EURANDOM PREPRINT SERIES
2011-019

Insensitivity and Stability of Random-Access Networks

P.M. van de Ven, J.S.H. van Leeuwaarden, S.C. Borst, A. Proutiére
ISSN 1389-2355

Insensitivity and Stability of Random-Access Networks

P.M. van de Ven[1,2] S.C. Borst[1,3]
J.S.H. van Leeuwaarden[1,2] A. Proutière[4]

Abstract

Random-access algorithms such as the Carrier-Sense Multiple-Access (CSMA) protocol provide a popular mechanism for distributed medium access control in large-scale wireless networks. In recent years, fairly tractable models have been shown to yield remarkably accurate throughput estimates for CSMA networks. These models typically assume that both the transmission durations and the back-off periods are exponentially distributed. We show that the stationary distribution of the system is in fact insensitive with respect to the transmission durations and the back-off times.

These models primarily pertain to a saturated scenario where nodes always have packets to transmit. In reality however, the buffers may occasionally be empty as packets are randomly generated and transmitted over time. The resulting interplay between the activity states and the buffer contents gives rise to quite complicated queueing dynamics, and even establishing the stability criteria is usually a serious challenge. We explicitly identify the stability conditions in a few relevant scenarios, and illustrate the difficulties arising in other cases.

1 Introduction

Emerging wireless mesh networks typically lack any centralized control entity for regulating access and coordinating transmissions. Instead, these networks vitally rely on the individual nodes to operate autonomously and to efficiently share the medium in a distributed fashion. This requires the nodes to schedule their individual transmissions and decide on the use of a shared medium based on knowledge that is locally available or only involves limited exchange of information. A popular mechanism for distributed medium access control is provided by the Carrier-Sense Multiple-Access (CSMA) protocol, various incarnations of which are implemented in IEEE 802.11 networks. In the CSMA

[1]Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands
[2]Eurandom, P.O. Box 513, 5600 MB Eindhoven, The Netherlands
[3]Alcatel-Lucent Bell Labs, P.O. Box 636, Murray Hill, NJ 07974-0636, USA
[4]Microsoft Research, 7JJ Thomson Avenue, Cambridge, CB30FB, UK

protocol each node attempts to access the medium after a certain back-off time, but nodes that sense activity of interfering nodes freeze their back-off timer until the medium is sensed idle.

While the CSMA protocol is fairly easy to understand at a local level, the interaction among interfering nodes gives rise to quite intricate behavior and complex throughput characteristics on a macroscopic scale. In recent years relatively parsimonious models have emerged that provide a useful tool in evaluating the throughput characteristics of CSMA-like networks. These models essentially assume that the interference constraints can be represented by a general conflict graph, and that the various nodes activate asynchronously at exponential rates whenever none of their neighbors are active. Such models were first pursued in the context of IEEE 802.11 systems by Wang & Kar [31], and further studied in that setting in [8, 9, 10], with several extensions and refinements in [7, 12, 24]. These models in fact long pre-date the IEEE 802.11 standard and were already considered in the 1980's [4, 5, 16, 19]. The model has strong connections with Markov random fields and migration processes, and can under certain assumptions be interpreted as a special instance of a loss network [15, 17, 18, 28, 32].

Although the representation of the IEEE 802.11 back-off mechanism in the above-mentioned models is far less detailed than in the landmark work of Bianchi [2], the general interference graph offers greater versatility and covers a broad range of topologies. Experimental results in [21] demonstrate that these models, while idealized, provide throughput estimates that match remarkably well with measurements in actual IEEE 802.11 systems.

The above-mentioned CSMA models typically assume that both the transmission durations and the back-off periods are exponentially distributed, in which case the model reduces to a loss network. As the relationship with loss networks suggests, it is easily verified that the stationary distribution of the joint activity process is in fact insensitive with respect to the distribution of the transmission times, i.e., holds for any distribution with the same mean. In contrast, loss networks are generally not insensitive with respect to the interarrival time distribution. So the impact of the distribution of the back-off periods is far less obvious, and it might seem that no insensitivity property should be expected to hold. However, closer inspection reveals that the strict equivalence with loss networks relies on the back-off periods being exponentially distributed. In order to see that, observe that in loss networks the arrival process is not affected by the occupancy state, whereas in the CSMA model the back-off process of a node is suspended when that node is active, and is possibly frozen by the activity of neighboring nodes. In case the back-off periods are exponentially distributed, back-off freezing does not affect the activity process, so the CSMA model is equivalent to a loss network. For generally distributed back-off periods this distinction does become relevant, and no direct analogy with loss networks applies. As we will show in the present paper, this difference plays a crucial role, and it turns out that the stationary distribution of the joint activity process in the CSMA network is insensitive with respect to the distribution of the back-off periods as well. This holds for any interference graph, and irrespective of whether or not the back-off process of a node is frozen by the activity of its neighbors.

