\lambda}$ is strictly outside \mathcal{C} , the process is transient. However, establishing this fact is beyond the scope of the present paper—it may be a subject of future work.

We present a proof of Theorem 1 in Sect. 5.2.

Corollary 1 *In a finite symmetric network, where the graph \mathcal{G} is $(m - 1)$ -regular (so that each node has degree $(m - 1)$) and $\lambda_i = \lambda$ for each i , the condition of Theorem 1 is equivalent to the requirement that $\lambda < 1/m$. In particular, for a network of nodes located on a circle with the same arrival intensity λ at each node, $\lambda < 1/3$ guarantees stability.*

The example of a network of nodes located on a circle illustrates well the fact that, due to the conservative nature of the algorithm, its stability set, $\lambda < 1/3$, is smaller than the maximal stability set $\lambda < 1/2$ achieved by, for instance, MaxWeight. It is also smaller than the set $\lambda < 2/5$ achieved by the less conservative protocol considered in [17].

We now prove Corollary 1. Indeed, if $\lambda < 1/m$, then the vector $(\lambda, \dots, \lambda)$ is component-wise upper-bounded by the vector $(1/m, \dots, 1/m)$ which belongs to the set \mathcal{C} trivially (one needs to take the vector $\bar{p} = (1, \dots, 1)$ to verify this). Assume now that for some \bar{p} the vector $(\lambda, \dots, \lambda)$ is component-wise smaller than a vector $\bar{\varphi}(\bar{p}) \in \mathcal{C}$. Then

$$\frac{1}{\lambda} > \sum_{j \in \mathcal{N}_i} \frac{p_j}{p_i}$$

for each i , and, if we add up these inequalities, we obtain

$$\frac{N}{\lambda} > \sum_i \sum_{j \in \mathcal{N}_i} \frac{p_j}{p_i} = \frac{1}{2} \sum_i \sum_{j \in \mathcal{N}_i} \left(\frac{p_j}{p_i} + \frac{p_i}{p_j} \right) \geq mN,$$

which implies $\lambda < 1/m$.

Remark 1 As our proof of Theorem 1 is based on fluid limits, its results are also valid for a continuous version of the model similar to that of [16]. We refer the reader to the Introduction for an explanation of the connection between models.

A finite or infinite network is called *rate-stable* if, w.p.1,

$$\lim_{k \rightarrow \infty} X_i(k)/k = 0, \quad \forall i,$$

for any initial state $\bar{X}(0)$ with all components being finite, $X_i(0) < \infty$.

Rate stability is a weaker property than stochastic stability. The following result gives a sufficient condition for rate stability.

Theorem 2 *The infinite or finite system is rate-stable if $\bar{\lambda} \leq \bar{\varphi}(\bar{p})$ for some \bar{p} such that $0 < c \leq p_i \leq 1$ for all i .*

A proof of Theorem 2 is given in Sect. 5.3.

3.2 Symmetric multi-hop networks with Geometric service requirements

In this section, we restrict our attention to a finite $(m - 1)$ -regular graph, $m \geq 2$. Assume that the access procedure is the same as before. Now, however, upon service, a message leaves the system with probability $1/k$, and goes to a neighbouring node with probability $(1 - 1/k) \frac{1}{m-1}$. Assume that the arrival rate into each node is λ/k , so that the total workload for each node is λ —this follows from standard rate-balance equations.

One can think of each message needing a Geometric($1/k$) number of successful transmissions to leave the system and, conditionally on not leaving the system upon a successful transmission, performing a simple random walk on the graph, i.e., choosing a neighbouring node uniformly at random.

Theorem 3 *Suppose that the system graph is finite $(m - 1)$ -regular, $m \geq 2$. Then, if $\lambda < 1/m$, the system is stochastically stable.*

We provide a proof of Theorem 3 in Sect. 6.

Remark 2 Similarly to Remark 1, our results in the multi-hop setting hold for a continuous-time version of the model.

4 Fluid limits

4.1 Fluid scaling

Our main results are based on the fluid-limit technique (see [4,14,19]). For the application of this technique to *discrete-time* processes, see, for example, [3,20].

We consider the Markov chain $\bar{X}(k)$, $k = 0, 1, 2, \dots$, and extend it to continuous time with the convention $\bar{X}(t) = \bar{X}(\lfloor t \rfloor)$. Consider a sequence of processes $\bar{X}^{(r)}(\cdot)$, indexed by $r \uparrow \infty$, and their *fluid-scaled* versions

$$\bar{x}^{(r)}(t) = \frac{\bar{X}^{(r)}(rt)}{r}, \quad t \geq 0.$$

4.2 Fluid sample paths

Fluid sample paths (FSP) are defined as possible limits of the realisations of $\bar{x}(\cdot)$, with common “driving” processes’ realisations, satisfying the functional strong law of large numbers. The definition is along the same lines as FSP definitions in other contexts (see, for example, the proof of [20, Theorem 2] for a context somewhat similar to ours). However, it is not quite standard, for two reasons. First, the service process construction (i.e., the procedure determining the transmission schedule in each slot) is not quite standard. Second, we need to define FSPs with a possibly countable set of component functions. For these reasons, we give a formal definition here. To improve the exposition, the definition will be for the single-hop case only. Extension to the multi-hop case is straightforward (we will comment on that later).

