Mean-field limits for multi-hop random-access networks

Mean-field limits for multi-hop random-access networks

F. Cecchi[a], P.M. Van de Ven[b], S. Shneer[c]
[a]Eindhoven University of Technology, Eindhoven, NL
[b]Centrum Wiskunde & Informatica (CWI), Amsterdam, NL
[c]Heriot-Watt University, Edinburgh, UK
f.cecchi@tue.nl, P.M.van.de.Ven@cwi.nl, v.shneer@hw.ac.uk

ABSTRACT

Recent years have seen wireless networks increasing in scale, interconnecting a vast number of devices over large areas. Due to their size these networks rely on distributed algorithms for control, allowing each node to regulate its own activity. A popular such algorithm is Carrier-Sense Multi-Access (CSMA), which is at the core of the well-known 802.11 protocol. Performance analysis of CSMA-based networks has received significant attention in the research literature in recent years, but focused almost exclusively on saturated networks where nodes always have packets available.

However, one of the key features of emerging large-scale networks is their ability to transmit packets across large distances via multiple intermediate nodes (multi-hop). This gives rise to vastly more complex dynamics, and to phenomena not captured by saturated models. Consequently, performance analysis of multi-hop random-access networks remains elusive. Based on the observation that emerging multi-hop networks are typically dense and contain a large number of nodes, we consider the mean-field limit of multi-hop CSMA networks. We show that the equilibrium point of the resulting initial value problem provides a remarkably accurate approximation for the pre-limit stochastic network in stationarity, even for sparse networks with few nodes. Using these equilibrium points we investigate the performance of linear networks under different back-off rates, which govern how fast each node transmits. We find the back-off rates which provide the best end-to-end throughput and network robustness, and use these insights to determine the optimal back-off rates for general networks. We confirm numerically the resulting performance gains compared to the current practice of assigning all nodes the same back-off rate.

Keywords

CSMA; Mean-field limits; Multi-hop networks; Random-access

1. INTRODUCTION

Wireless devices are increasingly part of every aspect of our lives, and are expected to spread even more, being at the heart of the so-called Internet of Things (IoT) [9]. The resulting increase in the number of wireless devices requires careful sharing of the available medium, in order to mitigate interference. This is done by the Medium access control (MAC) mechanism, which establishes the nodes that may transmit at any instant. Due to their size these networks rely on distributed algorithms for control, allowing each node to regulate its own activity.

A popular distributed MAC algorithm is CSMA (Carrier-Sense Multiple-Access), which underlies the IEEE 802.11 and 802.15.4 standards. The popularity of CSMA is due to its simplicity and efficiency, and its key features are its back-off and sensing mechanisms. Each node needs to obey a random back-off period between successive transmissions, and collisions are avoided by freezing the back-off process whenever the node senses potential interference.

Due to its widespread use, the CSMA algorithm has been extensively studied in the research literature in recent years. Most of this research assumes saturation (i.e., nodes always have packets available for transmission), see, e.g. [2, 23, 8]. The resulting models are simple yet elegant, and provide throughput estimates that match remarkably well with experimental results for saturated IEEE 802.11-based systems [16].

Emerging large-scale wireless networks violate this saturation assumption in two important ways: (i) traffic is intermittently generated over time, so nodes may be empty for a non-negligible amount of time (unsaturation); and (ii) packets may be transmitted over large distances through intermediate nodes (multi-hop). These two properties significantly complicate the dynamics of the network, since the packet arrival process produces queueing dynamics and the buffers may empty from time to time. Nodes without packets temporarily refrain from the medium contention and thus the activity process and the queueing dynamics of the buffer contents are strongly intertwined.

As a result, unsaturated CSMA models are drastically more complicated, and in general are untractable. Even the simplest instance of such model, where each node has its own arrival process, and packets are not forwarded (single-hop), only a few partial results are known, all in the case of a full interference graph (all nodes interfere with each other) [4, 20]. In [20] it is demonstrated that in an unsaturated network, even if the arrival rates at the individual nodes are less than their saturation throughput, the network may still be unstable. This counterintuitive result demonstrates that, due to the complex interactions between the nodes' buffer contents, the saturation throughput cannot be reliably used to explain the performance of unsaturated networks. In general, stability conditions do not admit a close-form expression [4].

IFIP WG 7.3 Performance 2017. Nov. 14-16, 2017, New York, NY USA

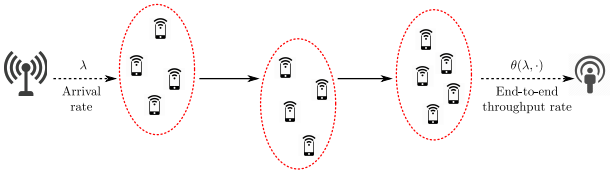

Figure 1: Example of a multi-hop network with many nodes, where packets are forwarded through the network.

Multi-hop networks are even more complex than single-hop networks, so only a few results are available. In [1] the authors studied a multi-hop CSMA network using ns-2 simulations and experiments, and showed that the end-to-end throughput may decrease as the exogenous arrival increases, due to congestion. This phenomenon is unique to multi-hop networks, and cannot be captured by saturated systems or single-hop models. In [15] stability and throughput results are obtained for multi-hop networks, under the assumption that a node freezes its arrival process during a back-off period, and in the regime that all nodes are unsaturated. The behavior of a three-node linear multi-hop network is computed in [18], but the approach used there cannot be extended to larger, more realistic networks.

In order to circumvent these complex and untractable models, we observe that emerging large-scale wireless networks are often dense, and consist of many nodes. This suggests a mean-field scaling regime, where nodes are grouped in classes, and the number of nodes grows large, see Figure 1. In this regime, any tagged node is subject to an averaged effect (the mean-field) obtained aggregating the influence of all other nodes. Thus, the high-dimensional stochastic process describing the queueing dynamics exhibits a state-space aggregation and is shown to converge to the solution of a deterministic initial-value problem (IVP). The equilibrium point of the IVP is often used to obtain accurate approximations for the performance metrics of the system with a finite number of nodes.

Such mean-field approach has recently been successfully applied to single-hop networks, see [3, 5, 6, 7]. In [7] the authors discuss sufficient conditions for asymptotic accuracy of such approximations. Via this approach, in [3] an approximation for the stability region of the system is obtained and in [5, 6] the stationary queue length and packet delay performance are discussed. In the context of random-access networks the mean-field regime not only provides analytical tractability, but is also highly relevant due to the envisioned massive numbers of IoT devices.

In the present paper we consider for the first time the mean-field limit for multi-hop CSMA networks. We show that the stochastic process converges to an initial-value problem (IVP) as the network is scaled in an appropriate way, and characterize the equilibrium point of that IVP. We show numerically that the equilibrium point provides an excellent approximation for the stationary behavior of the pre-limit stochastic network, even for small numbers of nodes. We then focus on linear networks, where we can solve the equilibrium points in closed form for certain back-off rates (i.e., the reciprocal of the mean duration of a back-off period). We show that the choice of back-off rates has significant impact on the end-to-end throughput, and determine which set of back-off rates gives the best performance. That is, they maximize the end-to-end throughput as well as providing robustness, meaning that the throughput does not decrease when the network is in overload. We then use these insights to design optimal back-off rates for general networks, and show numerically that these perform well.

It is worth remarking upon that our main result is not the proof of the mean-field convergence of the CSMA model to the IVP, which is an extension of an earlier result in [6]. Rather, our contribution lies in the characterization and utilization of the equilibrium points of the new IVP, which is particularly challenging due to (i) the increased interaction between nodes compared to single-hop networks; and (ii) the presence of local bottlenecks. Summarizing, our main contributions are:

- We characterize the equilibrium point of the mean-field limit, and show that this is an excellent approximation for the pre-limit network. This is the first accurate and closed-form approximation for such networks;
- We find the vector of back-off rates which is optimal in the mean-field regime, and show that this works very well for the pre-limit network as well.

The remainder of this paper is structured as follows. In Section 2 we introduce the stochastic model for multi-hop CSMA networks, and give some background on saturated CSMA models required for our analysis. In Section 3 we discuss the mean-field regime, and characterize the equilibrium points of the resulting IVP. These equilibrium points are further studied in Sections 4 and 5 for linear networks. In Section 4 we solve the equilibrium point in closed-form under various choices of back-off rates, and in Section 5 we determine the optimal back-off rate vectors. These results are then used in Section 6 to design good back-off rates for general networks. We conclude in Section 7, and outline some open problems and future research directions.

2. MODEL

We consider a CSMA network of N nodes grouped into C classes, and denote $\mathcal{C} = \{1, \ldots, C\}$. Class c contains N_c nodes, such that $N = \sum_c N_c$. The network structure is characterized by an undirected interference graph $(\mathcal{C}, \mathcal{E})$, with the vertices in the graph corresponding to the classes, and two classes share an edge if they interfere. Two nodes within the same class or neighboring classes are prevented from simultaneous activity.

The transmission of a packet in a class-c node requires an exponentially distributed random time with parameter 1. In between two consecutive transmissions, each node remains silent for a random back-off period. The duration of this period is exponentially distributed with parameter $\nu_c^{(N)} = \nu_c/N_c$, for some $\nu_c > 0$. We refer to $\nu_c^{(N)}$ as the back-off rate of a class-c node, and denote $\boldsymbol{\nu} = (\nu_1, \ldots, \nu_C)$. At the end of its back-off period, a class-c node senses the medium and either starts a transmission if it senses no activiy from neighboring nodes in the graph, or it starts new back-off period otherwise. Because at most one node can be active per class, we can represent the activity state of the network by

$$\boldsymbol{Y}^{(N)}(t) \in \Omega, \qquad \Omega = \{\boldsymbol{\omega} \in \{0,1\}^C : \omega_c \omega_{c'} = 0 \forall \{c, c'\} \in \mathcal{E}\},$$

where $Y_c^{(N)}(t) = 1$ if a node in class c is transmitting at time t, and $Y_c^{(N)}(t) = 0$ otherwise.

Each node is equipped with a buffer of infinite size, with $Q_{c,k}^{(N)}(t)$ denoting the queue length of the k-th node in class c at time t, not including the packet potentially in transmission. Here we denote $\boldsymbol{Q}^{(N)}(t) = (Q_{c,k}^{(N)}(t))_{c\in\mathcal{C};k=1,\ldots,N_c}$. The queue length of a node is increased by one whenever a new packet arrives, and is decreased by one whenever this node starts a new transmission. A transmission is always assumed to be successful. In multi-hop networks, packets arrive and need to be transmitted by a sequence of nodes before leaving the system. Packets arrive at class 1 according to an exogenous Poisson arrival process with rate λ, and arriving packets are assigned uniformly at random to one of the N_1 nodes. When a packet transmission is completed at a class-c node with $c < C$, it is forwarded to a node of class $c+1$ selected uniformly at random. When a class-C node completes a transmission, the packet leaves the system. As we shall see, choosing to route packets uniformly at random will result in a tractable limit that provides an accurate approximation for pre-limit models, even for few of nodes.

For every class c, we define the following subsets of Ω:

$$\Omega_{-c} = \{\boldsymbol{\omega} \in \Omega : \omega_c = 0, \omega_d = 0 \,\forall\, d \text{ s.t. } \{c,d\} \in \mathcal{E}\},$$
$$\Omega_{+c} = \{\boldsymbol{\omega} \in \Omega : \omega_c = 1\}.$$

This means that $\boldsymbol{\omega} \in \Omega_{-c}$ if and only if in the class activity state $\boldsymbol{\omega}$ none of the nodes belonging to class c or to a class interfering with class c are active, while $\boldsymbol{\omega} \in \Omega_{+c}$ if and only if a class-c node is active.

The process $(\boldsymbol{Q}^{(N)}(t), \boldsymbol{Y}^{(N)}(t))_{t\geq 0}$ is a Markov process, with the following transitions:

• A new packet arrives in the system and joins the buffer of the k-th node in class 1. This happens at rate λ/N_1. It generates the transition

$$Q_{1,k}^{(N)} \;\to\; Q_{1,k}^{(N)} + 1.$$

• The k-th node in class c completes a back-off period and begins a transmission. This happens at rate ν_c/N_c and only if $Q_{c,k}^{(N)} > 0$ and $\boldsymbol{Y}^{(N)} \in \Omega_{-c}$. It generates the transition

$$Q_{c,k}^{(N)} \;\to\; Q_{c,k}^{(N)} - 1, \qquad Y_c^{(N)} \;\to\; 1.$$

• A node in class $c < C$ completes a transmission and the packet joins the buffer of the k'-th node in class $c+1$. This happens at rate $1/N_{c+1}$ and only if $\boldsymbol{Y}^{(N)} \in \Omega_{+c}$. It generates the transition

$$Q_{c+1,k'}^{(N)} \;\to\; Q_{c+1,k'}^{(N)} + 1, \qquad Y_c^{(N)} \;\to\; 0.$$

• A node in class C completes a transmission and a packet leaves the system. This happens at rate 1 and only if $\boldsymbol{Y}^{(N)} \in \Omega_{+C}$. It generates the transition

$$Y_C^{(N)} \;\to\; 0.$$

Unfortunately, the complex interactions between the nodes make the process $(\boldsymbol{Q}^{(N)}(t), \boldsymbol{Y}^{(N)}(t))_{t\geq 0}$ intractable. Several partial attempts at understanding this process have been made in the past few years, but these are limited to small networks under specific assumptions [1, 18]. Instead we use mean-field limits to get a better understanding of the network. As we shall see this limit provides a remarkably good approximation of the network.

2.1 Preliminary results for saturated networks

Before we turn towards the mean-field limit, we first present some background on saturated networks, i.e., networks where all nodes have an infinite supply of packets available for transmission. These results will turn out to play a key role in the analysis of the mean-field limit. We denote the aggregate back-off rate of the nodes in class c by α_c and $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_C)$. Because nodes have an infinite supply of packets available they always compete for the medium, and the network is fully described by $\boldsymbol{Y}^{(N)}(t)$, which keeps track of whether or not each class has an active node.