The above-mentioned references further primarily focus on a saturated scenario where nodes always have packets pending for transmission. The throughput characteristics in such scenarios provide useful first-order estimates of the system performance. In reality however, buffer contents fluctuate as packets are accumulated and flushed over time, giving rise to queueing dynamics. In particular, the buffers may empty from time to time, and nodes will refrain from competition for the medium during these periods. The resulting interaction between the activity states and the buffer contents of the various nodes gives rise to quite intricate behavior. In particular, the queueing dynamics entail high-dimensional stochastic processes with infinite state spaces, which generally do not admit closed-form expressions for the stationary distribution. Even just establishing the existence of a stationary distribution, i.e., obtaining the stability conditions, is generally a challenging problem, and often is about as hard as determining the entire joint distribution of the buffer contents. In the present paper, we will identify necessary and sufficient conditions for stability in a complete interference graph, and illustrate the difficulties that arise for general interference graphs.

The above discussion pertains to the case where the mean back-off times and the activation rates are assumed to be fixed. When the activation rates are allowed to be adapted, there are simple necessary and sufficient conditions for stability to be achievable. Several authors have proposed clever backlog-based algorithms for adapting activation rates that achieve stability whenever feasible to do so at all [13, 14, 22, 23, 26, 27]. The case of fixed activation rates is nevertheless relevant, since in practice the adaptation of back-off parameters involves a wide range of non-trivial implementation issues (finite-range precision, communication overhead, information exchange), and hence it is important to gain insight in the achievable performance of non-adaptive algorithms.

The paper is organized as follows. In Section 2 we present a detailed model description and establish the insensitivity result for the stationary distribution of the model under saturated conditions. It is proved that the stationary distribution is insensitive with respect to both the transmission durations and back-off periods, regardless of whether or not the back-off process of a node is frozen while its neighbors are active. We then turn to a non-saturated scenario, and derive stability conditions for several cases of interest. Section 3 presents a necessary stability condition for general interference graphs. In Section 4 the stability region is characterized for full interference graphs, while Section 5 illustrates the difficulties that arise for partial interference graphs. In Section 6 we make some concluding remarks and identify various topics for future research.

2 Insensitivity of the saturated model

We consider a network of several nodes sharing a wireless medium according to a CSMA-type protocol. The network is described by an undirected graph

(V, E) where the set of vertices $V = \{1, \ldots, N\}$ represents the various nodes of the network and the set of edges $E \subseteq V \times V$ indicates which pairs of nodes interfere. In other words, nodes that are neighbors in the interference graph are prevented from simultaneous activity, and thus the independent sets correspond to the feasible joint activity states. A node is said to be blocked whenever the node itself or any of its neighbors is active, and unblocked otherwise.

In this section we consider a scenario where nodes are saturated, i.e., always have packets to transmit. The transmission times of node i are independent and generally distributed with mean $1/\mu_i$. After each transmission, a node starts a back-off period. The back-off periods of node i are independent and generally distributed with mean $1/\nu_i$. For compactness, denote $\sigma_i = \nu_i/\mu_i$.

When a node is in back-off, it becomes blocked when any of its neighbors activate. We distinguish two scenarios, depending on whether or not the back-off period is frozen in that case. In case the back-off period is frozen, it is resumed as soon as the node becomes unblocked again. When the back-off period of a node ends, it is unblocked and will start a transmission. In case the back-off does not get frozen, the back-off period may end while the node is blocked, in which case the node simply starts a new back-off period. Under the distributional assumptions below, the probability of the back-off periods of two nodes ending simultaneously is equal to 0. Note that in case the back-off periods are exponentially distributed, it does not matter whether or not the back-off periods are frozen when a node becomes blocked, nor does it matter whether they are resumed or resampled when a node becomes unblocked again. Equivalently, we could then think of the potential activation epochs of a node as occurring according to a Poisson process, and actual transmission periods starting whenever a potential activation event occurs while the node is unblocked.