We start with specifying a construction of the service process which is consistent with the model definition. Recall that the neighbourhood \mathcal{N}_i of each node i is finite. Suppose the nodes are indexed by $i = 1, 2, \dots$ (This indexing is arbitrary.) Given the queue lengths \bar{X} in a given time slot k , the transmission schedule in this slot is determined recursively as follows. Suppose, the mutual ranking $m(j) \in \mathcal{M}_j$ of the first j nodes is already determined; here \mathcal{M}_j is the set of all $j!$ permutations of $(1, \dots, j)$. The smaller the number (ranking level) of a node, the higher the ranking. (Note that for $j = 1$, the “mutual” ranking $m(1)$ of the single node 1 is deterministic, trivially equal to (1).) The following describes how the mutual ranking $m(j + 1) \in \mathcal{M}_{j+1}$ of the first $j + 1$ nodes is determined. Given $m(j)$, the ranking $m(j + 1)$ is random, but it is consistent with m_j . Namely, node $j + 1$ gets a random ranking level n in $\{1, \dots, j + 1\}$, and those nodes in $\{1, \dots, j\}$, whose ranking level in $m(j)$ was greater or equal to n , will have their ranking level shifted up by 1. The random ranking level n of node $j + 1$ is determined as follows. Associated with each index j and each ranking (permutation) $m \in \mathcal{M}_j$ and each time k , there is an independent r.v. $\kappa(j, m, k)$, uniformly distributed in $[0, 1)$. Let π_n , $n = 1, \dots, j + 1$, be, conditional on $(X_1, \dots, X_j, X_{j+1})$ and the mutual ranking $m \in \mathcal{M}_j$ of the nodes $1, \dots, j$, the probabilities that node $j + 1$ receives the mutual ranking level n among nodes $1, \dots, j + 1$. Then, node $j + 1$ receives mutual ranking level n if

$$\kappa(j, m, k) \in \left[\sum_{\ell=1}^{n-1} \pi_\ell, \sum_{\ell=1}^n \pi_\ell \right).$$

Note that for any node i there exists a finite j such that $\mathcal{N}_i \in \{1, \dots, j\}$. Therefore, even though this mutual ranking determination procedure has an infinite number of steps, i.e., it is a (random) sequence of mutual rankings $m(1), m(2), \dots$, the mutual

rankings within \mathcal{N}_i —and therefore the decision on whether or not i transmits—will be determined within a finite number of steps. We see that the service process is driven by a countable set of i.i.d. r.v. $\kappa(j, k, m)$ for all k, j , and $m \in \mathcal{M}_j$.

The driving processes for the arrivals are natural. Namely, a countable set of r.v. $\xi_i(k)$ for all k and i , giving the number of arrivals in node i at time k . The r.v. $\xi_i(k)$ are independent across all k , and identically distributed across k for each i .

The driving processes are easily seen to satisfy the following functional strong law of large numbers (FLLN) properties. W.p.1.,

$$\forall j \text{ and } \forall m \in \mathcal{M}_j,$$

$$\lim_{r \rightarrow \infty} (1/r) \sum_{k \leq rt} I\{\kappa(j, k, m) \leq u\} = ut, \quad \text{u.o.c. in } (u, t) \in [0, 1] \times [0, \infty), \quad (4)$$

and, $\forall i$,

$$\lim_{r \rightarrow \infty} (1/r) \sum_{k \leq rt} \xi_i(k) = \lambda_i t, \quad \text{u.o.c. in } t \in [0, \infty). \quad (5)$$

We can and do assume that all processes $\bar{X}^{(r)}(\cdot)$ (for all r), and their fluid-scaled versions $\bar{x}^{(r)}(\cdot)$, are constructed on a common probability space, defined by the driving processes $\{\kappa(j, k, m)\}$ and $\{\xi_i(k)\}$.

A vector-function $\bar{x}(t) = (x_i(t))$, $t \geq 0$, is called a *fluid sample path* (FSP), with initial state $\bar{x}(0)$ having finite components $x_i(0) < \infty$, if there exists a sequence $r \rightarrow \infty$ and a *realisation* of the driving processes $\{\kappa(j, k, m)\}$ and $\{\xi_i(k)\}$, such that the FLLN conditions (4) and (5) hold and the corresponding sequence of fluid-scaled process *realisations* is such that

$$x_i^{(r)}(t) \rightarrow x_i(t), \quad \text{u.o.c. in } t \in [0, \infty), \quad \forall i. \quad (6)$$

The above FSP definition is for the single-hop network, where each transmitted message leaves the system. Extending this definition to a multi-hop case is straightforward. An additional driving process $\zeta_i(\ell)$ consisting of uniformly distributed in $[0, 1)$ r.v. is defined: $\zeta_i(\ell)$ determines the routing of the ℓ th message departing node i . This driving process satisfies the FLLN, analogous to (4). An FSP is defined, again, as a limit of fluid-scaled trajectories corresponding to the driving processes' realisations satisfying the FLLN conditions. We omit further details.