It turns out that in this case $\{Y^{(N)}(t)\}$ in isolation is a Markov process, and has a product-form stationary distribution [2, 20, 23]:

$$\pi(\boldsymbol{\omega};\boldsymbol{\alpha}) = Z(\boldsymbol{\alpha})^{-1} \prod_{c=1}^{C} \alpha_c^{\omega_c}, \qquad \boldsymbol{\omega} \in \Omega, \tag{1}$$

where $Z(\boldsymbol{\alpha}) = \sum_{\boldsymbol{\omega}\in\Omega} \prod_{c=1}^{C} \alpha_c^{\omega_c}$ denotes the normalization constant.

We define

$$\pi(\Omega_{+c};\boldsymbol{\alpha}) = \sum_{\boldsymbol{\omega}\in\Omega_{+c}} \pi(\boldsymbol{\omega};\boldsymbol{\alpha}), \quad \pi(\Omega_{-c};\boldsymbol{\alpha}) = \sum_{\boldsymbol{\omega}\in\Omega_{-c}} \pi(\boldsymbol{\omega};\boldsymbol{\alpha}).$$

Similar to the unsaturated network discussed at the beginning of Section 2, we are interested in the throughput of the network $\bar{\boldsymbol{\theta}}(\boldsymbol{\alpha}) = (\bar{\theta}_1(\boldsymbol{\alpha}), \ldots, \bar{\theta}_C(\boldsymbol{\alpha}))$, defined as $\bar{\theta}_c(\boldsymbol{\alpha}) = \pi(\Omega_{+c};\boldsymbol{\alpha})$.

In [14] it has been shown that the range of the throughput mapping $\bar{\boldsymbol{\theta}}$ is equal to the convex hull Γ of all feasible activity vectors. Moreover, in [21] this mapping was shown to be globally invertible, so for every $\boldsymbol{\gamma} \in \text{int}(\Gamma)$, there exists a unique vector $\boldsymbol{\sigma}(\boldsymbol{\gamma}) \in \mathbb{R}_+^C$ such that

$$\bar{\boldsymbol{\theta}}(\boldsymbol{\sigma}(\boldsymbol{\gamma})) = \boldsymbol{\gamma}. \tag{2}$$

3. MEAN-FIELD ANALYSIS

Because the Markov process $(\boldsymbol{Q}^{(N)}(t), \boldsymbol{Y}^{(N)}(t))_{t\geq 0}$ defined in Section 2 is not tractable, we focus on the mean-field analysis of the model. That is, we let $N \to \infty$ and study the limiting process. In the remainder of this paper we shall demonstrate that this limit proves an excellent approximation for the pre-limit network, even for small N.

Before proceeding to this, we first introduce some additional notation. The queue length process $\boldsymbol{Q}^{(N)}(t)$ uniquely defines the population process $\boldsymbol{X}^{(N)}(t) = (X_{c,n}(t))_{c\in\mathcal{C};n\in\mathbb{N}_0}$, where

$$X_{c,n}^{(N)}(t) = \frac{1}{N_c} \sum_{k=1}^{N_c} \mathbb{1}\{Q_{c,k}^{(N)}(t) = n\}$$

denotes the fraction of class-c nodes with n packets in the buffer at time t. Observe that $(\boldsymbol{X}^{(N)}(t), \boldsymbol{Y}^{(N)}(t))_{t\geq 0}$ is still a Markov process with state space within $E^1 \times \Omega$, where

$$E^1 = \{\boldsymbol{x} = (x_{c,n})_{c\in\mathcal{C},n\in\mathbb{N}_0} : \sum_{n=0}^{\infty} x_{c,n} = 1,\ \forall c \in \mathcal{C}\}. \tag{3}$$

Given $\boldsymbol{x} \in E^1$, denote $\boldsymbol{x}_0 = \{x_{c,0}\}_{c\in\mathcal{C}} \in [0,1]^C$. Denote also by $\boldsymbol{e}$ the vector with every component equal to 1, and let $\boldsymbol{e}_{j,l}$ be the matrix with all 0 components except for a 1 at coordinate (j,l). We denote by $\cdot$ the component-wise multiplication, so that for instance $(\boldsymbol{e} - \boldsymbol{x}_0) \cdot \boldsymbol{\nu} = ((1 - x_{1,0})\nu_1, \ldots, (1 - x_{C,0})\nu_C)$

We are interested in the mean-field analysis of the model, i.e., we consider the regime where N grows to infinity and denote $p_c = \lim_{N\to\infty} N_c/N > 0$. We then obtain the following result.

THEOREM 1. *Assume* $\boldsymbol{X}^{(N)}(0) \xrightarrow{N\to\infty} \boldsymbol{x}^\infty \in E^1$. *Then the sequence of processes*

$$(\boldsymbol{X}^{(N)}(Nt)) \in D_{E^1}[0,\infty)$$

weakly converges to $\boldsymbol{x}(t) \in C_{E^1}[0,\infty)$ *which is determined by the unique solution of the initial-value problem*

$$\frac{d\boldsymbol{x}_c(t)}{dt} = H_c(\boldsymbol{x}(t)), \qquad \boldsymbol{x}(0) = \boldsymbol{x}^\infty, \tag{4}$$

where the function $\boldsymbol{H}(\cdot) = (H_1(\cdot), \dots, H_C(\cdot))$ *is defined by*

$$\begin{aligned} H_1(\boldsymbol{x}) = \frac{1}{p_1}\Big(&\lambda \sum_{n=0}^{\infty} x_{1,n}(\boldsymbol{e}_{1,n+1} - \boldsymbol{e}_{1,n}) \\ &+ \nu_1 \pi(\Omega_{-1}; (\boldsymbol{e} - \boldsymbol{x}_0)\cdot\boldsymbol{\nu}) \sum_{n=1}^{\infty} x_{1,n}(\boldsymbol{e}_{1,n-1} - \boldsymbol{e}_{1,n})\Big), \end{aligned} \tag{5}$$

$$\begin{aligned} H_c(\boldsymbol{x}) = \frac{1}{p_c}\Big(&\pi(\Omega_{+c-1}; (\boldsymbol{e} - \boldsymbol{x}_0)\cdot\boldsymbol{\nu}) \sum_{n=0}^{\infty} x_{c,n}(\boldsymbol{e}_{c,n+1} - \boldsymbol{e}_{c,n}) \\ &+ \nu_c \pi(\Omega_{-c}; (\boldsymbol{e} - \boldsymbol{x}_0)\cdot\boldsymbol{\nu}) \sum_{n=1}^{\infty} x_{c,n}(\boldsymbol{e}_{c,n-1} - \boldsymbol{e}_{c,n})\Big) \end{aligned} \tag{6}$$

for any $c > 1$.

PROOF. The proof of this theorem is analogous to that of [6, Thm. 1] for single-hop networks, with the class instantaneous arrival rate $(\lambda_1, \dots, \lambda_C)$ replaced by

$$\Big(\lambda, \mathbb{1}\{\boldsymbol{Y}^{(N)}(t) \in \Omega_{+1}\}, \dots, \mathbb{1}\{\boldsymbol{Y}^{(N)}(t) \in \Omega_{+(C-1)}\}\Big)$$

at time t. □

Note that $H_c(\boldsymbol{x})$ in (5) and (6) consist of two parts. The first summand describes the changes in the population process as a new packet arrives, while the second describes the changes as a back-off period is completed, i.e., a packet leaves the buffer. The impact of the activity process is captured by the stationary distribution of the saturated network with per class back-off rates $(\boldsymbol{e} - \boldsymbol{x}_0)\cdot\boldsymbol{\nu}$. In particular, in the mean-field regime the activity process reaches stationarity before the population process changes. For this reason, $\pi(\Omega_{-c}; (\boldsymbol{e} - \boldsymbol{x}_0)\cdot\boldsymbol{\nu})$ denotes the stationary measure of the activity states allowing class-c nodes to back-off, and similarly $\pi(\Omega_{+c}; (\boldsymbol{e} - \boldsymbol{x}_0)\cdot\boldsymbol{\nu})$ denotes the stationary measure of the activity states allowing class-c nodes to transmit. It is worth remarking that the interaction between different classes only depends on $\boldsymbol{x}$ through $\boldsymbol{x}_0$, which will turn out to be a crucial property in our analysis.

3.1 Equilibrium points

We are interested in the analysis of the model in stationarity, hence we now focus on the equilibrium points $\boldsymbol{x}^*$ of the differential equation (4). We stress that convergence to the equilibrium point is not addressed in this work. This is a challenging problem which has only been solved in single-hop networks for the case where $(\mathcal{C}, \mathcal{E})$ is the complete graph (so all classes interfere with each other) [6].

Before we discuss our characterization of the equilibrium points of (4), we first introduce the following system of load equations for our network. Denote $\boldsymbol{\rho} = (\rho_1, \dots, \rho_C)$, where the $\rho_c \in \mathbb{R}_+$ are a solution to

$$\rho_1 = \frac{\lambda}{\nu_1 \pi(\Omega_{-1}; \boldsymbol{\rho}\cdot\boldsymbol{\nu})}, \tag{7}$$

$$\rho_c = \frac{\pi(\Omega_{+(c-1)}; \boldsymbol{\rho}\cdot\boldsymbol{\nu})}{\nu_c \pi(\Omega_{-c}; \boldsymbol{\rho}\cdot\boldsymbol{\nu})}, \quad c = 2, \dots, C, \tag{8}$$

with $\boldsymbol{\rho} < \boldsymbol{e}$ componentwise. It turns out that the ρ_c can be interpreted as the load of each class in the equilibrium point: the right-hand side of both (7) and (8) consists of the quotient of the rate that traffic arrives into that class and the service rate of that class, assuming that each class operates as a single node in a saturated network with back-off rates $\boldsymbol{\rho}\cdot\boldsymbol{\nu} \le \boldsymbol{\nu}$. This intuition leads us to the following characterization of the equilibrium points.

PROPOSITION 1. $\boldsymbol{x}^* \in E^1$ *is an equilibrium point of* (4) *if and only if there exists a solution* $\boldsymbol{\rho}$ *to* (7) *and* (8) *with* $\boldsymbol{\rho} < \boldsymbol{e}$, *such that*

$$x^*_{c,n} = (1-\rho_c)\rho_c^n, \qquad c \in \mathcal{C},\ n \in \mathbb{N}_0. \tag{9}$$

PROOF. Let $\boldsymbol{x}^* \in E^1$ be an equilibrium point of (4), so it satisfies $H_c(\boldsymbol{x}^*) = 0, c = 1, \dots, C$. After some manipulations, this system of equations can be written as

$$-\lambda x^*_{1,n} + \nu_1 \pi(\Omega_{-1}; (\boldsymbol{e} - \boldsymbol{x}^*_0)\cdot\boldsymbol{\nu}) x^*_{1,n+1} = 0, \tag{10}$$

$$\begin{aligned} &-\pi(\Omega_{+(c-1)}; (\boldsymbol{e} - \boldsymbol{x}^*_0)\cdot\boldsymbol{\nu}) x^*_{c,n} \\ &\quad + \nu_c \pi(\Omega_{-c}; (\boldsymbol{e} - \boldsymbol{x}^*_0)\cdot\boldsymbol{\nu}) x^*_{c,n+1} = 0, \quad \forall\ c > 1, \end{aligned} \tag{11}$$

for $n \in \mathbb{N}_0$. Using successive substitution we can solve (10) as

$$x^*_{1,n} = x^*_{1,0}\rho_1^n, \qquad \rho_1 = \frac{\lambda}{\nu_1 \pi(\Omega_{-1}; (\boldsymbol{e} - \boldsymbol{x}^*_0)\cdot\boldsymbol{\nu})}, \tag{12}$$

for every $n \in \mathbb{N}_0$. Because $\boldsymbol{x}^* \in E^1$, we have that $\sum_n x^*_{1,n} = 1$. Plugging in (12), this can be written as

$$1 = \sum_{n=0}^{\infty} x^*_{1,0}\rho_1^n = \frac{x^*_{1,0}}{1-\rho_1}, \qquad \text{iff } \rho_1 < 1,$$

which yields $x^*_{1,0} = 1 - \rho_1$. Substituting this into (12) we obtain

$$x^*_{1,n} = (1-\rho_1)\rho_1^n, \qquad \forall\ n \in \mathbb{N}_0.$$

We can repeat the above steps for $c \ge 2$ using (11) to obtain

$$x^*_{c,n} = (1-\rho_c)\rho_c^n, \quad \rho_c = \frac{\pi(\Omega_{+(c-1)}; (\boldsymbol{e} - \boldsymbol{x}^*_0)\cdot\boldsymbol{\nu})}{\nu_c \pi(\Omega_{-c}; (\boldsymbol{e} - \boldsymbol{x}^*_0)\cdot\boldsymbol{\nu})}, \tag{13}$$

for every $c > 1$ and $n \in \mathbb{N}_0$, and thus it holds that

$$x^*_{c,0} = 1 - \rho_c, \qquad \forall\ c \in \mathcal{C}. \tag{14}$$

By (14) we see that $\boldsymbol{e} - \boldsymbol{x}^*_0 = \boldsymbol{\rho}$, so the ρ_c defined in (12) and (13) are indeed equivalent to (7) and (8), concluding one direction of the proof. The other direction can be verified by substituting the proposed solution (9) into (10) and (11). □

So we know that each solution of (7) and (8) generates one equilibrium point, and that for each equilibrium point the population process is geometric. Note that the equilibrium point (1) does not contain p_c, so it is independent of the pre-limit class sizes N_c.

One of the performance metrics of interest is the throughput of each class, defined as the rate at which it is transmitting packets. Since packets have unit length, this is equivalent to the fraction of time each class is transmitting. We know from Proposition 1 that in the mean-field limit the network behaves as a saturated CSMA networks with back-off rates $\boldsymbol{\rho}\cdot\boldsymbol{\nu}$. Thus, with $\bar{\theta}_c$ the throughput of a saturated network defined in Section 2.1, the multi-hop throughput θ_c can be written as

$$\theta_c(\boldsymbol{\rho};\lambda,\boldsymbol{\nu}) = \bar{\theta}_c(\boldsymbol{\rho}\cdot\boldsymbol{\nu}). \tag{15}$$

We are particularly interested in the stationary end-to-end throughput θ of the network, i.e., the throughput of class C:

$$\theta(\boldsymbol{\rho};\lambda,\boldsymbol{\nu}) = \theta_C(\boldsymbol{\rho};\lambda,\boldsymbol{\nu}).$$

We denote $\boldsymbol{\theta}(\boldsymbol{\rho};\lambda,\boldsymbol{\nu}) = (\theta_1(\boldsymbol{\rho};\lambda,\boldsymbol{\nu}),\dots,\theta_C(\boldsymbol{\rho};\lambda,\boldsymbol{\nu}))$, and when the solution $\boldsymbol{\rho}$ is unique, we remove $\boldsymbol{\rho}$ from the throughput notation.