Define $\Omega \subseteq \{0, 1\}^N$ as the set of all feasible joint activity states of the network, i.e., the incidence vectors of the independent sets of the interference graph. Let $X(t) \in \Omega$ represent the activity state of the network at time t, with $X_i(t)$ indicating whether node i is active at time t or not. Denote by $\pi(x) = \lim_{t\to\infty} P\{X(t) = x\}$ the limiting probability that the joint activity state is $x \in \Omega$.

We assume all back-off times and transmission durations to have phase-type distributions, and are interested in the limiting distribution of the Markov process that keeps track of the phase of each node. Using this limiting distribution, and the fact that phase-type distribution are dense in the space of all probability distributions with positive support, we will show that π is insensitive to the distributions of the back-off times and the transmission durations, and only depends on these distributions through the σ_i.

Let the back-off process of node i have a phase-type distribution with $m_i + 1$ phases, where states $1, \ldots, m_i$ are transient, and state $m_i + 1$ is absorbing. The corresponding starting probabilities are $\alpha_1, \ldots, \alpha_{m_i+1}$, and the transition rates are given by q_{kl}. Similarly, the transmission times of node i have a phase-type distribution with n_i+1 phases (with state n_i+1 absorbing), starting probabilities $\gamma_1, \ldots, \gamma_{n_i+1}$ and transition rates r_{kl}. Note that α_{m_i+1} (γ_{n_i+1}) represents the probability of a back-off period (transmission) of zero length.

Let β_k, $k = 1, \dots, m_i$, represent the fraction of time that the back-off process of a node is in phase k, and let η_k, $k = 1, \dots, n_i$, represent the fraction of time that the transmission process of a node is in phase k. The fractions β_k and η_k follow from the equations (8) and (9) in Section 7. We study the Markov process that keeps track of the activity of all nodes. Let $\omega = (\omega_1, \dots, \omega_N)$ denote the state of the system. We use the convention that $\omega_i = -k$ when node i is in back-off phase k, and $\omega_i = k$ when node i is in the k-th transmission phase. So ω_i assumes values in $\Omega_i = \{-m_i, \dots, -1\} \cup \{1, \dots, n_i\}$, and the state space Ω_{PH} of the Markov process of the joint activity state satisfies $\Omega_{\mathrm{PH}} \subseteq \Omega_1 \times \cdots \times \Omega_N$. The limiting distribution π_{PH} of this Markov process has the following product-form solution:

Lemma 1. *Let the back-off times and transmission durations have a phase-type distribution. Then, regardless of back-off freezing,*

$$\pi_{\mathrm{PH}}(\omega) = Z^{-1} \prod_{i:\omega_i \leq -1} \beta_{-\omega_i} \prod_{i:\omega_i \geq 1} \sigma_i \eta_{\omega_i}, \quad \omega \in \Omega_{\mathrm{PH}}, \tag{1}$$

where Z is the normalization constant.

The proof of Lemma 1 is presented in Section 7.2. In this proof we treat the freezing and non-freezing systems in parallel, and we demonstrate that freezing indeed has no impact on the stationary distribution.

Using Lemma 1 we can now show that the node activity is insensitive to the distributions of the back-off times and transmission durations.

Theorem 1. *Let the back-off times and transmission durations have a phase-type distribution. Then, regardless of back-off freezing,*

$$\pi(x) = Z^{-1} \prod_{i=1}^{N} \sigma_i^{x_i}, \quad x \in \Omega, \tag{2}$$

where Z is the normalization constant.

Proof. Denote by $\Omega_{\mathrm{PH}}(x)$ the set of all states $\omega \in \Omega_{\mathrm{PH}}$ that correspond to $x \in \Omega$, i.e.,

$$\Omega_{\mathrm{PH}}(x) = \{\omega \in \Omega_{\mathrm{PH}} \mid \forall i : \omega_i \leq -1 \text{ if } x_i = 0,\ \omega_i \geq 1 \text{ if } x_i = 1\}.$$

Then

$$\pi(x) = \sum_{\omega \in \Omega_{\mathrm{PH}}(x)} \pi_{\mathrm{PH}}(\omega) = Z^{-1} \sum_{\omega \in \Omega_{\mathrm{PH}}(x)} \prod_{i:\omega_i \leq -1} \beta_{\omega_i} \prod_{i:\omega_i \geq 1} \sigma_i \eta_{\omega_i} = Z^{-1} \prod_{i=1}^{N} \sigma_i^{x_i},$$

as $\sum_k \beta_k = 1$ and $\sum_k \eta_k = 1$. □

The result in Theorem 1 is well known for exponentially distributed back-off periods, and follows in that case from the connection with loss networks.