We emphasise again that the set of FSP components $x_i(\cdot)$ is finite for a finite system and countable for an infinite one. In either case, it is easy to see that all FSP components $x_i(\cdot)$ are Lipschitz *uniformly in i* .

4.3 Fluid-limit result

Lemma 1 *For either the single- or the multi-hop setting, the following holds. Suppose that for a given sequence $r \rightarrow \infty$, the initial states of the fluid-scaled process are such that $\limsup_{r \rightarrow \infty} x_i^{(r)}(0) < \infty$ for each i . Then, w.p.1., any subsequence of $\bar{x}^{(r)}(\cdot)$*

contains a further subsequence such that $x_i^{(r)}(t) \rightarrow x_i(t)$, u.o.c., for each i , where $\bar{x}(\cdot)$ is an FSP.

The proof of Lemma 1 is very standard and is omitted here. For example, it can follow the same lines as that of [20, Theorem 2], for a model close to ours.

4.4 Stability

For a *finite* system, to establish stochastic stability (positive recurrence) of the Markov chain $\{\bar{X}(k)\}_{k \geq 0}$, it suffices to prove (see [14]) that for some $T > 0$ and $\varepsilon > 0$ any sequence of processes $\bar{X}^{(r)}(\cdot)$, with $\|\bar{X}^{(r)}(0)\| = r$, is such that

$$\limsup_{r \rightarrow \infty} \mathbb{E} \frac{1}{r} \|\bar{X}^{(r)}(rT)\| \leq 1 - \varepsilon.$$

It is a standard result when applying the fluid-limit technique [4, 14, 19] that, for the above to hold, it is sufficient to show that for some $\varepsilon > 0$ and $T > 0$, any FSP with $\|x(0)\| = 1$ is such that

$$\|\bar{x}(T)\| \leq 1 - \varepsilon. \tag{7}$$

For a *finite or infinite* system, to show rate stability, it suffices to prove (see Lemma 6) that any FSP starting from the zero initial state (all $x_i(0) = 0$) stays in the zero state at all times: $x_i(t) = 0, t \geq 0$. This follows from the rate-stability definition and Lemma 1.

5 Single-hop network

This section is devoted to the proofs of our main results for single-hop networks. In Sect. 5.1, we establish further properties of fluid sample paths. Sect. 5.2 is devoted to the proof of Theorem 1, and Sect. 5.3 is devoted to the proof of Theorem 2.

5.1 FSP properties

We start by establishing properties of the FSP dynamics in Lemma 2.

Lemma 2 *Any FSP in the single-hop case satisfies the following conditions:*

$$x_i(t) > 0 \text{ implies } x_i'(t) = \lambda_i - \varphi_i(\bar{x}(t)), \text{ for almost all } t \geq 0, \tag{8}$$

$$\left[x_i(t) = 0 \text{ and } \sum_{j \in \mathcal{N}_i} x_j(t) > 0 \right] \text{ implies } \frac{d^+}{dt} x_i(t) = \lambda_i. \tag{9}$$

In particular, by property (9), any FSP is such that, for a fixed i , $\sum_{j \in \mathcal{N}_i} x_j(t) > 0$ implies that $x_i(\tau) > 0$ for all $\tau > t$ sufficiently close to t . Furthermore, if the network

is finite, any FSP is such that $\sum_i x_i(t) > 0$ implies that $x_i(\tau) > 0, \forall i$, for all $\tau > t$ sufficiently close to t .

A proof of Lemma 2 may be given following the same lines as that of [20, Theorem 2], for a model close to ours, and we omit it here.

5.2 Proof of Theorem 1

We present two different proofs of Theorem 1. One proof of (7), and then of Theorem 1, follows from the following much more general result.

For a function $\bar{\psi} = \bar{\psi}(\bar{p})$, mapping a finite-dimensional positive orthant $\mathbb{R}_+^N, N < \infty$, into itself, define

$$\mathcal{D} = \left\{ \bar{\lambda} \in \mathbb{R}_+^N : \bar{\lambda} \leq \bar{\psi}(\bar{p}) \text{ for some } \bar{p} \right\}. \tag{10}$$

Lemma 3 Consider a family of Lipschitz trajectories $\bar{x}(t), t \geq 0$, in $\mathbb{R}_+^N, N < \infty$, which satisfy the following conditions:

$$x_i(t) > 0 \text{ implies } x_i'(t) = \lambda_i - \psi_i(\bar{x}(t)), \text{ for almost all } t \geq 0, \text{ for any } i, \tag{11}$$

$$\sum_i x_i(t) > 0 \text{ implies } x_i(\tau) > 0, \text{ for all } i, \text{ for all } \tau > t \text{ sufficiently close to } t, \tag{12}$$

where the function $\bar{\psi}$ is such that:

- (A) each ψ_i is non-increasing in x_j for all $j \neq i$;
- (B) each ψ_i is 0-homogeneous, i.e., $\psi_i(s\bar{x}) = \psi_i(\bar{x})$ for all $s > 0$ and for all \bar{x} .