Using the characterization presented in Proposition 1 we can gain a better understanding of the equilibrium points. Recall that Γ denotes the capacity region of the corresponding saturated network, defined in Section 2.1.

PROPOSITION 2. *Given network structure $(\mathcal{C},\mathcal{E})$, arrival rate λ, and back-off rates $\boldsymbol{\nu}$ we have that:*

(i) If $\lambda\boldsymbol{e}\notin\mathrm{Int}(\Gamma)$, then (4) has no equilibrium points

(ii) If $\lambda\boldsymbol{e}\in\mathrm{Int}(\Gamma)$, there exists at most one equilibrium point of (4).

(iii) This equilibrium point is such that all nodes have the same throughput $\theta_c(\lambda,\boldsymbol{\nu})=\lambda$ for every $c\in\mathcal{C}$.

PROOF. Since the equilibrium points of (4) are characterized by solutions to (7) and (8) with $\boldsymbol{\rho}<\boldsymbol{e}$, we are interested in such solutions. Using that $\pi(\Omega_{+c};\boldsymbol{\rho}\cdot\boldsymbol{\nu}) = \nu_c\rho_c\pi(\Omega_{-c};\boldsymbol{\rho}\cdot\boldsymbol{\nu})$, we can rearrange (7) and (8) to obtain

$$\pi(\Omega_{+c};\boldsymbol{\rho}\cdot\boldsymbol{\nu}) = \lambda, \qquad \forall\, c\in\mathcal{C}. \tag{16}$$

So if an equilibrium point exists it satisfies (16), proving item (iii).

Recall from Section (2.1) that $\pi(\Omega_{+c};\boldsymbol{\rho}\cdot\boldsymbol{\nu}) = \bar{\theta}_c(\boldsymbol{\rho}\cdot\boldsymbol{\nu})$, the throughput of class c in the fully saturated network with back-off vector $\boldsymbol{\rho}\cdot\boldsymbol{\nu}$. Thus we know from [21] that there exist a unique solution $\boldsymbol{\rho}\in(0,\infty)^C$ such that $\bar{\boldsymbol{\theta}}(\boldsymbol{\rho}\cdot\boldsymbol{\nu})=\lambda\boldsymbol{e}$ for any $\lambda\boldsymbol{e}\in\Gamma$. This proves item (ii). Finally, item (i) is proven by the observation that by [21], there is no solution to $\bar{\boldsymbol{\theta}}(\boldsymbol{\rho}\cdot\boldsymbol{\nu})=\lambda\boldsymbol{e}$ outside of Γ. □

3.2 Partial equilibrium points

Propositions 1 and 2 characterize the limiting behavior of (4) in the case that $\boldsymbol{\rho}<\boldsymbol{e}$, i.e., all classes are stable and all nodes have the same throughput. However, from a practical perspective we are particularly interested in the case where $\rho_c\geq 1$ for some c, so one or more classes are overloaded. These kinds of local bottlenecks are typical for multi-hop networks, and are the cause of many performance issues. As we shall see, by taking into account the possibility of local bottlenecks, the analysis of the equilibrium point becomes much more complex and is no longer equivalent to single-hop networks.

When $\rho_c>1$, the drift of the queue length at class-c nodes is positive, and as a result $x_{c,n}(t)\to 0$ $(t\to\infty)$ for all $n\in\mathbb{N}_0$. In particular, we have that $x_{c,0}\to 0$ for these unstable classes, so all nodes in this class have packets available for transmission.

To accommodate analysis of these partially stable systems, we introduce the concept of partial equilibrium points, where certain classes may be in equilibrium in the traditional sense, while other classes are unstable, so all their nodes have infinite backlog. We define $\boldsymbol{\rho}^- = (\rho_1^-,\dots,\rho_C^-)$, where $\rho_c^- = \min\{1,\rho_c\}$, and modify the load equations (7) and (8) as

$$\rho_1 = \frac{\lambda}{\nu_1\pi(\Omega_{-1};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu})}, \tag{17}$$

$$\rho_c = \frac{\pi(\Omega_{+(c-1)};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu})}{\nu_c\pi(\Omega_{-c};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu})}, \quad c=2,\dots,C. \tag{18}$$

We then define partial equilibrium points as follows.

DEFINITION 1. *We say $\boldsymbol{x}^*$ is a partial equilibrium point of (4) if there exists a $\boldsymbol{\rho}$-solution to (17) and (18) such that*

$$x^*_{c,n} = (1-\rho_c)\rho_c^n, \quad \text{if } \rho_c<1,$$
$$x^*_{c,n} = 0, \quad \text{if } \rho_c\geq 1.$$

From Definition 1 we see that partial equilibrium points are fully characterized by the $\boldsymbol{\rho}$ solutions to (17) and (18). We say that this point has a load configuration $\boldsymbol{s}=(s_1,\dots,s_C)$, where $s_c=\mathtt{U}$ if $\rho_c\leq 1$, and $s_c=\mathtt{S}$ otherwise.

To accommodate the search for partial equilibrium points, let us rewrite these equations slightly. Define $\rho_c^+ = \min\{1,\rho_c^{-1}\}$, and observe from the stationary distribution of the saturated network (1) that

$$\pi(\Omega_{-c};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu}) = (\nu_c\rho_c^-)^{-1}\pi(\Omega_{+c};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu}).$$

Substituting this into (17) and (18) and rearranging, we obtain

$$\pi(\Omega_{+1};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu}) = \lambda\rho_1^-/\rho_1,$$
$$\pi(\Omega_{+c};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu}) = \pi(\Omega_{+(c-1)};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu})\rho_c^-/\rho_c, \quad c=2,\dots,C,$$

It is readily seen that $\rho_c^-/\rho_c = \rho_c^+$, and successive substitution yields

$$\pi(\Omega_{+c};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu}) = \pi(\Omega_{+(c-1)};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu})\rho_c^+$$
$$= \pi(\Omega_{+(c-2)};\boldsymbol{\rho}^-\cdot\boldsymbol{\nu})\rho_c^+\rho_{c-1}^+ = \cdots = \lambda\prod_{d=1}^{c}\rho_d^+, \tag{19}$$

which is equivalent to (17) and (18).

The right-hand side of (19) is equal to the total arrival rate times the fraction of load which is transmitted by classes $1,\dots,c$, which corresponds to the throughput of class c itself. On the other hand, the left-hand side can be interpreted as the fraction of time that class-c is transmitting. So (19) states that the flow out of class c has to be equal to the fraction of time the class is active.

Note that the equilibrium points from Proposition 1 are a special case of the partial equilibrium points introduced in Definition 1, with all $\rho_c<1$. The throughput for the partial equilibrium point can be obtained similarly to (15), with $\boldsymbol{\rho}$ replaced by $\boldsymbol{\rho}^-$. In the remainder of this paper we are only interested in partial equilibrium points rather than the 'pure' equilibrium points discussed in Proposition 1. Going forward, we refer to them simply as equilibrium points.

Finding equilibrium points by solving (19) is not straightforward, it is in fact as difficult as solving a nonlinear complementarity problem (NLCP), for which even conditions for

existence and uniqueness of solutions are not yet known in general [11]. However, we can show that (19) always has a solution, thus ensuring the existence of a equilibrium point of (4) for any network.

PROPOSITION 3. *For any network structure* $(\mathcal{C}, \mathcal{E})$*, arrival rate* λ*, and back-off rates* $\boldsymbol{\nu}$*, there exists a solution to* (19)*.*

The proof of this proposition can be found in Section A.1.

In addition to the throughput θ introduced in Section 3.1, our second performance metric of interest is how much traffic the network can sustain, measured as the maximum arrival rate such that the network is still stable:

$$\lambda^*(\boldsymbol{\nu}) = \sup_{\lambda>0} \Big\{ \exists \text{ solution } \boldsymbol{\rho} \text{ to (19) s.t. } \boldsymbol{\rho} \leq \boldsymbol{e} \Big\}.$$

We can characterize λ^* as follows.

PROPOSITION 4. *It holds that*

$$\lambda^*(\boldsymbol{\nu}) = \sup_{\lambda>0} \Big\{ \boldsymbol{\sigma}(\lambda \boldsymbol{e}) \leq \boldsymbol{\nu} \Big\}.$$

PROOF. This result immediately follows from the definition of $\boldsymbol{\sigma}(\cdot)$ in (2), in fact

$$\begin{aligned}\lambda^*(\boldsymbol{\nu}) &= \sup_{\lambda>0} \Big\{ \exists\, \boldsymbol{\rho} \leq \boldsymbol{e} \text{ s.t. } \lambda \boldsymbol{e} = \bar{\boldsymbol{\theta}}(\boldsymbol{\rho} \cdot \boldsymbol{\nu}) \Big\} \\ &= \sup_{\lambda>0} \Big\{ \exists\, \boldsymbol{\rho} \leq \boldsymbol{e} \text{ s.t. } \boldsymbol{\sigma}(\lambda \boldsymbol{e}) = \boldsymbol{\rho} \cdot \boldsymbol{\nu} \Big\} = \sup_{\lambda>0} \Big\{ \boldsymbol{\sigma}(\lambda \boldsymbol{e}) \leq \boldsymbol{\nu} \Big\}.\end{aligned}$$

□

It turns out that the equilibrium point from Definition 1 provides an excellent approximation for the pre-limit (stochastic) CSMA network in stationarity. To illustrate this, consider a linear network with three classes where class 2 interferes with both classes 1 and 3, with $\nu_c = 6$, $c = 1, 2, 3$. In Figure 2(a) we plot the end-to-end throughput against the arrival rate, for both the mean-field limit and for the stochastic network (obtained through simulation of the model in Section 2) for $N_c = 1, 10, 100$. Note that the mean-field solution reflects the decrease in end-to-end throughput when the network is in overload, and that the point of overload λ^* is almost identical for the mean-field and the stochastic cases. In Figure 2(b) we take $\lambda = 0.5$, so that $\boldsymbol{\rho} = (0.6009,\ 1.3838,\ 0.2171)$. We plot the fraction of nodes with 0 and 4 packets for both classes 1 and 3, for both the mean-field system (dashed) and the stochastic system (marked). The x-axis shows N_c for the stochastic system, and we see that the geometric queue-length distribution seen in the mean-field limit accurately captures the behavior of the stochastic model, even for moderate N_c.

4. PERFORMANCE ANALYSIS OF LINEAR NETWORKS

In this section we use the equilibrium points of the mean-field limit (4) to gain a better understanding of the network. We restrict ourselves to linear networks with nearest-neighbor blocking where class-c nodes prevent nodes in classes $c \pm 1$ to activate, i.e., $\mathcal{C} = \{1, \ldots, C\}$ and $\mathcal{E} = \{\{c, c+1\}\}_{c=1,\ldots,c-1}$. Special attention will be given to the influence of the back-off rates on the performance of the system. For convenience, we introduce $R(\lambda, \boldsymbol{\nu})$ as the set of solutions to (19) for a given λ and $\boldsymbol{\nu}$. We consider various cases where

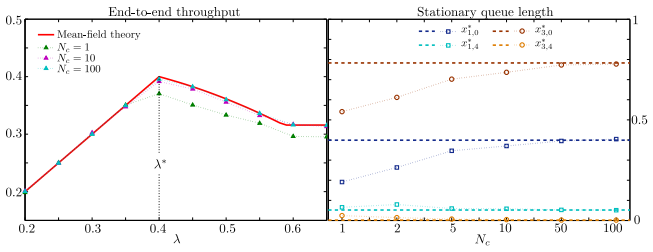

Figure 2: Comparing the equilibrium point of the mean-field limit with stochastic CSMA networks.

we show that the equilibrium point of (4) is unique, and can be solved explicitly.

In Section 4.1 we consider a linear network with general back-off rates, and derive $\lambda^*(\boldsymbol{\nu})$ as well as the maximum end-to-end throughput achievable. In the subsequent sections we provide a more detailed analysis for linear networks in the following specific cases. In Sections 4.2 and 4.3 we consider a network with C classes with back-off rates $\boldsymbol{\nu}^u = (\nu, \ldots, \nu)$ and $\boldsymbol{\nu}^f = (\nu, \nu(1+\nu), \ldots, \nu(1+\nu), \nu)$ respectively, and in Section 4.4 a 3-class network with general back-off rates, in oversaturation ($\lambda \to \infty$).

4.1 Maximal throughput in linear networks

We start by rewriting (19) for linear networks with general back-off rates. These networks are acyclic and we may exploit [12, Thm. 1], which concerned saturated CSMA networks. Specifically, $\lambda \prod_{d=1}^{c} \rho_d^+$ can be interpreted as the target throughput of the c-th class of nodes, whose back-off rate is $\nu_c \rho_c^-$. Thanks to this interpretation, $\boldsymbol{\rho}$ solves (19) if and only if it solves

$$\nu_1 \rho_1^- = \frac{\lambda \rho_1^+}{1 - \lambda \rho_1^+ (1 + \rho_2^+)}, \tag{20}$$

$$\nu_c \rho_c^- = \frac{\lambda \prod_1^c \rho_d^+ (1 - \lambda \prod_1^c \rho_d^+)}{\big(1 - \lambda \prod_1^{c-1} \rho_d^+ (1 + \rho_c^+)\big)\big(1 - \lambda \prod_1^c \rho_d^+ (1 + \rho_{c+1}^+)\big)}, \quad c = 2, \ldots, C-1, \tag{21}$$

$$\nu_C \rho_C^- = \frac{\lambda \prod_1^C \rho_d^+}{1 - \lambda \prod_1^{C-1} \rho_d^+ (1 + \rho_C^+)}. \tag{22}$$

Using this refined characterization we can study the $\lambda^*(\boldsymbol{\nu})$ for linear networks with general back-off rates. In the argument we will use the following lemma, which follows immediately from [12, Thm. 1] with target throughput $\lambda \boldsymbol{e}$. Recall that $\boldsymbol{\sigma}$ is the throughput inverse for the saturated CSMA model, introduced in Section 2.1.