For generally distributed back-off periods and back-off freezing, partial proof arguments are presented in [21]. In the case without back-off freezing, the insensitivity result may be directly proven by representing the dynamics as those observed in an Engset network as considered in [3], Section 5.3. The Engset network is constructed from the interference graph (V, E): the links in the Engset network are of unit capacity and correspond to the undirected edges in the graph. Each node $i \in V$ is then represented as a customer in the Engset network which, when active, simultaneously uses links $\{i, j\}$ for all j such that $\{i, j\} \in E$. Each customer alternates between active and inactive phases. Finally, without back-off freezing, a customer who wishes to become active, is blocked if one of the required links is occupied (i.e., the corresponding node has an active neighbor), in which case the customer starts a new inactivity period (back-off). This corresponds to the so-called *jump-over* retrial behavior considered in [3], and insensitivity follows. In the case of back-off freezing, it does not seem possible to apply such a representation, and our proof of insensitivity is new.

3 A necessary stability condition

In the previous section, we considered a scenario with saturated nodes that always have packets to transmit. We now turn to a scenario where packets are generated over time and buffers may occasionally be empty. Specifically, packets arrive at node i according to a renewal process with mean interarrival time $1/\lambda_i$. Nodes compete for access to the medium as before, with the modification that when unblocked nodes have no packets to transmit at the time a back-off period ends, they simply start a new back-off period. Once a packet has been transmitted, it leaves the system. The results in this section are valid irrespective of whether or not the back-off process is frozen.

Denote by $\rho_i = \lambda_i/\mu_i$ the traffic intensity at node i, so ρ_i is the fraction of time that this node has to be active in order to sustain the arrival rate λ_i. Define θ_i as the throughput of node i, i.e., the expected number of transmissions per unit of time, and denote by τ_i the fraction of time that node i is active, so $\theta_i = \mu_i \tau_i$. Denote by $\tau_i^* = \sum_{x \in \Omega : x_i = 1} \pi(x)$ and $\theta_i^* = \mu_i \tau_i^*$ the fraction of time node i is active and throughput of node i, respectively, in the regime where all nodes are saturated. We have by definition $\theta_i \leq \lambda_i$, with equality when node i is stable.

The next proposition provides a simple necessary condition for stability.

Proposition 1. *If $\lambda_i > \theta_i^*$ for all $i = 1, \ldots, N$, then all the nodes are unstable.*

Proof. We will show that all the nodes are unstable in the sense that each of the associated queues only empties finitely often, but it can in fact be established that the queue of node i grows in a linear manner at rate $\lambda_i - \theta_i^*$ in the long run. For convenience, we restrict the proof to a Poisson arrival process, but the arguments extend to general renewal arrival processes.

The key observation is that once a queue empties, with non-zero probability the system may enter a state with all queues non-empty. Since all queues have positive drift in saturated conditions, all queues remain non-empty with non-zero probability. Thus, every time a queue empties, it may never do so again with probability bounded away from zero, and hence the queue will only empty finitely often.

In order to formalize the above observation, let $Q_j(t)$ be the number of packets pending for transmission or in the process of being transmitted at node j at time t. Let $T_{i,n}$ be the time that the queue of node i empties for the n-th time. Let $U_{i,n} := \inf\{t \geq T_{i,n} : Q_j(t) \geq 1 \text{ for all } j = 1, \ldots, N\}$ be the first time after $T_{i,n}$ when all queues are non-empty. It is easily verified that there exists $b_1 > 0$ such that $\mathbb{P}\{U_{i,n} < T_{i,n+1}\} > b_1$ for all n.