Assume that $\bar{\lambda}$ is such that $\bar{\lambda} < \bar{v}$ for some $\bar{v} \in \mathcal{D}$. Then for any constants $0 < \delta < K < \infty$, there exists $T > 0$ such that, for any such trajectory with $\|x(0)\| = K$,

$$\|x(T)\| \leq \delta.$$

Proof of Lemma 3 We now prove Lemma 3. Fix a vector \bar{p} such that

$$\bar{v} \leq \bar{\psi}(\bar{p})$$

for every i . Note that, as $\lambda_i > 0$ for each i , necessarily $p_i > 0$ for each i . Consider the function

$$F(\bar{y}) = \max_i \left(\frac{y_i}{p_i} \right).$$

For ease of notation, in the rest of the proof we drop the index t and make the dependence of $\bar{\psi}$ on $\bar{x}(t)$ implicit.

Denote

$$\mathcal{K} = \left\{ k : k \in \arg \max_i \left(\frac{x_i}{p_i} \right) \right\}.$$

The function $\max_i \{x_i / p_i\}$ is Lipschitz, because all $x_i(\cdot)$ are Lipschitz. The time points t , where the derivatives of all x_i and of $\max_i \{x_i / p_i\}$ exist, are called regular. Almost all points (with respect to Lebesgue measure) are regular. Then, due to [5, Lemma 2.8.6], $\left(\frac{x_k}{p_k}\right)' = \left(\frac{x_l}{p_l}\right)'$ at any regular point of F for any $k, l \in \mathcal{K}$. The derivative of the function F at a regular point is then

$$(F(\bar{x}))' = \frac{1}{p_k}(\lambda_k - \psi_k)$$

with an arbitrary $k \in \mathcal{K}$. Note that, as $\frac{x_k}{p_k} \geq \frac{x_j}{p_j}$ for any $k \in \mathcal{K}$ and for any j , due to property (A),

$$\begin{aligned} \psi_k &= \psi_k(\bar{x}) \geq \psi_k\left(\frac{x_k p_1}{p_k}, \dots, \frac{x_k p_{k-1}}{p_k}, x_k, \frac{x_k p_{k+1}}{p_k}, \dots, \frac{x_k p_N}{p_k}\right) \\ &= \psi_k\left(\frac{x_k p_1}{p_k}, \dots, \frac{x_k p_{k-1}}{p_k}, \frac{x_k p_k}{p_k}, \frac{x_k p_{k+1}}{p_k}, \dots, \frac{x_k p_N}{p_k}\right) \\ &= \psi_k(\bar{p}) \geq v_k, \end{aligned}$$

where in the last step we used property (B).

Noting that there exists $\varepsilon > 0$ such that $\lambda_i < v_i - \varepsilon$ for every i , we obtain

$$(F(\bar{x}))' = \frac{1}{p_k}(\lambda_k - v_k) + \frac{1}{p_k}(v_k - \psi_k) < -\frac{\varepsilon}{p_k}.$$

This implies that $(F(\bar{x}))'$ is negative and bounded away from 0, as long as $F(\bar{x})$ is positive and bounded away from 0. This concludes the proof of Lemma 3. \square

Theorem 1 follows from Lemma 3 as the rates $\bar{\varphi}$ clearly satisfy conditions (A) and (B).

Remark 3 Lemma 3 is rather general and relates to so-called cooperative dynamical systems (see, for example, [7,18]). We believe that this result is interesting on its own and may have other applications. This result also allows us, in an obvious fashion, to obtain stability conditions for networks with a more general notion of neighbourhood considered in [16]. For other specific examples when cooperative dynamic systems arise in the analysis of scheduling in communication systems, see, for example, [10,11].

We also present a different proof of Theorem 1, which is specific to our model and is based on a global optimality of the rates $\bar{\varphi}$. We think that this optimality is interesting on its own as it is an important structural property of the rates, and as it provides an example of a decentralised algorithm which maximises a global utility function. We also present a simple proof showing stability of algorithms maximising utility functions over a set, without requiring convexity of the set.

The remainder of this second proof consists of the following steps, which correspond to two lemmas below:

1. We show that the FSPs of this model are such that the “service rates” the nodes receive are utility maximising (in fact 2-fair) in the set \mathcal{C} . This property is proved in Lemma 4;
2. The property (7) of the FSPs (which is sometimes referred to as stability of FSPs), and hence the stability of the underlying Markov chain, follows from the utility-maximisation property of the “service rates.” We only need to note that this fact is usually proved for convex sets of possible rates, whereas our set \mathcal{C} is not convex. However, convexity is in fact not needed in the stability proof, and we provide a proof for any sets, based on the proof of [20, Theorem 2]; this is done in Lemma 5.

Lemma 4 For any \bar{x} with $x_i > 0$ for all i , the rates $\bar{\varphi} = \bar{\varphi}(\bar{x})$ are 2-fair in the set \mathcal{C} (see the Introduction of this paper or, for example, [1] for definition of α -fairness).