LEMMA 1. *For any* $\lambda < \frac{1}{2}$*, it holds that*

$$\boldsymbol{\sigma}(\lambda \boldsymbol{e}) = \Big(\frac{\lambda}{1-2\lambda}, \frac{\lambda(1-\lambda)}{(1-2\lambda)^2}, \ldots, \frac{\lambda(1-\lambda)}{(1-2\lambda)^2}, \frac{\lambda}{1-2\lambda} \Big). \tag{23}$$

Observe that, due to Equations (20)-(22), for any solution $\boldsymbol{\rho} \in R(\lambda, \boldsymbol{\nu})$ with $\boldsymbol{\rho} \leq \boldsymbol{e}$ it necessarily holds that $\boldsymbol{\rho} \cdot \boldsymbol{\nu} = \boldsymbol{\sigma}(\lambda \boldsymbol{e})$. The theorem below follows immediately from Proposition 4 and (23).

THEOREM 2. *Given a C-node linear network with back-off rates $\boldsymbol{\nu}$, it holds that*

$$\lambda^*(\boldsymbol{\nu}) = \min\Big\{\frac{\nu_1}{1+2\nu_1}, \frac{\nu_C}{1+2\nu_C}, \frac{1}{2}-\frac{1}{2\sqrt{1+4\nu_c}}, 1<c<C\Big\}.$$

As discussed in the introduction and Section 3, the throughput of these linear multi-hop networks decreases in the offered load once a class becomes saturated. This is formalized in the the following theorem, which states that λ^* is in fact the maximum end-to-end throughput of the network.

THEOREM 3. *Given a C-node linear network with back-off rates $\boldsymbol{\nu}$, it holds that*

$$\max_{\lambda\geq 0}\theta\big(\boldsymbol{\rho};\lambda,\boldsymbol{\nu}\big) = \lambda^*(\boldsymbol{\nu}), \qquad \forall\, \boldsymbol{\rho}\in R(\lambda,\boldsymbol{\nu}).$$

The proof of Theorem 3 is provided in Appendix A.2. Theorems 2 and 3 immediately imply the following result.

COROLLARY 1. *Given a C-node linear network, it holds that*

$$\lim_{\boldsymbol{\nu}\to\infty}\Big(\max_{\lambda\geq 0}\theta\big(\boldsymbol{\rho};\lambda,\boldsymbol{\nu}\big)\Big) = \frac{1}{2}, \qquad \forall\, \boldsymbol{\rho}\in R(\lambda,\boldsymbol{\nu}).$$

This shows that if unrestricted in the choice of back-off rates, one can make any linear network throughput-optimal and maximally stable. However, the result does not provide any insight into what happens as λ grows larger than $\lambda^*(\boldsymbol{\nu})$, where the throughput could drastically drop. The next section provides an example of such a situation (see Corollary 2). In general we see that both λ^* and the behavior for $\lambda>\lambda^*$ strongly depend on $\boldsymbol{\nu}$.

4.2 Linear networks with uniform back-off rates

In this section we fix $\nu>0$ and let the vector of back-off rates be $\boldsymbol{\nu}^u=(\nu,\ldots,\nu)$. First, it follows from Theorem 2 that in this case

$$\lambda^*(\boldsymbol{\nu}^u) = \frac{1}{2}-\frac{1}{2\sqrt{1+4\nu}}.$$

We introduce the following constant:

$$\lambda^{**}(\nu) = \frac{\nu(1+2\nu)}{(1+3\nu)(1+\nu)}.$$

It is readily verified that $\lambda^*(\boldsymbol{\nu}^u)\leq\lambda^{**}(\nu)$ for every $\nu>0$. The following theorem, whose proof is provided in Appendix A.3, completely characterizes the stationary performance of a linear network with these uniform back-off rates.

THEOREM 4. *Given a C-node linear network with back-off rates $\boldsymbol{\nu}^u=(\nu,\ldots,\nu)$ and arrival rate λ, there exists a unique solution $\boldsymbol{\rho}\in R(\lambda,\boldsymbol{\nu}^u)$. For $\lambda\leq\lambda^*(\boldsymbol{\nu}^u)$, it holds that $\boldsymbol{s}=(\mathtt{U},\ldots,\mathtt{U})$ and*

$$\boldsymbol{\sigma}(\lambda\boldsymbol{e}) = \boldsymbol{\nu}^u\cdot\boldsymbol{\rho}.$$

For $\lambda^*(\boldsymbol{\nu}^u)<\lambda\leq\lambda^{**}(\nu)$, it holds that $\boldsymbol{s}=(\mathtt{U},\mathtt{S},\mathtt{U},\ldots,\mathtt{U})$ and

$$\rho_1 = \frac{\lambda(1-\lambda)}{\nu(1-\lambda-\lambda c(\lambda,\nu))}, \quad \rho_C = \frac{\lambda c(\lambda,\nu)(1-\lambda c(\lambda,\nu))}{\nu(1-2\lambda c(\lambda,\nu))},$$
$$\rho_2 = \frac{1}{c(\lambda,\nu)}, \qquad \rho_c = \frac{(1-\lambda-\lambda c(\lambda,\nu))}{(1-2\lambda c(\lambda,\nu))}, \quad c=3,\ldots,C-1,$$
$$c(\lambda,\nu) = \frac{1+(3-2\lambda)\nu-\sqrt{(1-2\lambda)^2\nu^2+2\nu+1}}{2(1+2\lambda\nu)}.$$

For $\lambda>\lambda^{**}(\nu)$, it holds that $\boldsymbol{s}=(\mathtt{S},\mathtt{S},\mathtt{U},\ldots,\mathtt{U})$ and

$$\rho_1 = \frac{\lambda(1+\nu)(1+3\nu)}{\nu(1+2\nu)}, \qquad \rho_2 = \frac{1+2\nu}{1+\nu},$$
$$\rho_C = \frac{1}{1+\nu}, \qquad \rho_c = \frac{1+2\nu}{(1+\nu)^2}, \quad c=3,\ldots,C-1.$$

The result of Theorem 4 describes the behavior of the system as λ increases, and can be understood as follows.

- Because classes $2,\ldots,C-1$ face more competition than the outer classes, the middle classes saturate first, at $\lambda=\lambda^*(\boldsymbol{\nu}^u)$.
- From this point onwards, class 2 acts as a bottleneck, reducing the arrival rate at the successive classes of nodes whose load decreases as λ increases. Specifically, note that $c(\lambda,\nu)$ decreases in λ and consequently ρ_c decreases for $c=3,\ldots,C$. On the other hand, as λ grows, ρ_1 increases and when the arrival rate matches $\lambda^{**}(\nu)$, class 1 also saturates.
- For $\lambda>\lambda^{**}(\nu)$ the system is oversaturated and increasing the load only further decreases ρ_1.

Since the end-to-end stationary throughput is given by $\theta\left(\lambda;\boldsymbol{\nu}^u\right)=\lambda\prod_{c=1}^{C}\rho_c^+$, the next corollary follows from Theorem 4.

COROLLARY 2. *Given a C-node linear network with back-off rates $\boldsymbol{\nu}^u$ and arrival rate λ, it holds that*

$$\theta\left(\lambda;\boldsymbol{\nu}^u\right) = \begin{cases}\lambda & \lambda\in[0,\lambda^*(\boldsymbol{\nu}^u)]\\ c(\lambda,\nu) & \lambda\in(\lambda^*(\boldsymbol{\nu}^u),\lambda^{**}(\nu)]\\ \frac{\nu}{1+3\nu} & \lambda>\lambda^{**}(\nu).\end{cases}$$

The corollary above implies that

$$\lim_{\nu\to\infty}\theta\big(\lambda,\boldsymbol{\nu}^u(\nu)\big) = \frac{1}{3}, \qquad \forall\,\lambda\geq\lim_{\nu\to\infty}\lambda^{**}(\nu)=\frac{2}{3}.$$

This illustrates the drop in the end-to-end throughput of the linear network with uniform back-off rates when one of the classes saturates.

4.3 Linear networks with fair back-off rates

In this section we fix a $\nu>0$ and consider a linear network with back-off rates given by $\boldsymbol{\nu}^f=(\nu,\nu(1+\nu),\ldots,\nu(1+\nu),\nu)$. It is known (see [21]) that these back-off rates guarantee fairness (i.e. equal throughputs of all nodes) in the saturated single-hop linear network. Hence, we refer to these as the 'fair' rates, and we conjecture that they may provide good performance in the multi-hop setting as well. This will be given a rigorous meaning and shown in Section 5.

It follows from Theorem 2 that in this case

$$\lambda^*(\boldsymbol{\nu}^f) = \frac{\nu}{1+2\nu} = \frac{1}{2}-\frac{1}{2\sqrt{1+4\nu(1+\nu)}}.$$

The following theorem completely characterizes the stationary performance of a linear network with fair back-off rates.

THEOREM 5. *Given a C-node linear network with fair back-off rates $\boldsymbol{\nu}^f$ and arrival rate λ, there exists a unique solution $\boldsymbol{\rho}\in R(\lambda,\boldsymbol{\nu}^f)$. For $\lambda\leq\lambda^*(\boldsymbol{\nu}^f)$, it holds that $\boldsymbol{s}=(\mathtt{U},\ldots,\mathtt{U})$ and*

$$\rho_1=\rho_C=\frac{\lambda}{\nu(1-2\lambda)}, \qquad \rho_c=\frac{\lambda(1-\lambda)}{\nu(1+\nu)(1-2\lambda)^2}$$

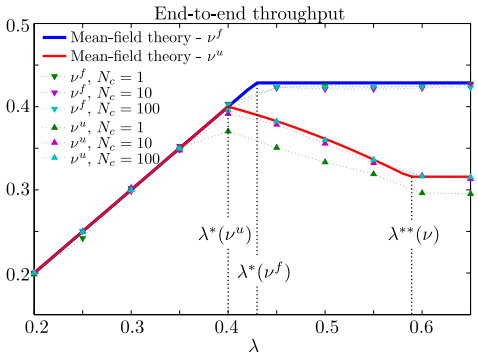

Figure 3: Throughput $\theta(\lambda, \boldsymbol{\nu})$ as a function of λ for $\boldsymbol{\nu} = \boldsymbol{\nu}^u, \boldsymbol{\nu}^f$ with $C = 3$. The mean-field approximations are plotted in solid blue and red. Simulation results for $N_c = 1, 10, 100$ nodes per class are marked.

with $c = 2, \ldots, C-1$. For $\lambda > \lambda^*(\boldsymbol{\nu}^f)$, it holds that $\boldsymbol{s} = (\mathtt{S}, \mathtt{U}, \ldots, \mathtt{U})$ and

$$\rho_1 = \frac{\lambda(1+2\nu)}{\nu}, \qquad \rho_c = 1, \quad 2 \le c \le C.$$

The argument of the proof follows the same structure as the proof of Theorem 4 and is not presented in this paper due to space restrictions.

Intuitively, when the back-off rates are equal to $\boldsymbol{\nu}^f$, all the classes saturate simultaneously when the arrival rate is equal to $\lambda^*(\boldsymbol{\nu}^f)$. This is possible since classes with more neighbors have shorter back-off periods. Then, as λ grows further, class 1 acts as a bottleneck, but since the system is already in oversaturation, this does not affect the throughput of the nodes in the system and only ρ_1 further increases.

Theorem 5 allows us to characterize the end-to-end stationary throughput $\theta(\lambda; \boldsymbol{\nu}^f)$.

COROLLARY 3. *Given a C-node linear network with back-off rates $\boldsymbol{\nu}^f$ and arrival rate λ, it holds that*

$$\theta\left(\lambda; \boldsymbol{\nu}^f\right) = \begin{cases} \lambda, & \lambda \le \lambda^*(\boldsymbol{\nu}^f), \\ \frac{\nu}{1+2\nu}, & \lambda > \lambda^*(\boldsymbol{\nu}^f). \end{cases}$$

Uniform vs optimal back-off rates. In Figure 3 we plot the end-to-end throughput of a 3-class linear network against the offered load at class 1. We compare the uniform and fair rates, with $\boldsymbol{\nu}^u = (6, 6, 6)$ and $\boldsymbol{\nu}^f = (3, 12, 3)$. We show both the results from the mean-field limit (from Corollaries 2 and 3) and simulation results for the network with $N_c = 1, 10, 100$ nodes per class. Note that $\sum_c \boldsymbol{\nu}_c^u = \sum_c \boldsymbol{\nu}_c^f = 18$ and therefore it seems reasonable to compare the performance of the networks with these back-off vectors.

The figure provides further confirmation that the equilibrium point of the mean-field limit provides a good approximation for the pre-limit network, even for small N_c. We also see that carefully selecting the back-off rates yields a relevant improvement both in the throughput and the stability performance. In fact, from from Corollaries 2 and 3 we see that $\lambda^*(\boldsymbol{\nu}^f) = 3/7 > 2/5 = \lambda^*(\boldsymbol{\nu}^u)$, and $\theta(\lambda; \boldsymbol{\nu}^f) \ge \theta(\lambda; \boldsymbol{\nu}^u)$ for every $\lambda \ge 0$.

4.4 Linear networks in oversaturation

In this section we restrict ourselves to the case with $C = 3$, and let $\lambda \to \infty$ so class 1 is saturated. This allows us to find the equilibrium point for general $\boldsymbol{\nu}$. We stress that even a simple case like this one is only partially tractable for finite N, see [18].

Let us introduce the following subsets of $\mathbb{R}^3$,

$$S_1 = \Big\{\frac{\nu_2}{1+\nu_1} \ge \nu_1,\ \nu_1 \le \nu_3\Big\}, \quad S_2 = \Big\{\frac{\nu_2}{1+\nu_1} < \nu_1,\ \frac{\nu_2}{1+\nu_1} \le \nu_3\Big\},$$
$$S_3 = \Big\{\frac{\nu_2}{1+\nu_3} \ge \nu_1,\ \nu_1 > \nu_3\Big\}, \quad S_4 = \Big\{\frac{\nu_2}{1+\nu_3} < \nu_1,\ \frac{\nu_2}{1+\nu_1} > \nu_3\Big\},$$

which provide a partition of $\mathbb{R}^3_+$.