Let $A_j(s,t)$ be the number of packet arrivals at node j during the time interval $[s,t]$ and $B_j(s,t)$ the number of packet transmissions at node j during the time interval $[s,t]$, so

$$Q_j(t) = Q_j(s) + A_j(s,t) - B_j(s,t). \tag{3}$$

Moreover, denote by $B_j^*(U_{i,n}, t)$ the number of packet transmissions at node j during the time interval $[U_{i,n}, t]$ in a modified version of the network where the various nodes are in the exact same state at time $U_{i,n}$ and are all assumed to be saturated from that time onward. Define

$$V_{i,j,n} := \sup\{t : Q_j(s) \geq 1 \ \forall s \in [U_{i,n}, t]\},$$
$$W_{i,j,n} := \sup\{t : A_j(U_{i,n}, s) - B_j^*(U_{i,n}, s) \geq 0 \ \forall s \in [U_{i,n}, t]\},$$

and denote $V_{i,n} := \min_{j=1,\ldots,N} V_{i,j,n}$ and $W_{i,n} := \min_{j=1,\ldots,N} W_{i,j,n}$.

By definition, $Q_j(t) \geq 1$ for all $t \in [U_{i,n}, V_{i,n}]$, $j = 1, \ldots, N$. Thus for all nodes j we have $B_j(U_{i,n}, t) = B_j^*(U_{i,n}, t)$ for all $t \in [U_{i,n}, V_{i,n}]$. From (3) with $s = U_{i,n}$ we see

$$Q_j(t) \geq 1 + A_j(U_{i,n}, t) - B_j(U_{i,n}, t) = 1 + A_j(U_{i,n}, t) - B_j^*(U_{i,n}, t), \quad t \in [U_{i,n}, V_{i,n}],$$

so $V_{i,n} \geq W_{i,n}$. Since $\frac{1}{t}\mathbb{E}\{A_j(U_{i,n}, U_{i,n} + t)\} = \lambda_j > \theta_j^* = \lim_{t\to\infty} \frac{1}{t}\mathbb{E}\left\{B_j^*(U_{i,n}, U_{i,n} + t)\right\}$, it follows that there exists $b_2 > 0$ such that $\mathbb{P}\{V_{i,n} = \infty\} \geq \mathbb{P}\{W_{i,n} = \infty\} > b_2$ for all n.

In conclusion, the probability that the queue of node i never empties again after it has emptied for the n-th time, is bounded from below by $b = b_1 b_2 > 0$. Thus, the total expected number of times that the queue of node i empties is bounded from above by $\sum_{n=0}^{\infty}(1-b)^n = 1/b$, which means that it only empties finitely often with probability 1. □

Proposition 1 establishes a connection between the throughput in the saturated model and stability in the non-saturated model. Recall that the saturation throughput follows directly from $\theta_i^* = \mu_i \sum_{x\in\Omega:x_i=1} \pi(x)$, with $\pi(\cdot)$ the limiting distribution in Theorem 1.

It might seem natural that a dual property to Proposition 1 holds as well, i.e., all the nodes are stable if $\lambda_i < \theta_i^*$ for all $i = 1, \ldots, N$. It is indeed the case that at least one of the nodes must be stable, as otherwise the network behaves as in the saturated regime, and each node i would have a throughput $\theta_i = \theta_i^*$. This contradicts the fact that $\theta_i \leq \lambda_i < \theta_i^*$ for all nodes. However, it is not the case in general that all the nodes are stable if $\lambda_i < \theta_i^*$ for all $i = 1, \ldots, N$. In order to see that, we next consider an illustrative example.

3.1 Example: ring topology

Consider a 4-node ring topology, i.e., $N = 4$ and $E = \{\{1,2\},\{2,3\},\{3,4\},\{4,1\}\}$. Suppose that $\nu_i \equiv \nu$ and $\mu_i \equiv 1$, so that $\theta_i^* \equiv \theta^* = \nu(1+\nu)/(1+4\nu+2\nu^2)$. Also, let $\lambda_1 = \lambda_2 = \lambda_3 \equiv \lambda < \theta^*$, and $\lambda_4 = 0$. Assume that the arrival processes are Poisson and that the back-off periods and transmission durations are exponentially distributed.

First observe that both nodes 1 and 3 must be stable. In case either of these nodes were unstable, the fraction of time it would be active is bounded from below by

$$(1-\tau_2)\frac{\nu}{1+\nu} \geq (1-\lambda)\frac{\nu}{1+\nu} \geq \left(1 - \frac{\nu(1+\nu)}{1+4\nu+2\nu^2}\right)\frac{\nu}{1+\nu} > \frac{\nu(1+\nu)}{1+4\nu+2\nu^2} > \lambda,$$

which yields a contradiction.