Proof of Lemma 4 Indeed, due to the definition of the set \mathcal{C} , for any $\bar{\mu} \in \mathcal{C}$,

$$\sum_i x_i \left(\frac{\mu_i}{x_i}\right)^{-1} \geq \sum_i x_i \left(\frac{p_i}{(\sum_{j \in \mathcal{N}_i} p_j)x_i}\right)^{-1}$$

for the corresponding vector \bar{p} . Hence, it is sufficient to show that

$$\sum_i x_i \left(\frac{\varphi_i}{x_i}\right)^{-1} \leq \sum_i x_i \left(\frac{p_i}{(\sum_{j \in \mathcal{N}_i} p_j)x_i}\right)^{-1}$$

for all vectors \bar{p} .

Note that the LHS of the above is equal to $\sum_i x_i \sum_{j \in \mathcal{N}_i} x_j = \sum_i x_i^2 + \sum_i \sum_{j \in \mathcal{N}_i, j \neq i} x_i x_j$. Consider now

$$\begin{aligned} \sum_i x_i \left(\frac{p_i}{(\sum_{j \in \mathcal{N}_i} p_j)x_i}\right)^{-1} &= \sum_i x_i^2 \left(1 + \sum_{j \in \mathcal{N}_i, j \neq i} \frac{p_j}{p_i}\right) \\ &= \sum_i x_i^2 + \frac{1}{2} \sum_i \sum_{j \in \mathcal{N}_i, j \neq i} \left(x_i^2 \frac{p_j}{p_i} + x_j^2 \frac{p_i}{p_j}\right). \end{aligned}$$

For any i and j ,

$$x_i^2 \frac{p_j}{p_i} + x_j^2 \frac{p_i}{p_j} \geq 2x_i x_j,$$

and equality is possible if and only if $x_i^2 \frac{p_j}{p_i} = x_j^2 \frac{p_i}{p_j}$, which is equivalent to $\frac{p_i}{x_i} = \frac{p_j}{x_j}$. Therefore, we obtain

$$\sum_i x_i \left(\frac{\sum_{j \in \mathcal{N}_i} p_j}{p_i x_i}\right)^{-1} \geq \sum_i x_i^2 + \sum_i \sum_{j \in \mathcal{N}_i, j \neq i} x_i x_j,$$

and equality is possible if and only if $\frac{p_i}{x_i} = \frac{p_j}{x_j}$ for all i and j . This implies that $\frac{p_i}{x_i}$ has to be a constant for each i , as the graph is connected. This concludes the proof of Lemma 4. \square

The result of Theorem 1 now follows from stability of FSPs under α -fair algorithms (see [1,6]). One only needs to note that such proofs are usually given for convex sets, but convexity is not in fact needed, and stability may be proved following the lines of the proof of Theorem 2 in [20, Section 8], where it was given in the case $\alpha = 1$. The proof for far more general rate allocations is essentially the same and we provide it here for completeness.

Lemma 5 *Let \mathcal{C} be a compact coordinate-convex subset of \mathbb{R}_+^N . Let $h_i : [0, \infty) \rightarrow \mathbb{R}$, for each i , be an increasing, differentiable, concave function (the case when $h_i(y) \downarrow -\infty$ as $y \downarrow 0$ is allowed). Let $g_i : [0, \infty) \rightarrow [0, \infty)$, for each i , be a continuous non-decreasing function such that $g_i(0) \geq 0$ and $g_i(y) > 0$ for $y > 0$.*

Consider a family of Lipschitz trajectories $\bar{x}(t)$, $t \geq 0$, in \mathbb{R}_+^N , which satisfy the following conditions:

$$x_i(t) > 0 \text{ implies } x'_i(t) = \lambda_i - \psi_i(\bar{x}(t)), \text{ for almost all } t \geq 0, \tag{13}$$

$$\sum_i x_i(t) > 0 \text{ implies } x_i(\tau) > 0, \text{ for all } i, \text{ for all } \tau > t \text{ sufficiently close to } t, \tag{14}$$

where the rates $\bar{\psi}$ satisfy

$$\bar{\psi} \in \arg \max_{\mu \in \mathcal{C}} \sum_i g_i(x_i) h_i(\mu_i). \tag{15}$$

Assume that $\bar{\lambda}$ is such that $\bar{\lambda} < \bar{v}$ for some $\bar{v} \in \mathcal{C}$. Then, for any constants $0 < \delta < K < \infty$, there exists $T > 0$ such that, for any such trajectory with $\|x(0)\| = K$,

$$\|x(T)\| \leq \delta.$$

Proof of Lemma 5 In this proof, we will drop the index t for ease of notation. We will also write simply ψ_i , with its dependence on $\bar{x}(t)$ being implicit.

Property (14) implies that $x_i(t) > 0$, for all i , for $t \in (0, \theta)$, where θ is the first time (if any) when all $x_i(t)$ “hit” 0 simultaneously. Consider a trajectory $\bar{x}(\cdot)$ in this interval $(0, \theta)$.