THEOREM 6. *Consider a 3-class linear network with back off rates $\boldsymbol{\nu} \in \mathbb{R}^3_+$, in oversaturation there exists a unique solution $\boldsymbol{\rho}$ for (19) and*

$$\boldsymbol{\nu} \in S_1 \ \Rightarrow\ \boldsymbol{s} = (\mathtt{S}, \mathtt{U}, \mathtt{U}), \qquad \boldsymbol{\nu} \in S_2 \ \Rightarrow\ \boldsymbol{s} = (\mathtt{S}, \mathtt{S}, \mathtt{U}),$$
$$\boldsymbol{\nu} \in S_3 \ \Rightarrow\ \boldsymbol{s} = (\mathtt{S}, \mathtt{U}, \mathtt{S}), \qquad \boldsymbol{\nu} \in S_4 \ \Rightarrow\ \boldsymbol{s} = (\mathtt{S}, \mathtt{S}, \mathtt{S}).$$

The detailed description of the unique solution $\boldsymbol{\rho}$ for (19) is provided in the proof given in Appendix A.4. Hence, the following corollary immediately follows.

COROLLARY 4. *Consider a 3-class linear network with back off rates $\boldsymbol{\nu} \in \mathbb{R}^3_+$, then*

$$\lim_{\lambda\to\infty} \theta\left(\lambda; \boldsymbol{\nu}\right) = \begin{cases} \frac{\nu_1}{2\nu_1+1}, & \boldsymbol{\nu} \in S_1, \\ \frac{\nu_2}{1+\nu_1+2\nu_2}, & \boldsymbol{\nu} \in S_2, \\ \frac{\nu_3(1+\nu_1)}{(1+2\nu_1)(1+\nu_3)}, & \boldsymbol{\nu} \in S_3, \\ \frac{\nu_3(1+\nu_1)}{1+\nu_1+\nu_2+\nu_3+\nu_1\nu_3}, & \boldsymbol{\nu} \in S_4. \end{cases}$$

5. OPTIMAL BACK-OFF RATES FOR LINEAR NETWORKS

In Section 4 we provided a detailed analysis of linear networks with different back-off rates. In particular, in Section 4.3 we observed that the 'fair' back-off rates (i.e., those rates that provide strict fairness in the setting where all nodes are saturated) outperform uniform back-off rates, both in the sense that the network can sustain higher loads, and that it behaves better in overload. In this section we investigate what are the 'best' back-off rates. We still restrict ourselves to linear networks with nearest-neighbor blocking, although many of the results derived here hold more generally. It is worth repeating that we believe a unique fixed point exists for all λ and $\boldsymbol{\nu}$, i.e., $|R(\lambda, \boldsymbol{\nu})| = 1$, although we cannot prove this. We now introduce two desirable properties of back-off rates: robustness and throughput optimality.

5.1 Robust back-off rates

In Section 4.2, it is shown that for uniform back-off rates, the throughput of a linear network decreases when it is in overload. This is detrimental to the network performance, and we aim to find so-called 'robust' back-off rates such that the throughput performance of the system in overload is identical to the case with arrival rate $\lambda^*(\boldsymbol{\nu})$.

DEFINITION 2. *The back-off rates $\boldsymbol{\nu}$ are robust if and only if for every $\lambda \ge \lambda^*(\boldsymbol{\nu})$ there exists a solution $\boldsymbol{\rho} \in R(\lambda, \boldsymbol{\nu})$ such that $\theta(\boldsymbol{\rho}; \lambda, \boldsymbol{\nu}) = \lambda^*(\boldsymbol{\nu})$.*

We first present the following sufficient condition for back-off rates to be robust.

PROPOSITION 5. *If* $\nu_1 \leq \nu_C$ *and*

$$\frac{\nu_1}{1+2\nu_1} \leq \frac{1}{2} - \frac{1}{2\sqrt{1+4\nu_c}}, \qquad c = 2, \ldots, C-1,$$

then the back-off rates $\boldsymbol{\nu}$ *are robust.*

PROOF. Note that $\tilde{\boldsymbol{\rho}}(\lambda) \in R(\lambda, \boldsymbol{\nu})$ for every $\lambda \leq \lambda^*(\boldsymbol{\nu})$, where $\tilde{\boldsymbol{\rho}}(\lambda) \cdot \boldsymbol{\nu} = \boldsymbol{\sigma}(\lambda \boldsymbol{e})$ and $\boldsymbol{\sigma}(\lambda \boldsymbol{e})$ as in Lemma 1. Due to the hypothesis, it holds that

$$\tilde{\boldsymbol{\rho}}\big(\lambda^*(\boldsymbol{\nu})\big) = \Big(1, \tilde{\rho}_2\big(\lambda^*(\boldsymbol{\nu})\big), \ldots, \tilde{\rho}_C\big(\lambda^*(\boldsymbol{\nu})\big)\Big) \leq \boldsymbol{e}.$$

Now, for every $\lambda \geq \lambda^*(\boldsymbol{\nu})$, define $\boldsymbol{\rho}(\lambda)$ as

$$\boldsymbol{\rho}\big(\lambda\big) = \Big(\frac{\lambda}{\lambda^*(\boldsymbol{\nu})}, \tilde{\rho}_2\big(\lambda^*(\boldsymbol{\nu})\big), \ldots, \tilde{\rho}_C\big(\lambda^*(\boldsymbol{\nu})\big)\Big),$$

and we now show that $\boldsymbol{\rho}\big(\lambda\big) \in R(\lambda, \boldsymbol{\nu})$ and $\theta(\boldsymbol{\rho}(\lambda); \lambda, \boldsymbol{\nu}) = \lambda^*(\boldsymbol{\nu})$. For $c = 1$, it holds that

$$\lambda \rho_1(\lambda)^+ = \frac{\lambda}{\rho_1(\lambda)} = \lambda^*(\boldsymbol{\nu}) = \bar{\theta}_1(\tilde{\boldsymbol{\rho}}(\lambda^*(\boldsymbol{\nu}))^- \cdot \boldsymbol{\nu}) = \bar{\theta}_1(\boldsymbol{\rho}(\lambda)^- \cdot \boldsymbol{\nu}),$$

and for $c > 1$, it holds that

$$\lambda \prod_{d=1}^{c} \rho_d(\lambda)^+ = \lambda^*(\boldsymbol{\nu}) = \bar{\theta}_c(\tilde{\boldsymbol{\rho}}(\lambda^*(\boldsymbol{\nu}))^- \cdot \boldsymbol{\nu}) = \bar{\theta}_c(\boldsymbol{\rho}(\lambda)^- \cdot \boldsymbol{\nu}),$$

where the last equality is due to the definition of $\boldsymbol{\rho}(\lambda)$. Hence, $\boldsymbol{\rho}(\lambda) \in R(\lambda, \boldsymbol{\nu})$ and yields throughput $\theta(\boldsymbol{\rho}(\lambda); \lambda, \boldsymbol{\nu}) = \bar{\theta}_C(\boldsymbol{\rho}(\lambda)^- \cdot \boldsymbol{\nu}) = \lambda^*(\boldsymbol{\nu})$. □

We would like to point out that robust back-off rates are also *fair* in the sense that the various classes transmit for the same fraction of time with the system in equilibrium. In Proposition 2 we observed that when the network is stable, i.e., $\lambda \leq \lambda^*(\boldsymbol{\nu})$, the network self-regulates and behaves in a fair way, in fact $\theta_c(\lambda, \boldsymbol{\nu}) = \lambda$, for every class $c \in \mathcal{C}$. In general, the fairness of the network breaks for $\lambda > \lambda^*(\boldsymbol{\nu})$, however, that is not the case when the back-off rates are robust. Consider $\lambda > \lambda^*(\boldsymbol{\nu})$ and the solution $\boldsymbol{\rho} \in R(\lambda, \boldsymbol{\nu})$ presented in the proof of Proposition 5. It holds that

$$\theta_c(\boldsymbol{\rho}; \lambda, \boldsymbol{\nu}) = \bar{\theta}_c(\boldsymbol{\rho}^- \cdot \boldsymbol{\nu}) = \lambda^*(\boldsymbol{\nu}), \qquad \forall\, c \in \mathcal{C}.$$

5.2 Throughput-optimal back-off rates

In Corollary 1 we observed that the maximum throughput achievable approaches 1/2 when all the coordinates of $\boldsymbol{\nu}$ grow to infinity. However, in practice the back-off rates are bounded by physical constraints. Otherwise, with non-negligible probability, there will be arbitrarily small gaps between the ends of back-off periods of neighboring nodes, which may lead to collisions. We model this as an upper bound on the aggregate back-off rates allocated among the nodes. Putting an upper bound on the back-off rate of each node would lead to similar results as those presented below.

Hence, we impose a budget $V > 0$, and consider the set of admissible back-off rates $\boldsymbol{\nu}$ to be $\mathcal{V} = \{\boldsymbol{\nu} \in \mathbb{R}_+^C : \sum_{c=1}^C \nu_c \leq V\}$. Among the admissible back-off rates, we can identify those for which the maximum throughput is the highest.

DEFINITION 3. *The back-off rates* $\boldsymbol{\nu}'$ *are throughput-optimal in the set* $\mathcal{V}$ *if and only if* $\lambda^*(\boldsymbol{\nu}') = \max_{\boldsymbol{\nu} \in \mathcal{V}} \{\lambda^*(\boldsymbol{\nu})\}$.

Denote by $\lambda^* = \max_{\boldsymbol{\nu} \in \mathcal{V}} \{\lambda^*(\boldsymbol{\nu})\}$ the maximum throughput achievable over all back-off rates and arrival rates. Let us now consider the fair back-off rates already introduced in Section 4.3, and recall $\boldsymbol{\nu}^f(\nu) = (\nu, \nu(1+\nu), \ldots, \nu(1+\nu), \nu)$, $\nu > 0$. Define ν^* as the unique positive solution of $V = ||\boldsymbol{\nu}^f(\nu^*)||_1$, and $\boldsymbol{\nu}^* = \boldsymbol{\nu}^f(\nu^*)$, i.e., the highest possible fair rates allowed by the budget. We can show that $\boldsymbol{\nu}^*$ is both robust and throughput-optimal.

THEOREM 7. *The back-off rates* $\boldsymbol{\nu}^*$ *are both robust and throughput-optimal.*

PROOF. The robustness follows immediately from Proposition 5. In fact, $\nu_1^* = \nu_C^*$ and

$$\frac{\nu_1^*}{1+2\nu_1^*} = \frac{\nu}{1+2\nu} = \frac{1}{2} - \frac{1}{2\sqrt{1+4\nu(\nu+1)}} = \frac{1}{2} - \frac{1}{2\sqrt{1+4\nu_c^*}},$$

for $c = 2. \ldots, C-1$.

It remains to be proved that $\lambda^*(\boldsymbol{\nu}^*) = \lambda^*$. Consider $\boldsymbol{\rho} \in R(\lambda^*, \boldsymbol{\nu})$ with $\boldsymbol{\nu} \in \mathcal{V}^*$ and $\boldsymbol{\rho} \leq \boldsymbol{e}$. Define $\tilde{\boldsymbol{\nu}} = \boldsymbol{\rho} \cdot \boldsymbol{\nu}$. We will show that $\tilde{\boldsymbol{\nu}} = \boldsymbol{\nu}^*$. Since $\tilde{\boldsymbol{\nu}} \leq \boldsymbol{e}$, we have $\bar{\theta}_c(\boldsymbol{e} \cdot \tilde{\boldsymbol{\nu}}) = \bar{\theta}_c(\boldsymbol{\rho} \cdot \boldsymbol{\nu}) = \lambda^*$ for all $c \in \mathcal{C}$, so that $\tilde{\boldsymbol{\nu}} \in \mathcal{V}^*$. Since $\boldsymbol{e} \in R(\lambda^*, \tilde{\boldsymbol{\nu}})$, there necessarily exist $\nu > 0$ such that $\tilde{\boldsymbol{\nu}} = \boldsymbol{\nu}^{f,\nu}$, so

$$\lambda^*(\boldsymbol{\nu}^*) = \lambda^*(\boldsymbol{\nu}^f(\nu^*)) \geq \lambda^*(\boldsymbol{\nu}^f(\nu)) = \lambda^*(\tilde{\boldsymbol{\nu}}) = \lambda^*.$$

Since $\boldsymbol{\nu}^* \in \mathcal{V}$, we also have $\lambda^*(\boldsymbol{\nu}^*) \leq \lambda^*$, and therefore $l(\nu^*) = \lambda^*$. □

As a consequence of the above result, we deduce that for every arrival rate λ, the back-off rates $\boldsymbol{\nu}^*$ outperform any other $\boldsymbol{\nu} \in \mathcal{V}$. In fact, for any λ there exists $\boldsymbol{\rho}^* \in R(\lambda, \boldsymbol{\nu}^*)$ such that for every $\boldsymbol{\rho} \in R(\lambda, \boldsymbol{\nu}^*)$, if $\lambda < \lambda^*$ it holds that

$$\theta\big(\boldsymbol{\rho}^*; \lambda, \boldsymbol{\nu}^*\big) = \lambda \geq \theta\big(\boldsymbol{\rho}; \lambda, \boldsymbol{\nu}\big),$$

while, if $\lambda \geq \lambda^*$ we have that

$$\theta\big(\boldsymbol{\rho}^*; \lambda, \boldsymbol{\nu}^*\big) = \lambda^* \geq \lambda^*(\boldsymbol{\nu}) \geq \theta\big(\boldsymbol{\rho}; \lambda, \boldsymbol{\nu}\big).$$

6. GENERAL NETWORKS

The detailed analysis of linear networks in Sections 4 and 5 highlights the importance of the choice of back-off rates when operating a multi-hop network. In this section we exploit the results and intuition from these sections to find optimal back-off rate vectors for general networks, subject to a budget for the aggregate rates. We then compare this with the uniform back-off rates and we show substantial gains achieved both in terms of throughput and stability.