A 'busy period' of node 2 is said to begin when node 2 starts a transmission after at least one transmission of nodes 1 and/or 3 and to end when node 2 completes a transmission which is followed by at least one transmission of nodes 1 and/or 3. The time period U between two successive busy periods of node 2 is referred to as an 'idle period' of node 2. Denote by V the amount of time from the start of a busy period of node 2 until the first packet arrival at node 1 or 3 and by W the possible remaining transmission time of node 2 after V. The number of further transmissions of node 2 after W during that busy period is bounded from above by a geometrically distributed random variable with parameter 1/2. Thus the total amount of time T that node 2 is active during the busy period satisfies $\mathbb{E}\{T\} \leq \mathbb{E}\{V\} + \mathbb{E}\{W\} + 1 \leq \frac{1}{2\lambda} + 2$. Now distinguish two cases: (i) $\mathbb{E}\{U\} \geq \mathbb{E}\{T\} + 1$; and (ii) $\mathbb{E}\{U\} \leq \mathbb{E}\{T\} + 1$, and denote by Z the amount of time that node 2 is idle during the busy period. In case (i), we have

$$\tau_2 = \frac{\mathbb{E}\{T\} - \mathbb{E}\{Z\}}{\mathbb{E}\{T\} + \mathbb{E}\{U\} + \mathbb{E}\{Z\}} \leq \frac{\mathbb{E}\{T\}}{\mathbb{E}\{T\} + \mathbb{E}\{U\}} \leq \frac{\frac{1}{2\lambda} + 2}{\frac{1}{\lambda} + 5} = \frac{1+4\lambda}{2+10\lambda}.$$

It is easily verified that the latter upper bound is less than λ when $\lambda > (1 + \sqrt{11})/10$, and thus node 2 must be unstable when $\lambda > (1+\sqrt{11})/10$.

In case (ii), note that by the time the last transmission of either node 1 or node 3 during an idle period of node 2 ends, one of these two has been inactive for at least an expected amount of time 1. Thus the fraction of time υ that either nodes 1 and 2 or nodes 2 and 3 are inactive, is bounded from below as

$$\upsilon \geq \frac{1 + \mathbb{E}\{Z\}}{\mathbb{E}\{T\} + \mathbb{E}\{U\} + \mathbb{E}\{Z\}} \geq \frac{1}{\mathbb{E}\{T\} + \mathbb{E}\{U\}} \geq \frac{1}{\frac{1}{\lambda} + 5} = \frac{\lambda}{1+5\lambda},$$

where the second inequality follows from the fact that $\mathbb{E}\{T\} + \mathbb{E}\{U\} \geq 1$. Denote by υ_{12} and υ_{23} the fraction of time that nodes 1 and 2 are inactive, and the fraction of time that nodes 2 and 3 are inactive, respectively. Then

$$\min\{\tau_1+\tau_2, \tau_2+\tau_3\} = \min\{1-\upsilon_{12}, 1-\upsilon_{23}\} = 1-\max\{\upsilon_{12}, \upsilon_{23}\} \leq 1-\frac{1}{2}(\upsilon_{12}+\upsilon_{23}) \leq 1-\frac{1}{2}\upsilon \leq \frac{2+9\lambda}{2+10\lambda}.$$

It is easily verified that the latter upper bound is less than 2λ when $\lambda > (5 + \sqrt{185})/40$. Since nodes 1 and 3 are stable, it follows that node 2 must be unstable when $\lambda > (5 + \sqrt{185})/40$.

We conclude that in either case node 2 is unstable when $\lambda > (5+\sqrt{185})/40 < 0.47$, regardless of the value of ν. However, θ^* approaches $1/2$ for ν sufficiently large, and hence the condition $\lambda < \theta^*$ is not sufficient for stability.