Note that there exists $\varepsilon > 0$ such that $\lambda_i < v_i - \varepsilon$ for each i . Denote

$$F(\bar{v}) = \sum_{i=1}^N G_i(v_i) h'_i(v_i),$$

with $G_i(z) = \int_0^z g_i(s)ds$, and note that

$$\begin{aligned}
 (F(\bar{x}))' &= \sum_{i=1}^N h'_i(v_i)g_i(x_i)(\lambda_i - \psi_i) \\
 &= \sum_{i=1}^N h'_i(v_i)g_i(x_i)(\lambda_i - v_i) + \sum_{i=1}^N h'_i(v_i)g_i(x_i)(v_i - \psi_i) \\
 &< -\varepsilon \sum_{i=1}^N h'_i(v_i)g_i(x_i) + \sum_{i=1}^N h'_i(v_i)g_i(x_i)(v_i - \psi_i). \tag{16}
 \end{aligned}$$

As $\bar{v} \in C$, due to (15),

$$\sum_{i=1}^N g_i(x_i)h_i(\psi_i) \geq \sum_{i=1}^N g_i(x_i)h_i(v_i).$$

Using this and the concavity of the functions h_i , we have

$$0 \leq \sum_{i=1}^N g_i(x_i) (h_i(\psi_i) - h_i(v_i)) \leq \sum_{i=1}^N g_i(x_i)h'_i(v_i)(\psi_i - v_i).$$

This, together, with (16), implies that

$$(F(\bar{x}))' < -\varepsilon \sum_{i=1}^N h'_i(v_i)g_i(x_i).$$

We can now conclude that $(F(\bar{x}))'$ is negative and bounded away from 0, as long as $F(\bar{x})$ is positive and bounded away from 0. Indeed, assume $F(\bar{x}) \geq c_1 > 0$. Then there exists i such that $h'_i(v_i)G_i(x_i) \geq c_1/N$ and hence $G_i(x_i) \geq c_2 > 0$ with obvious notation for c_2 . Therefore, $x_i \geq G_i^{-1}(c_2)$, where G_i^{-1} is the inverse of G_i , which is a well-defined strictly increasing function due to our conditions on the functions g_i . Then

$$(F(\bar{x}))' < -\varepsilon \sum_{i=1}^N h'_i(v_i)g_i(x_i) \leq -\varepsilon h'_i(v_i)g_i(x_i) \leq -\varepsilon h'_i(v_i)g_i(G_i^{-1}(c_2)) = c_3 < 0,$$

as g_i is non-decreasing. This concludes the proof. □

It is easy to see that the condition (9) implies (14) as the graph \mathcal{G} is connected. Indeed, assume that $x_i(t) = 0$ for some i but $\sum_i x_i(t) > 0$. Let us first assume that there exists $j \in \mathcal{N}_i$ such that $x_j(t) > 0$. Then, due to condition (9), $x_i(\tau) > 0$ for all τ sufficiently close to t . If $x_j(t) = 0$ for all $j \in \mathcal{N}_i$, due to the fact that \mathcal{G} is connected, there exists v such that $x_v(t) > 0$ and such that there exists a path

$i = v_1, v_2, \dots, v_l = v$ between nodes i and v (i.e., there is an edge between v_d and v_{d+1} for all $d = 1, \dots, l - 1$). We can then use the argument above to show that $x_{v_d}(\tau) > 0$ for all τ sufficiently close to t , for all d .

The result of Theorem 1 now follows if we take $g_i(y) = y^2$ and $h_i(y) = -y^{-1}$.

5.3 Proof of Theorem 2

Recall that Theorem 2 is for the infinite, as well as finite, systems. It suffices to prove the following.

Lemma 6 Any FSP starting from the zero initial state, $x_i(0) = 0$ for all i , stays in the zero state at all times: $x_i(t) = 0, t \geq 0$, for all i .

Proof of Lemma 6 Consider any FSP with the zero initial state. Denote $s(t) = \sup_i x_i(t)/p_i$. This function is Lipschitz, because all $x_i(\cdot)$ are uniformly Lipschitz and all $0 < c \leq p_i \leq 1$. Time points t where all derivatives $x'_i(t)$ and $s'(t)$ exist are called regular. Almost all points t (with respect to Lebesgue measure) are regular. We will show that at any regular point $t, s'(t) \leq 0$. This will imply that $s(\cdot)$ cannot escape from 0. Suppose not, and, at some regular point $t, s'(t) = \eta > 0$. If this is true, then there exists a positive function $\delta_1 = \delta_1(\delta) \downarrow 0$ as $\delta \downarrow 0$, such that the following hold for any sufficiently small $\delta > 0$: (a) there exists i such that the increment

$$x_i(t + \delta)/p_i - x_i(t)/p_i \geq (\eta/2)\delta, \tag{17}$$

(b) $|x_i(\xi)/p_i - s(t)| < \delta_1$ for all $\xi \in [t, t + \delta]$, (c) $x_j(\xi)/p_j - s(t) < \delta_1$ for all $j \in \mathcal{N}_i$ and all $\xi \in [t, t + \delta]$. If we consider such an i , we observe that for any regular $\xi \in [t, t + \delta]$, thanks to properties (b) and (c),

$$\varphi_i(\bar{x}) > \frac{(s(t) - \delta_1)p_i}{\sum_{j \in \mathcal{N}_i} (s(t) + \delta_1)p_j} = \frac{s(t) - \delta_1}{s(t) + \delta_1} \varphi_i(\bar{p}).$$

Hence, as $\lambda_i < \varphi_i(\bar{p})$, we see that $x'_i(\xi)/p_i \leq \epsilon = \epsilon(\delta_1)$, where $\epsilon(\delta_1) \downarrow 0$ as $\delta_1 \downarrow 0$. Therefore, for a sufficiently small δ and a corresponding $i, x'_i(\xi)/p_i \leq \eta/3$ for all regular $\xi \in [t, t + \delta]$. This contradicts (17). \square

6 Symmetric multi-hop networks with Geometric service requirements: Proof of Theorem 3

As in the single-hop case, we first present a lemma on the conditions any FSP satisfies.