6.1 Heuristic optimal back-off rates

We have seen in Sections 4 and 5 that for linear networks, all fair rates are robust, and that the highest fair rates allowed by our budget are also throughput-optimal. The intuition behind this is that under the fair rates all classes saturate simultaneously, which both ensures robustness, and that the point of saturation λ^* is as high as possible.

This seems to suggest that for general networks, too, the highest possible fair rates are both robust and throughput-optimal. Although we cannot prove this, we can compute these rates numerically, and show that they perform very well compared to uniform rates. Recall that $\boldsymbol{\sigma}$ denotes the inverse throughput mapping for saturated networks, so that for any $\lambda > 0$ such that $\lambda \boldsymbol{e} \in \text{int}(\Gamma)$, the fair rates such that each class has throughput λ are given by $\boldsymbol{\sigma}(\lambda \boldsymbol{e})$. The key is then to find the highest possible such λ that can be afforded from our budget:

$$(P) \; : \quad \max\big\{\lambda : V \geq ||\boldsymbol{\sigma}(\lambda \boldsymbol{e})||_1\big\}. \tag{24}$$

In general networks, explicit formulae for $\boldsymbol{\nu}^{\lambda}$ are not available. Instead, in Section 6.2 we solve (24) by bisection, where the $\boldsymbol{\sigma}(\lambda \boldsymbol{e})$ in each iteration is solved numerically or through approximations described briefly below. In Section 6.2 we will discuss the performance of multi-hop networks with back-off rates $\boldsymbol{\nu}^*$ computed via the different methods.

6.1.1 *Exact method*

In [21] a numerical algorithm is presented that computes $\boldsymbol{\sigma}(\gamma)$. The main drawback of such algorithm is that it needs to know all the independent sets of the interference graph whose number grows exponentially in C. We stress that this problem is NP-Hard and the method works only when C is reasonably small.

6.1.2 *Approximation methods*

Due to the computational complexity of the exact method, there has been recently a lot of interest in obtaining accurate approximation for the back-off rates yielding a target throughput, see [12, 13, 19, 24]. The most popular methods are based on free energy approximations due to Bethe [24] and Kikuchi [12, 19] which are shown to be exact for acyclic and chordal networks respectively. In [13] a theoretical framework covering both methods is presented.

6.2 Numerical examples

In this section, we test the performance of the heuristic method presented above and compare it with the uniform back-off rates. We aim to demonstrate that the back-off rates obtained by solving (24) exhibit great performance not only in linear networks, but in networks with general interference graphs as well.

6.2.1 *Small network*

Consider the 5-class multi-hop network in Figure 4(a). We fix $V = 10$, and as the network is small, we compute $\boldsymbol{\nu}^*$ via the numerical method described in [21] and obtain $\boldsymbol{\nu}^* = [1.075, 3.642, 1.76, 1.763, 1.76]$ On the other hand, the uniform back-off rates are $\boldsymbol{\nu}^u$ with $\nu_c^u = 2$ for each class.

In Figure 5 we present the end-to-end throughput performance for the system with the different back-off rates and both 1 and 100 nodes per class. First of all, we observe the striking similarity of these two plots, the qualitative behavior of the end-to-end throughput as a function of the arrival rate λ is nearly identical. For this small network, the fixed-point iterative method applied to (25) converges quickly for every choice of λ and both $\boldsymbol{\nu}^*$ and $\boldsymbol{\nu}^u$. We compare the simulated results with the throughput expected from the equilibrium point of the mean-field differential equation. There is a slight difference in the end-to-end throughput for $N_c = 1$ and $\boldsymbol{\nu} = \boldsymbol{\nu}^u$ for $\lambda > 0.3$, but the accuracy of the mean-field approximation is nevertheless impressive.

In Figure 4(b) we show the throughput of the different classes with arrival rate $\lambda = 0.5$. We display the scenario with a single node per class, the case with 100 nodes per class exhibits similar behavior. Note that the system initialized with $\boldsymbol{\nu} = \boldsymbol{\nu}^*$ is extremely fair even in such an oversaturated regime, corroborating the idea that the heuristic method we proposed is robust. When $\boldsymbol{\nu} = \boldsymbol{\nu}^u$, the throughput of class-1 nodes is very high (it matches the arrival rate $\lambda = 0.5$), however the packets are then stuck in the buffer of class-2 nodes which are not equipped to sustain the incoming flow. Note that since class 1 is still unsaturated, a further worsening of the performance is foreseeable as the arrival rate grows larger.

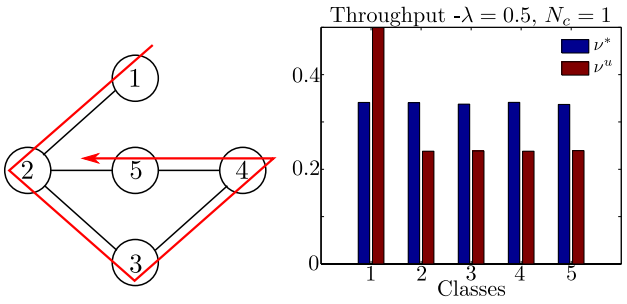

Figure 4: (a) 5-class network, (b) Per-class throughput with back-off rates $\boldsymbol{\nu}^*, \boldsymbol{\nu}^u$ and $N_c = 1$ node per class.

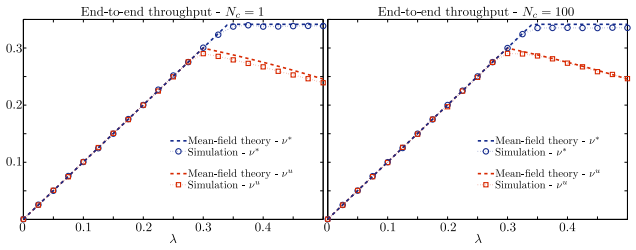

Figure 5: End-to-end throughput performance as a function of λ for back-off rates $\boldsymbol{\nu}^*, \boldsymbol{\nu}^u$. (a) $N_c = 1$ (b) $N_c = 100$.

6.2.2 *Large network*

Consider now a large network with 20 classes with average number of neighbors per class equal to 10.4 and maximum clique dimension equal to 6. We fix $V = 100$, however the graph is too large to use the exact method described in [21], we thus rely on Kikuchi's method described in [12] and obtain the back-off rates $\boldsymbol{\nu}^K$. The graph considered is not chordal, but we expect Kikuchi's approximation to be fairly accurate since there are no long cycles without internal chords. The uniform back-off rates $\boldsymbol{\nu}^u$ have $\nu_c^u = 5$ for every $c = 1, \ldots, 20$.

We simulate the network with $N_c = 100$ nodes per class.

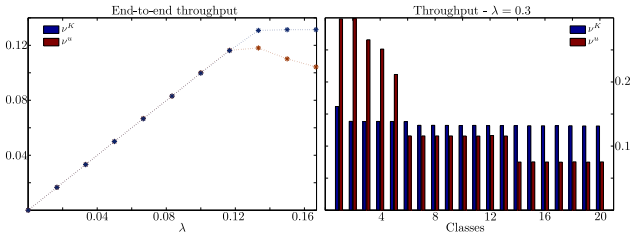

Figure 6: 20-class network with $N_c = 100$ nodes per class. (a) End-to-end throughput performance as a function of λ for back-off rates $\boldsymbol{\nu}^K$ and $\boldsymbol{\nu}^u$. (b) Per-class throughput performance for back-off rates $\boldsymbol{\nu}^K$ and $\boldsymbol{\nu}^u$ with arrival rate $\lambda = 0.3$.

In Figure 6(a) we show the end-to-end throughput performance of the system as λ increases. Note that up to arrival rate $\lambda = 0.12$ the networks with both back-off rates are stable and the end-to-end throughput rate coincides with the arrival rate. As λ grows further, in the system with uniform back-off rates some classes saturate and act as bottlenecks, since the network is not in oversaturation, i.e., class 1 is still stable, the performance drastically deteriorates. To better appreciate this, in Figure 6(b) we show the throughput of the various classes with arrival rate $\lambda = 0.3$. Observe that the throughput profile for the network with back-off rates $\boldsymbol{\nu}^K$ is quite fair and $\theta_1(\lambda, \boldsymbol{\nu}^K) - \theta_{20}(\lambda, \boldsymbol{\nu}^K)$ is relatively small. The throughput profile for the network with back-off rates $\boldsymbol{\nu}^u$ is very different. Nodes $3, 4, 5, 14$ are clearly bottlenecks and block a large fraction of the incoming flow. The difference $\theta_1(\lambda, \boldsymbol{\nu}^u) - \theta_{20}(\lambda, \boldsymbol{\nu}^u)$ is now very large. Note that class-1 is still unsaturated, hence as λ grows further we expect the end-to-end throughput to exhibit further deterioration. Already with $\lambda = 0.3$, the network with $\boldsymbol{\nu} = \boldsymbol{\nu}^K$ transmits roughly twice the number of packets that the network with $\boldsymbol{\nu} = \boldsymbol{\nu}^u$ does.

7. CONCLUSION AND OUTLOOK

In this paper we considered random-access networks where packets may be forwarded between nodes before eventually leaving the network. This multi-hop structure causes various performance issues, and makes the exact analysis of the underlying stochastic process intractable. In this paper we considered the mean-field regime, where the number of nodes grows large. We showed that in this regime the stochastic process converges to an IVP, and introduced the concept of partial equilibrium points, where certain classes of nodes may be in a stable equilibrium, while other classes are unstable. The mean-field regime is inspired by emerging, highly dense random-access networks such as mesh networks and the device-to-device mode in 5G, but we showed numerically that the equilibrium points of the mean-field limit also provide an excellent approximation for sparse networks with fewer nodes.

We then used the mean-field limit to better understand multi-hop networks, and to improve their performance. We first considered linear networks, and showed that the end-to-end throughput as a function of the offered load is the highest at the point when the first class of nodes saturates, and that this saturation point and the behavior in overload depend on the back-off rates of the nodes. We fully characterized the behavior at the equilibrium points for uniform back-off rates and 'fair' back-off rates, i.e., those that provide equal throughputs in the case all nodes are saturated. The uniform back-off rates turned out out to display a decrease in throughput when the network is overloaded, while the fair rates are robust, in the sense that the throughput remains the same in overload. We then looked at the problem of finding the best back-off rates given a certain total budget (representing for instance physical constraints on the back-off processes), and proved that the fair back-off rates are indeed optimal in the sense that they are both robust and give the largest stability region. We then used these results to devise a heuristic method for choosing back-off rates in general multi-hop networks, and we showed numerically that these rates perform well compared to uniform back-off rates.

7.1 Outlook

Although multi-hop random-access networks are becoming increasingly widespread and form the backbone of future technologies such as mesh networks, IoT and 5G, their performance is still poorly understood. The results presented in this paper may be viewed as a first step towards developing tools for the performance analysis and optimization of such networks. Below we outline some possible future research directions.

Rigorous treatment of partial stability. Although the partial equilibrium points defined in Definition 1 are numerically verified to be the correct choice, we would like to derive them in a more rigorous manner, analogous to the 'pure' equilibrium points from Proposition 1. This may be done by relying on the literature on partially stable ODEs (cf. [22]), which covers ODEs that are stable only in certain dimensions.

Extension beyond linear networks. Sections 4 and 5 are concerned with the partial equilibrium points for linear networks with nearest neighbor blocking. Many of the results here can be readily extended to general linear networks, but the key challenge will be to extend it beyond linear networks. For instance the proof of Theorem 3 relies on the Markov random field representation of the network, which cannot be used in general.

Networks with multiple traffic flows. Here we limited ourselves to a single flow of traffic traversing the network, but the derivation of the mean-field limit (4) can be readily extended to networks with multiple traffic flows, although the derivation of the partial equilibrium points (19) becomes more complex. This raises several very interesting new questions, for instance whether the uniqueness of $\boldsymbol{\rho}$-solutions to (4) we observed numerically still holds for two or more flows, or if there are multiple configurations possible with different bottlenecks and a dominant flow.

Adaptive CSMA algorithms in multi-hop networks. The past decade has seen the emergence of a class of adaptive CSMA algorithms that allow nodes in saturated or single-hop networks to find their optimal back-off rate in a distributed manner, see, e.g., [14, 10, 17]. It is unclear how well these algorithms work in multi-hop networks, and how the resulting coupling of departure and arrival processes of neighboring nodes affects their performance. The mean-field framework here could be extended to state-dependent back-off rates (see [6]), which would allow insights in these adaptive algorithms in a more realistic multi-hop setting.

8. REFERENCES

[1] A. Aziz, S. Shneer, P. Thiran (2013). Wireless multi-hop networks beyond capacity. In: *Proc. LANMAN.*

[2] R.R. Boorstyn, A. Kershenbaum, B. Maglaris, V. Sahin (1987). Throughput analysis in multihop CSMA packet radio networks. *IEEE Trans. Commun.*, **35**, 267–274.

[3] C. Bordenave, D. McDonald, A. Proutiere (2008). Performance of random medium access control, an asymptotic approach. *ACM SIGMETRICS Perf. Eval. Rev.*, **36 (1)**, 1–12.

[4] F. Cecchi, S.C. Borst, J.S.H. van Leeuwaarden (2014). Throughput of CSMA networks with buffer dynamics. *Perf. Eval.*, **79**, 216–234.

[5] F. Cecchi, S.C. Borst, J.S.H. van Leeuwaarden, P.A. Whiting (2016). CSMA networks in a many-sources regime: A mean-field approach. In: *Proc. IEEE Infocom.*

[6] F. Cecchi, S.C. Borst, J.S.H. van Leeuwaarden, P.A. Whiting (2016). Mean-field limits for large-scale random-access networks. arXiv preprint arXiv:1611.09723.

[7] J. Cho, J.Y. Le Boudec, Y. Jiang (2012). On the asymptotic validity of the decoupling assumption for analyzing 802.11 MAC protocol. *IEEE Trans. Inf. Theo.*, **58 (11)**, 6879–6893.

[8] M. Durvy, O. Dousse, P. Thiran (2009). Self-organization properties of CSMA/CA systems and their consequences on fairness. *IEEE Trans. Inf. Theo.*, **55 (3)**, 931–943.