Figure 1 shows the stability region (obtained by simulation) of the 4-node ring network as a function of λ_2 and λ_4, for $\mu = 1$ and $\nu = 10$. For arrival rates outside of the area demarcated by the solid line, at least one node is unstable. The area enclosed by the dashed line represents all rates such that $\max\{\lambda_2, \lambda_4\} < \theta^* = 110/241 \approx 0.456$. We compare $\lambda_1 = \lambda_3 = 0.3$ (Figure 1(a)) and $\lambda_1 = \lambda_3 = 0.45$ (Figure 1(b)), so in both cases we have $\max\{\lambda_1, \lambda_3\} < \theta^*$. In Figure 1(a) we see that $\max\{\lambda_2, \lambda_4\} < \theta^*$ is sufficient for stability. However, we see from Figure 1(b) that as λ_1 and λ_3 increase, the stability region decreases in size and no longer includes all $\max\{\lambda_2, \lambda_4\} < \theta^*$.

Observe that the range of values of λ_2 that stabilize the system grows with λ_4, so nodes 2 and 4 benefit from each other's activity. This was already hinted at in the above instability example, where it was shown that node 2 may become unstable when node 4 is removed from the network. Indeed, increasing activity from nodes 2 and 4 forces nodes 1 and 3 to operate in a more efficient fashion, i.e., simultaneous activity, thus increasing spatial reuse.

4 Stability for full interference graphs

In Section 3 we saw that although necessary stability conditions can be obtained by considering the saturated model, the case of a ring network already gives rise to intricate stability conditions. The case of a full interference graph is considerably simpler, and the stability conditions can be explicitly derived, as is shown next. In the remainder, we restrict ourselves to the case of back-off freezing.

We assume that all nodes mutually interfere, so at most one node can be active at any time. Without loss of generality assume that the nodes are ordered such that

$$\frac{\lambda_1}{\nu_1} \leq \frac{\lambda_2}{\nu_2} \leq \cdots \leq \frac{\lambda_N}{\nu_N}.$$

Denote

$$\hat{\tau}_i = \frac{\sigma_i}{1 + \sum_{j=i}^{N} \sigma_j} \Big(1 - \sum_{j=1}^{i-1} \rho_j\Big).$$

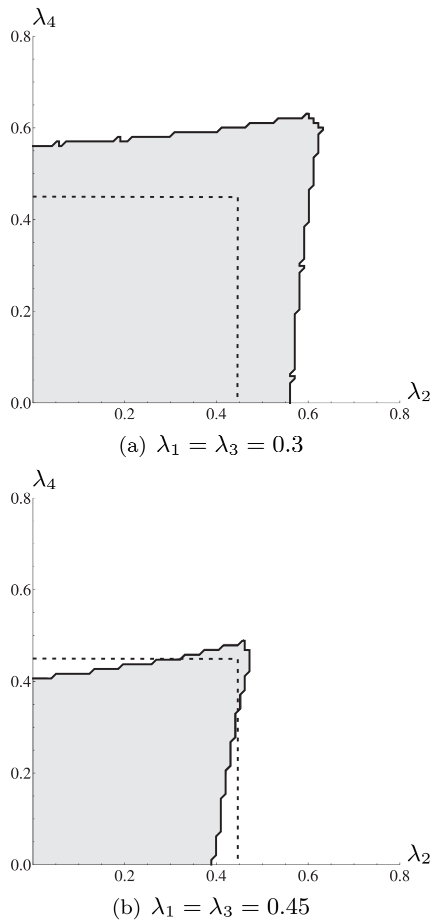

Figure 1: The stability region of a 4-node ring network.

These $\hat{\tau}_i$ may be interpreted as the fraction of time that node i is active, assuming that nodes $1, \ldots, i-1$ are all stable, while nodes $i, \ldots, N$ are all saturated. Also, define $i_{\max} = \max\{i \in \{1, \ldots, N\} : \rho_i < \hat{\tau}_i\}$, with the convention that $i_{\max} = 0$ when $\rho_i > \hat{\tau}_i$ for all $i = 1, \ldots, N$, and assume $\rho_{i_{\max}+1} > \hat{\tau}_{i_{\max}+1}$ in case $i_{\max} < N$. The interpretation of $\hat{\tau}_i$ suggests that node $i_{\max}$ is the stable node with the highest index, as will be shown in the following theorem.