Lemma 7 Any FSP for a multi-hop symmetric network satisfies the following conditions:

$$x_i(t) > 0 \text{ implies}$$

$$x'_i(t) = \lambda_i - \varphi_i(\bar{x}(t)) + \frac{1 - 1/k}{m - 1} \sum_{j \in \mathcal{N}_i, j \neq i} \varphi_j(\bar{x}(t)), \text{ for almost all } t \geq 0, \quad (18)$$

$$\left[x_i(t) = 0 \text{ and } \sum_{j \in \mathcal{N}_i} x_j(t) > 0 \right] \text{ implies } \frac{d^+_i}{dt} x_i(t) \geq \lambda_i. \quad (19)$$

In particular, by property (19), any FSP for a finite network is such that $\sum_i x_i(t) > 0$ implies that $x_i(\tau) > 0$, $\forall i$, for all $\tau > t$ sufficiently close to t .

Once again, a proof of Lemma 7 may be given following the same lines as that of [20, Theorem 2], for a model close to ours, and we omit it here. The last term on the RHS of (18) may be explained as follows: from each $j \in \mathcal{N}_i$, $j \neq i$ messages leave at rate φ_j , a proportion $1 - 1/k$ of these do not leave the system, and a further proportion $1/(m - 1)$ of those that do not leave the system choose node i as their destination.

We will omit the index t in the remainder of the proof. Fix $\varepsilon > 0$ such that $\lambda + \varepsilon < 1/m$ and consider the Lyapunov function $\frac{1}{2} \sum_i x_i^2$, whose drift is

$$\begin{aligned} & \sum_i x_i \left(\frac{\lambda}{k} - \varphi_i + \sum_{j \in \mathcal{N}_i, j \neq i} (1 - 1/k) \frac{1}{m - 1} \varphi_j \right) \\ &= \frac{\lambda}{k} \sum_i x_i - \frac{1}{k} \sum_i x_i \varphi_i + \sum_i x_i \left(-(1 - 1/k) \varphi_i + \sum_{j \in \mathcal{N}_i, j \neq i} (1 - 1/k) \frac{1}{m - 1} \varphi_j \right) \\ &< -\frac{\varepsilon}{k} \sum_i x_i + \frac{1}{k} \left(\frac{1}{m} \sum_i x_i - \sum_i x_i \varphi_i \right) \\ &+ (1 - 1/k) \sum_i \varphi_i \left(-x_i + \frac{1}{m - 1} \sum_{j \in \mathcal{N}_i, j \neq i} x_j \right) \\ &= -\frac{\varepsilon}{k} \sum_i x_i + \frac{1}{k} \left(\frac{1}{m} \sum_i x_i - \sum_i x_i \varphi_i \right) \\ &+ (1 - 1/k) \sum_i \varphi_i \left(-x_i + \frac{1}{m - 1} \left(\frac{x_i}{\varphi_i} - x_i \right) \right) \\ &= -\frac{\varepsilon}{k} \sum_i x_i + \frac{1}{k} \left(\frac{1}{m} \sum_i x_i - \sum_i x_i \varphi_i \right) \\ &+ (1 - 1/k) \left(-\frac{m}{m - 1} \sum_i x_i \varphi_i + \frac{1}{m - 1} \sum_i x_i \right) \\ &= -\frac{\varepsilon}{k} \sum_i x_i + \left(\frac{1}{k} + (1 - 1/k) \frac{m}{m - 1} \right) \left(\frac{1}{m} \sum_i x_i - \sum_i x_i \varphi_i \right) \end{aligned}$$

$$= -\frac{\varepsilon}{k} \sum_i x_i + \frac{1}{m-1} \left(m - \frac{1}{k}\right) \left(\frac{1}{m} \sum_i x_i - \sum_i x_i \varphi_i\right).$$

Let $S = \sum_i x_i$. Note that

$$\frac{1}{S} \sum_i x_i \varphi_i = \sum_i \frac{x_i}{S} \frac{1}{\frac{\sum_{j \in \mathcal{N}_i} x_j}{x_i}} \geq \frac{1}{\sum_i \frac{x_i}{S} \frac{\sum_{j \in \mathcal{N}_i} x_j}{x_i}} = \frac{S}{\sum_i \sum_{j \in \mathcal{N}_i} x_j} = \frac{1}{m},$$