[9] D. Evans (2011). The internet of things. *How the next evolution of the internet is changing everything*, Whitepaper, Cisco (IBSG).

[10] J. Ghaderi, R. Srikant (2010). On the design of effcient CSMA algorithms for wireless networks. In: *Proc. IEEE Dec. Contr.*, 954–959.

[11] P.T. Harker, J.S. Pang (1990). Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. *Math. Progr.*, **48 (1-3)**, 161–220.

[12] B. van Houdt (2016). Explicit back-off rates for achieving target throughputs in CSMA/CA networks. *IEEE/ACM Trans. Netw.*, **25 (2)**, 765–778.

[13] B. van Houdt (2017). Free Energy Approximations for CSMA networks. arXiv preprint arXiv:1703.10500.

[14] L. Jiang, J. Walrand (2010). A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Trans. Netw.*, **18 (3)**, 960–972.

[15] R. Laufer, L. Kleinrock (2016). The capacity of wireless CSMA/CA networks. *IEEE/ACM Trans. Netw.*, **24 (3)**, 1518–1532.

[16] S.C. Liew, C.H. Kai, J. Leung, B. Wong (2010). Back-of-the-envelope computation of throughput distributions in CSMA wireless networks. *IEEE Trans. Mob. Comp.* **9 (9)**, 1319–1331.

[17] S. Rajagopalan, D. Shah, J. Shin (2009). Network adiabatic theorem: an efficient randomized protocol for content resolution. In: *Proc. ACM SIGMETRICS Perf. Eval. Rev.*, 133–144.

[18] S. Shneer, P.M. van de Ven (2015). Stability and instability of individual nodes in multi-hop wireless CSMA/CA networks. In: *Proc. ACM SIGMETRICS Perf. Eval. Rev.*, **43 (2)**, 19–21.

[19] P.S. Swamy, V.P.K. Bellam, R.K. Ganti, K. Jagannathan (2016). Efficient CSMA based on Kikuchi approximation. In: *Proc. IEEE SPCOM.*

[20] P.M. van de Ven, S.C. Borst, J.S.H. van Leeuwaarden, A. Proutière (2010). Insensitivity and stability of random-access networks. *Perf. Eval.*, **67 (11)**, 1230–1242.

[21] P.M. van de Ven, A.J.E.M. Janssen, J.S.H. van Leeuwaarden, S.C. Borst (2011). Achieving target throughputs in random-access networks. *Perf. Eval.*, **68 (11)**, 1103–1117.

[22] V. Vorotnikov (2012). Partial stability and control. Springer Science & Business Media

[23] X. Wang, K. Kar (2005). Throughput modeling and fairness issues in CSMA/CA based ad-hoc networks. In: *Proc. IEEE Infocom.*

[24] S. Yun, J. Shin, Y. Yi (2015). CSMA using the Bethe approximation: scheduling and utility maximization. *IEEE Trans. Inf. Theo.*, **61 (9)**, 4776–4787.

APPENDIX

A. REMAINING PROOFS

A.1 Proof of Proposition 3

PROOF. Define the continuous function $\boldsymbol{F} : \mathbb{R}_+^C \to \mathbb{R}_+^C$ where

$$F_c(\boldsymbol{\rho}) = \frac{\lambda \prod_{d=1}^{c-1} \rho_d^+}{\nu_c \pi(\Omega_{-c}; \boldsymbol{\rho}^- \cdot \boldsymbol{\nu})}, \tag{25}$$

and observe that (19) can we rewritten to $F_c(\boldsymbol{\rho}) = \rho_c$. Define the convex and bounded set $\mathcal{U} \subseteq \mathbb{R}_+^C$ as

$$\mathcal{U} = \{\boldsymbol{\rho} \in \mathbb{R}_+^C : 0 \leq \rho_c \leq U_c, \, \forall c \in \mathcal{C}\}, \qquad U_c = \frac{\lambda}{\nu_c \pi(\mathbf{0}; \boldsymbol{\nu})}.$$

Below we show that $\boldsymbol{F}(\cdot)$ maps $\mathbb{R}_+^C$ in $\mathcal{U}$. Hence the result then follows from Brouwer's Theorem which establishes that there must exist a solution for $\boldsymbol{F}(\boldsymbol{\rho}) = \boldsymbol{\rho}$ in $\mathcal{U}$.

Given $\boldsymbol{\rho} \in \mathbb{R}_+^C$, we have to show that $0 \leq F_c(\boldsymbol{\rho}) \leq U_c$ for every $c \in \mathcal{C}$. Observe that $0 \leq F_c(\boldsymbol{\rho})$ clearly follows from the non-negativity of all the terms on the right hand side of (25). On the other hand, since $\boldsymbol{\rho}^+ \leq \boldsymbol{e}$ and $\mathbf{0} \in \Omega_{-c}$, it holds that

$$F_c(\boldsymbol{\rho}) \leq \frac{\lambda}{\nu_c \pi(\Omega_{-c}; \boldsymbol{\rho}^- \cdot \boldsymbol{\nu})} \leq \frac{\lambda}{\nu_c \pi(\mathbf{0}; \boldsymbol{\rho}^- \cdot \boldsymbol{\nu})}.$$

We conclude that $F_c(\boldsymbol{\rho}) \leq U_c$ since $\boldsymbol{\rho}^- \geq \mathbf{0}$ and thus

$$\frac{1}{\pi(\mathbf{0}; \boldsymbol{\rho}^- \cdot \boldsymbol{\nu})} = \sum_{\boldsymbol{\omega} \in \Omega} \prod_{c \in \mathcal{C}} \left(\nu_c \rho_c^-\right)^{\omega_c} \leq \sum_{\boldsymbol{\omega} \in \Omega} \prod_{c \in \mathcal{C}} \nu_c^{\omega_c} = \frac{1}{\pi(\mathbf{0}; \boldsymbol{\nu})}.$$

□

This proof also indicates that solutions $\boldsymbol{\rho}$ to (19) can be found through fixed-point iteration of (25). Although we cannot show that this iteration converges, numerically we see this always is the case.

A.2 Proof of Theorem 3

PROOF. Let us look at all the possible scenarios where the back-off rate of class c is equal to $\rho_c(\lambda, \boldsymbol{\nu})^- \nu_c = \overline{\nu}_c \leq \nu_c$. The idea is to show that for any set of throughputs $\bar{\theta}_1(\bar{\boldsymbol{\nu}}), \ldots, \bar{\theta}_C(\bar{\boldsymbol{\nu}})$ such that $\bar{\theta}_c(\bar{\boldsymbol{\nu}}) \geq \bar{\theta}_{c+1}(\bar{\boldsymbol{\nu}})$ we will have $\bar{\theta}_C(\bar{\boldsymbol{\nu}}) \leq \lambda^*(\boldsymbol{\nu})$. If we show that, we will show that the throughputs after saturation cannot be bigger than $\lambda^*(\boldsymbol{\nu})$ which is an achievable throughput before saturation. Let us fix $\bar{\boldsymbol{\nu}}$ and drop it from the notation.

We use formulae from the proof of [21, Prop. 3]. In particular, consider a saturated network with C nodes and back-off rates $\bar{\boldsymbol{\nu}}$, and denote $a_c = \mathbb{P}\{Y_c = 0 | Y_{c-1} = 0\}$ and $\psi_c = \mathbb{P}\{Y_{c-1} = 0\}$ with $c = 2, \ldots, C$. Since $\bar{\theta}_c = \mathbb{P}\{Y_{c-1} = 1\}$, we have $\bar{\theta}_c = (1 - a_c)\psi_c$ and $\bar{\theta}_c = 1 - \psi_{c+1}$

Consider now $c = C$, it holds that

$$\bar{\theta}_C = (1 - a_C)\psi_C = (1 - a_C)(1 - \bar{\theta}_{C-1}) \leq (1 - a_C)(1 - \bar{\theta}_C),$$

which leads to

$$\bar{\theta}_C \le \frac{1-a_C}{2-a_C}, \qquad a_C \le \frac{1-2\bar{\theta}_C}{1-\bar{\theta}_C}.$$

Since $1-a_C = \overline{\nu}_C a_C$, i.e., $a_C = 1/(1+\overline{\nu}_C)$, we get

$$\bar{\theta}_C \le \frac{\overline{\nu}_C}{1+2\overline{\nu}_C} \le \frac{\nu_C}{1+2\nu_C}. \tag{26}$$

Consider now any $c = 2, \dots, C-1$. From [21, Eqn. (34)] we have

$$1 - a_c = \overline{\nu}_c a_c a_{c+1}, \tag{27}$$

and hence,

$$\bar{\theta}_c = (1-a_c)\psi_c = (1-a_c)(1-\bar{\theta}_{c-1}) \le (1-a_c)(1-\bar{\theta}_c),$$

which leads to

$$\bar{\theta}_c \le \frac{1-a_c}{2-a_c}, \qquad a_c \le \frac{1-2\bar{\theta}_c}{1-\bar{\theta}_c}. \tag{28}$$

From (27), we have $a_c = (1+\overline{\nu}_c a_{c+1})^{-1}$, which we can plug into the inequality for the throughput to have

$$\begin{aligned}\bar{\theta}_{c+1} \le \bar{\theta}_c \le \frac{1-a_c}{2-a_c} &= \frac{\overline{\nu}_c a_{c+1}}{1+2\overline{\nu}_c a_{c+1}} \\ &\le \frac{(1-2\bar{\theta}_{c+1})\overline{\nu}_c}{1-\bar{\theta}_{c+1}+2(1-2\bar{\theta}_{c+1})\overline{\nu}_c}.\end{aligned}$$

This leads to a quadratic inequality implying

$$\bar{\theta}_C \le \theta_{c+1} \le \frac{1-\sqrt{\frac{1}{4\overline{\nu}_c+1}}}{2} \le \frac{1}{2} - \frac{1}{2\sqrt{1+4\nu_c}}. \tag{29}$$

It remains to deal with the case $c = 1$. For that, [21, Eqn. (33)] gives us

$$\bar{\theta}_2 \le \bar{\theta}_1 = \frac{\overline{\nu}_1 a_2}{1+\overline{\nu}_1 a_2} \le \frac{\overline{\nu}_1(1-2\bar{\theta}_2)}{1-\bar{\theta}_2+\overline{\nu}_1(1-2\bar{\theta}_2)},$$

from (28). The last inequality gives us $\bar{\theta}_C \le \bar{\theta}_2 \le \frac{\overline{\nu}_1}{1+2\overline{\nu}_1} \le \frac{\nu_1}{1+2\nu_1}$.

Combining the last inequality with (26) and (29), we see $\bar{\theta}_C(\bar{\boldsymbol{\nu}}) \le \lambda^*(\boldsymbol{\nu})$ for every $\bar{\boldsymbol{\nu}} \le \boldsymbol{\nu}$. □

A.3 Proof of Theorem 4

Proof. In the proof we will use the characterization in (20)-(22) with $\nu_c = \nu$ for $c = 1, \dots, C$.

We first show that the configurations different from $(\mathtt{U}, \dots, \mathtt{U})$, $(\mathtt{U}, \mathtt{S}, \mathtt{U} \dots, \mathtt{U})$, and $(\mathtt{S}, \mathtt{S}, \mathtt{U} \dots, \mathtt{U})$ do not admit a feasible solution.

- $\boldsymbol{s}$ such that $s_C = \mathtt{S}$. Consider the ratio between (22) and (20). It holds that

$$\begin{aligned}\frac{1}{\rho_1^-} &= \frac{\lambda\prod_1^C \rho_d^+}{(1-\lambda\prod_1^{C-1}\rho_d^+(1+\rho_C^+))}\frac{(1-\lambda\rho_1^+(1+\rho_2^+))}{\lambda\rho_1^+} \\ &= \frac{(1-\lambda\rho_1^+(1+\rho_2^+))\prod_2^C\rho_d^+}{(1-\lambda\prod_1^{C-1}\rho_d^+(1+\rho_C^+))},\end{aligned}$$

which yields a contradiction since

$$\begin{aligned}\prod_2^C \rho_d^+ &= \frac{1-\lambda\prod_1^{C-1}\rho_d^+(1+\rho_C^+)}{\rho_1^-(1-\lambda\rho_1^+(1+\rho_2^+))} \\ &\ge \frac{1-\lambda\prod_1^{C-1}\rho_d^+(1+\rho_C^+)}{1-\lambda\rho_1^+(1+\rho_2^+)} \ge 1.\end{aligned}$$

- $\boldsymbol{s}$ such that $s_c = \mathtt{S}$ for $c = 3, \dots, C-1$. Consider the ratio between the equations in (21) for the c-th and the 2nd class. Simple algebra yields

$$\frac{1}{\rho_2^-} = \frac{\prod_3^c \rho_d^+(1-L_c)(1-L_1-L_2)(1-L_2-L_3)}{(1-L_2)(1-L_{c-1}-L_c)(1-L_c-L_{c+1})},$$

where $L_{c'} = \lambda\prod_1^{c'}\rho_d^+$, and observe that $L_d \ge L_{d+1}$ for every $d = 1, \dots, C-1$. Hence,

$$\begin{aligned}\prod_3^c \rho_d^+ &= \frac{1}{\rho_2^-}\frac{(1-L_2)(1-L_{c-1}-L_c)(1-L_c-L_{c+1})}{(1-L_c)(1-L_1-L_2)(1-L_2-L_3)} \\ &\ge \underbrace{\frac{(1-L_{c-1}-L_c)}{(1-L_2-L_3)}}_{C_1}\underbrace{\frac{(1-L_2)(1-L_c-L_{c+1})}{(1-L_c)(1-L_1-L_2)}}_{C_2}.\end{aligned}$$

Observe now that $C_1 \ge 1$ iff $L_2 + L_3 \ge L_{c-1} + L_c$, which holds since $c \ge 3$. Similarly $C_2 \ge 1$ iff $L_1(1-L_c) \ge L_{c+1}(1-L_2)$ which holds since $L_1 \ge L_{c+1}$ and $L_2 \ge L_c$. Therefore, this yields a contradiction since $\prod_3^c \rho_d^+ \le \rho_c^+ < 1$.