due to convexity of the function $1/x$. It now follows that

$$\left(\frac{1}{2} \sum_i x_i^2\right)' < -\frac{\varepsilon}{k} \sum_i x_i.$$

From here the FSP property (7), and then Theorem 3, follows. □

7 Open problems

Our main result in the multi-hop setting only concerns symmetric (in terms of arrival intensities as well as routing) networks. We expect similar results to hold in greater generality. An interesting example is the following: assume that the graph is a circle, arrival intensities into each node are constant and equal to λ/k but the routing is not symmetric. Upon successful transmission, a message leaves the system with probability $1/k$ or goes to its neighbour on the right with probability $1 - 1/k$. Rate-balance equations imply that the total workload of each node is λ , and we expect that the condition $\lambda < 1/3$ guarantees stability in this model, as well as in the model with symmetric routing covered by Theorem 3. In fact, we conjecture that the same Lyapunov function as the one used in the proof of Theorem 3 has a negative drift in this scenario as well. Simple calculus shows that to prove this, one needs to show that

$$\sum_{i=1}^N \frac{x_i(x_i - x_{i+1})}{x_{i-1} + x_i + x_{i+1}} \geq 0$$

for all vectors \bar{x} , with the conventions that $x_0 = x_N$ and $x_{N+1} = x_1$.

We have ample numerical evidence in support of this hypothesis but currently lack a proof. If the inequality above is proved, it will imply, furthermore, that the condition $\lambda < 1/3$ guarantees stability on the circle topology with arrival intensities equal to λ/k for each node, every message leaving the system upon a successful transmission with probability $1/k$ and for arbitrary (but the same for all nodes) routing to neighbours if the message does not leave the system.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution,

and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Bonald, T., Massoulié, L.: Impact of fairness on Internet performance. *ACM SIGMETRICS Perform. Eval. Rev.* **29**, 82–91 (2001)
2. Bramson, M., D'Auria, B., Walton, N.: Proportional switching in FIFO networks (2017). *Oper. Res.* **65**(2), 469–513 (2017)
3. Coffman Jr., E.G., Stolyar, A.L.: Bandwidth packing. *Algorithmica* **29**, 70–88 (2001)
4. Dai, J.G.: On the positive Harris recurrence for open multiclass queueing networks: a unified approach via fluid limit models. *Ann. Appl. Probab.* **5**, 49–77 (1995)
5. Dai, J.G.: *Stability of Fluid and Stochastic Processing Networks*, vol. 9. MaPhySto Miscellanea Publication, Aarhus (1999)
6. De Veciana, G., Konstantopoulos, T., Lee, T.: Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Trans. Netw.* **9**, 2–14 (2001)
7. Hirsch, M.W.: Systems of differential equations that are competitive and cooperative: convergence almost everywhere. *SIAM J. Math. Anal.* **16**, 423–439 (1985)
8. Kelly, F.P., Maulloo, A.K., Tan, D.K.H.: Rate control for communication networks: shadow prices, proportional fairness and stability. *J. Oper. Res. Soc.* **49**, 237–252 (1998)
9. Kleinrock, L., Tobagi, F.: Packet switching in radio channels: part I-carrier sense multiple-access modes and their throughput-delay characteristics. *IEEE Trans. Comm.* **23**(12), 1400–1416 (1975)
10. Kushner, H., Whiting, P.: Convergence of proportional-fair sharing algorithms under general conditions. *IEEE Trans. Wirel. Commun.* **3**, 1250–1259 (2004)
11. Li, B., Srikant, R.: Queue-proportional rate allocation with per-link information in multihop wireless networks. *Queueing Syst.* **83**, 329–359 (2016)
12. Mo, J., Walrand, J.: Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Netw.* **8**, 556–567 (2000)
13. Roberts, J., Massoulié, L.: Bandwidth sharing and admission control for elastic traffic. *Telecommun. Syst.* **15**, 185–201 (2000)
14. Rybko, A.N., Stolyar, A.L.: Ergodicity of stochastic processes describing the operation of open queueing networks. *Probl. Inf. Transm.* **28**, 199–220 (1992)
15. Sankararaman, A., Baccelli, F.: Spatial birth-death wireless networks. *IEEE Trans. Inf. Theory* **63**(6), 3964–3982 (2017)
16. Sankararaman, A., Baccelli, F., Foss, S.: Interference queueing networks on grids (2019). [arXiv:1710.09797v4](https://arxiv.org/abs/1710.09797v4)
17. Shneer, S., Stolyar, A.L.: Stability conditions for a discrete-time decentralised medium access algorithm. *Ann. Appl. Probab.* **28**(6), 3600–3628 (2018)
18. Smith, H.L.: *Monotone Dynamical Systems: An Introduction to Competitive and Cooperative Systems*, vol. 41. American Mathematical Society, AMS Math. Surveys and Monographs, Providence (1995)
19. Stolyar, A.L.: On the stability of multiclass queueing networks: a relaxed sufficient condition via limiting fluid processes. *Markov Process. Relat. Fields* **1**(4), 491–512 (1995)
20. Stolyar, A.L.: Dynamic distributed scheduling in random access networks. *J. Appl. Probab.* **45**(2), 297–313 (2008)
21. Tassioulas, L., Ephremides, A.: Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Autom. Control* **37**(12), 1936–1948 (1992)
22. Walton, N.: Concave switching in single and multihop networks. *Queueing Syst.* **81**(2), 265–299 (2015)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.