- $\boldsymbol{s}$ such that $s_1 = \mathtt{S}$ and $s_2 = \mathtt{U}$. Consider the ratio between the equations in (20) and (21) for $c = 2$. It holds that

$$\frac{1}{\rho_2^-} = \frac{(1-\lambda\rho_1^+\rho_2^+(1+\rho_3^+)}{\rho_2^+(1-\lambda\rho_1^+\rho_2^+)},$$

which yields a contradiction since

$$\rho_2^+ = \rho_2^-\frac{(1-\lambda\rho_1^+\rho_2^+(1+\rho_3^+)}{(1-\lambda\rho_1^+\rho_2^+)} \le \frac{1-\lambda\rho_1^+\rho_2^+-\lambda\rho_1^+\rho_2^+\rho_3^+}{1-\lambda\rho_1^+\rho_2^+} < 1.$$

At this point, we prove that for each of the feasible states there exists a solution $\boldsymbol{\rho}$ which corresponds to the one given in the theorem.

- $\boldsymbol{s} = (\mathtt{U}, \dots, \mathtt{U})$: Immediately from Lemma 1.
- $\boldsymbol{s} = (\mathtt{U}, \mathtt{S}, \mathtt{U}, \dots, \mathtt{U})$: From Equation (21) for $c = 2$, we obtain that

$$\nu = \frac{\lambda\rho_2^+(1-\lambda\rho_2^+)}{(1-\lambda-\lambda\rho_2^+)(1-2\lambda\rho_2^+)},$$

which yields

$$\rho_2^+ = \frac{(3-2\lambda)\nu+1 \pm \sqrt{(1-2\lambda)^2\nu^2+2\nu+1}}{2(1+2\lambda\nu)} = c_\pm(\lambda,\nu).$$

It holds that $c_+(\lambda,\nu) \ge c_-(\lambda,\nu)$ for every $\lambda, \nu \ge 0$, that both are strictly decreasing in λ, and that

$$\rho_2^+ = c_+(\lambda,\nu) = 1 \iff \lambda = \frac{1}{2} + \frac{1}{2}\sqrt{\frac{1}{1+4\nu}},$$

$$\rho_2^+ = c_-(\lambda,\nu) = 1 \iff \lambda = \frac{1}{2} - \frac{1}{2}\sqrt{\frac{1}{1+4\nu}}.$$

From Equations (20),(22), and (21) for $c \ne 2$, we obtain

$$\begin{aligned}\rho_1^- &= \frac{\lambda}{\nu(1-\lambda-\lambda\rho_2^+)} = \frac{\lambda}{\nu(1-\lambda-\lambda c_\pm(\lambda,\nu))}, \\ \rho_c^- &= \frac{\lambda\rho_2^+(1-\lambda\rho_2^+)}{\nu(1-2\lambda\rho_2^+)^2} = \frac{\lambda\rho_2^+(1-\lambda\rho_2^+)(1-\lambda-\lambda\rho_2^+)}{\nu(1-2\lambda\rho_2^+)^2(1-\lambda-\lambda\rho_2^+)} \\ &= \frac{1-\lambda-\lambda c_\pm(\lambda,\nu)}{1-2\lambda c_\pm(\lambda,\nu)}, \\ \rho_C^- &= \frac{\lambda\rho_2^+}{\nu(1-2\lambda\rho_2^+)} = \frac{\lambda c(\lambda,\nu)}{\nu(1-2\lambda c_\pm(\lambda,\nu))}.\end{aligned}$$

Consider $c_\pm(\lambda,\nu) = c_+(\lambda,\nu)$, in order to have $\rho_2^+ \leq 1$ it requires that $\lambda > 1/2$, however

$$\rho_1^- = \frac{2\lambda(1+2\nu)}{\nu((1-2\lambda)(1+\nu) - \sqrt{1+2\nu+(1-2\lambda)^2\nu^2})} < 0.$$

Hence, the only feasible solution is given by $c_\pm(\lambda,\nu) = c_-(\lambda,\nu)$.

• $\boldsymbol{s} = (\mathtt{S},\mathtt{S},\mathtt{U},\ldots,\mathtt{U})$: From Equations (20) and (21) for $c = 2$, we obtain that

$$\nu = \frac{\lambda\rho_1^+}{1-\lambda\rho_1^+-\lambda\rho_1^+\rho_2^+} = \frac{\lambda\rho_1^+\rho_2^+(1-\lambda\rho_1^+\rho_2^+)}{(1-\lambda\rho_1^+-\lambda\rho_1^+\rho_2^+)(1-2\lambda\rho_1^+\rho_2^+)},$$

which yields

$$\rho_2^+ = \frac{1+\nu}{1+2\nu}, \qquad \rho_1^+ = \frac{\nu(1+2\nu)}{\lambda(1+4\nu+3\nu^2)} = \frac{\nu(1+2\nu)}{\lambda(1+\nu)(1+3\nu)}.$$

Hence, from (22) we have that

$$\rho_C^- = \frac{\lambda\rho_1^+\rho_2^+}{\nu(1-2\lambda\rho_1^+\rho_2^+)} = \frac{\frac{\nu}{1+3\nu}}{\nu(1-2\frac{\nu}{1+3\nu})} = \frac{1}{1+\nu},$$

while from (21) for $c = 3,\ldots,C-1$, we obtain

$$\rho_c^- = \frac{\lambda\rho_1^+\rho_2^+(1-\lambda\rho_1^+\rho_2^+)}{\nu(1-2\lambda\rho_1^+\rho_2^+)^2} = \frac{\frac{\nu}{1+3\nu}(1-\frac{\nu}{1+3\nu})}{\nu(1-2\frac{\nu}{1+3\nu})^2} = \frac{1+2\nu}{(1+\nu)^2}.$$

The last step of the proof consists in showing that the value of λ uniquely identifies the load configuration, and hence the solution $\boldsymbol{\rho} \in R(\lambda,\boldsymbol{\nu}^u)$. Consider the unsaturated configuration, we require $\lambda < 1/2$ otherwise $\rho_1^- < 0$. Hence, consider $c \in 2,\ldots,C-1$ and $c' \in \{1,C\}$,

$$\rho_c^- = \frac{1}{\nu}\frac{\lambda(1-\lambda)}{(1-2\lambda)^2} = \rho_{c'}^-\frac{1-\lambda}{1-2\lambda} \geq \rho_{c'}^-, \qquad \forall\lambda < \frac{1}{2}.$$

A solution in load configuration $(\mathtt{U},\ldots,\mathtt{U})$ exists if and only if $\lambda \leq \lambda^{(1)}$ the load at which $\rho_c^- = 1$ for $c \in 2,\ldots,C-1$.

Consider now the load configuration $(\mathtt{S},\mathtt{S},\mathtt{U}\ldots,\mathtt{U})$, it holds that $\rho_1^+ < 1$ iff $\lambda > \lambda^{(2)}$. Moreover, for every $\lambda > \lambda^{(2)}$ it holds that $\rho_2^+ < 1$ and $\rho_c^- \leq 1$ for every $c = 3,\ldots,C$.

Consider now the load configuration $(\mathtt{U},\mathtt{S},\mathtt{U}\ldots,\mathtt{U})$, it is sufficient to show that both for $\lambda \leq \lambda^{(1)}$ and for $\lambda > \lambda^{(2)}$, the solution is unfeasible. In fact, since a solution needs to exists even for $\lambda \in (\lambda^{(1)},\lambda^{(2)}]$ it necessarily is in load configuration $(\mathtt{U},\mathtt{S},\mathtt{U}\ldots,\mathtt{U})$. We can conclude since $\rho_2^+ < 1$ iff $\lambda > \lambda^{(1)}$ and $\rho_1^- \leq 1$ iff $\lambda \leq \lambda^{(2)}$. □

A.4 Proof of Theorem 6

Before showing this we need the following auxiliary result, showing that each possible oversaturated load configuration has a unique solution, whose admissibility depends on the subset S_j such that $\boldsymbol{\nu} \in S_j$.

Proposition 6. *For every load configuration* $\boldsymbol{s} \in \mathtt{S} \times \{\mathtt{U},\mathtt{S}\}^2$ *there exists a unique solution* $\boldsymbol{\rho}$ *for* (19). *If* $\boldsymbol{s} = (\mathtt{S},\mathtt{U},\mathtt{U})$ *it holds that*

$$\rho_1^+ = \frac{\nu_1}{\lambda(2\nu_1+1)}, \qquad \rho_2^- = \frac{\nu_1+\nu_1^2}{\nu_2}, \qquad \rho_3^- = \frac{\nu_1}{\nu_3},$$

if $\boldsymbol{s} = (\mathtt{S},\mathtt{S},\mathtt{U})$ *it holds that*

$$\rho_1^+ = \frac{\nu_1+\nu_1^2+\nu_1\nu_2}{\lambda(1+\nu_1+2\nu_2)(1+\nu_1)},$$
$$\rho_2^+ = \frac{\nu_2(1+\nu_1)}{\nu_1(1+\nu_1+\nu_2)}, \qquad \rho_3^- = \frac{\nu_2}{\nu_3(1+\nu_1)},$$

if $\boldsymbol{s} = (\mathtt{S},\mathtt{U},\mathtt{S})$ *it holds that*

$$\rho_1^+ = \frac{\nu_1}{\lambda(1+2\nu_1)}, \quad \rho_2^- = \frac{\nu_1(1+\nu_3)}{\nu_2}, \quad \rho_3^+ = \frac{\nu_3(1+\nu_1)}{\nu_1(1+\nu_3)}.$$

if $\boldsymbol{s} = (\mathtt{S},\mathtt{S},\mathtt{S})$ *it holds that*

$$\rho_1^+ = \frac{\nu_1(1+\nu_3)}{\lambda(1+\nu_1+\nu_2+\nu_3+\nu_1\nu_3)},$$
$$\rho_2^+ = \frac{\nu_2}{\nu_1(1+\nu_3)}, \qquad \rho_3^+ = \frac{\nu_3(1+\nu_1)}{\nu_2}.$$

Proof. Consider Equations (20)-(22) with $C = 3$ and set $\rho_1^- = 1$. We separately look at the four oversaturation possibilities.

• $\boldsymbol{s} = (\mathtt{S},\mathtt{U},\mathtt{U})$, i.e., $\rho_2^+,\rho_3^+ = 1$,

$$(20) = (22) \;\Rightarrow\; \nu_1 = \nu_3\rho_3^- \;\Rightarrow\; \rho_3^- = \frac{\nu_1}{\nu_3}.$$
$$(20) = (21) \;\Rightarrow\; \rho_2^- = \frac{\nu_1+\nu_1^2}{\nu_2}.$$
$$(21) \;\Rightarrow\; \lambda\rho_1^+ = \frac{\nu_1}{2\nu_1+1}.$$

• $\boldsymbol{s} = (\mathtt{S},\mathtt{S},\mathtt{U})$, i.e., $\rho_2^-,\rho_3^+ = 1$,

$$(21) = (22) \;\Rightarrow\; \nu_2 = \nu_3\rho_3^- + \nu_1\nu_3\rho_3^- \;\Rightarrow\; \rho_3^- = \frac{\nu_2}{\nu_3(1+\nu_1)}.$$
$$(21)/(20) \;\Rightarrow\; \rho_2^+ = \frac{\nu_2}{\nu_1+\nu_1\nu_3\rho_3^-} = \frac{\nu_2(1+\nu_1)}{\nu_1+\nu_1^2+\nu_1\nu_2}.$$
$$(21) \;\Rightarrow\; \rho_1^+ = \frac{\nu_1+\nu_1^2+\nu_1\nu_2}{\lambda(1+\nu_1+2\nu_2)(1+\nu_1)}.$$

• $\boldsymbol{s} = (\mathtt{S},\mathtt{U},\mathtt{S})$, i.e., $\rho_2^+,\rho_3^- = 1$,

$$(20) = (21) \;\Rightarrow\; \nu_1+\nu_1\nu_3 = \nu_2\rho_2^- \;\Rightarrow\; \rho_2^- = \frac{\nu_1(1+\nu_3)}{\nu_2}.$$
$$(22)/(20) \;\Rightarrow\; \rho_3^+ = \frac{\nu_3(1+\nu_1)}{\nu_1(1+\nu_3)}.$$
$$(20) \;\Rightarrow\; \lambda\rho_1^+ = \frac{\nu_1}{1+2\nu_1}.$$

• $\boldsymbol{s} = (\mathtt{S},\mathtt{S},\mathtt{S})$, i.e., $\rho_2^-,\rho_3^- = 1$,

$$(22)/(21) \;\Rightarrow\; \rho_3^+ = \frac{\nu_3+\nu_1\nu_3}{\nu_2}.$$
$$(21)/(20) \;\Rightarrow\; \rho_2^+ = \frac{\nu_2}{\nu_1+\nu_1\nu_3}.$$
$$(20) \;\Rightarrow\; \rho_1^+ = \frac{\nu_1+\nu_1\nu_3}{\lambda(1+\nu_1+\nu_2+\nu_3+\nu_1\nu_3)}.$$

□

We now present the proof of Theorem 6

Proof. It is a simple exercise to show that the sets S_1, S_2, S_3 and S_4 form a partition of $\mathbb{R}_+^3$. Now, observe that these solutions are feasible if and only if $\rho_c^+ < 1$ and $\rho_c^- \leq 1$ for $c = 1,2,3$. However, since $\lambda \gg 0$, it holds that $\rho_1^+ < 1$. Hence, so as for the solution with $\boldsymbol{s} = (\mathtt{S},\mathtt{U},\mathtt{U})$ to be feasible, we require

$$\frac{\nu_1+\nu_1^2}{\nu_2} \leq 1, \;\; \frac{\nu_1}{\nu_3} \leq 1,$$

which corresponds to $\boldsymbol{\nu} \in S_1$. Similarly, we can show that the solutions with $\boldsymbol{s} = (\mathtt{S},\mathtt{S},\mathtt{U})$, $\boldsymbol{s} = (\mathtt{S},\mathtt{U},\mathtt{S})$, and $\boldsymbol{s} = (\mathtt{S},\mathtt{S},\mathtt{S})$, are feasible if and only if $\boldsymbol{\nu} \in S_2$, $\boldsymbol{\nu} \in S_3$, and $\boldsymbol{\nu} \in S_4$ respectively. □