Theorem 2. *Nodes* $1, \ldots, i_{\max}$ *are stable, while nodes* $i_{\max} + 1, \ldots, N$ *are unstable.*

Proof. For compactness, denote by τ_0 the fraction of time that all nodes are inactive. As noted in Section 3, we have that $\theta_i \leq \lambda_i$, with equality when node i is stable, and thus $\tau_i \leq \rho_i$, with equality when node i is stable. In view of the back-off freezing, the back-off process of node i is only running when all nodes are inactive, and hence we have $\tau_0 = \gamma_i/\nu_i$, with γ_i the expected number of back-offs of node i per unit of time. (Without back-off freezing, this relationship still holds for exponential back-off time distributions, but for general back-off time distributions there does not seem to be a simple connection between γ_i and τ_0 in that case.) By definition, the probability that node i has a packet to transmit when a back-off period ends, equals $p_i = \theta_i/\gamma_i$. Combining these two relationships, we obtain the identity $\theta_i = p_i\nu_i\tau_0$ and thus $\tau_i = p_i\sigma_i\tau_0$. In particular, $\tau_i = \sigma_i\tau_0$ when node i is unstable. Hence $\rho_i \leq \sigma_i\tau_0$, i.e., $\lambda_i \leq \nu_i\tau_0$, when node i is stable, while $\rho_i \geq \sigma_i\tau_0$, i.e., $\lambda_i \geq \nu_i\tau_0$ when node i is unstable. It follows that the set of stable nodes is of the form $\{1,\dots,i^*\}$ for some $i^* \in \{0,\dots,N\}$. It remains to be shown that $i^* = i_{\max}$.

First observe that $\sum_{i=0}^N \tau_i = 1$, $\tau_i = \rho_i$ for all $i = 1,\dots,i^*$, and $\tau_i = \sigma_i\tau_0$ for all $i = i^*+1,\dots,N$. This yields

$$\tau_0 = \frac{1}{1+\sum_{j=i^*+1}^N \sigma_j}\Big(1 - \sum_{i=1}^{i^*}\rho_i\Big).$$

Further observe the equivalence relation

$$\rho_i > \frac{\sigma_i}{1+\sum_{j=i}^N \sigma_j}\Big(1 - \sum_{j=1}^{i-1}\rho_j\Big) \Longleftrightarrow \rho_i + \rho_i\sum_{j=i}^N \sigma_j > \sigma_i - \sigma_i\sum_{j=1}^{i-1}\rho_j$$
$$\Longleftrightarrow \rho_i + \rho_i\sum_{j=i+1}^N \sigma_j > \sigma_i - \sigma_i\sum_{j=1}^{i}\rho_j \Longleftrightarrow \rho_i > \frac{\sigma_i}{1+\sum_{j=i+1}^N \sigma_j}\Big(1 - \sum_{j=1}^{i}\rho_j\Big).$$

Since the nodes are indexed such that $\rho_i/\sigma_i = \lambda_i/\nu_i \leq \rho_{i+1}/\sigma_{i+1} = \lambda_{i+1}/\nu_{i+1}$, we obtain the property

$$\rho_i > \hat{\tau}_i \Longrightarrow \rho_{i+1} > \hat{\tau}_{i+1}. \tag{4}$$

Now suppose that $0 \leq i^* < i_{\max}$. The fact that node $i^*+1 \leq N$ is unstable means that

$$\rho_{i^*+1} > \tau_{i^*+1} = \sigma_{i^*+1}\tau_0 = \frac{\sigma_{i^*+1}}{1+\sum_{j=i^*+1}^N \sigma_j}\Big(1 - \sum_{i=1}^{i^*}\rho_i\Big) = \hat{\tau}_{i^*+1}.$$

Property (4) then implies that $\rho_i > \hat{\tau}_i$ for all $i = i^*+1,\dots,N$, which contradicts $i_{\max} \geq i^*+1$, and hence we must have $i^* \geq i_{\max}$.

The fact that node i^* is stable means that

$$\rho_{i^*} = \tau_{i^*} \leq \sigma_{i^*}\tau_0 = \frac{\sigma_{i^*}}{1+\sum_{j=i^*+1}^N \sigma_j}\Big(1 - \sum_{i=1}^{i^*}\rho_i\Big).$$

Property (4) then implies that

$$\rho_{i^*} \leq \frac{\sigma_{i^*}}{1+\sum_{j=i^*}^N \sigma_j}\Big(1 - \sum_{i=1}^{i^*-1}\rho_i\Big) = \hat{\tau}_{i^*},$$

and hence we must have $i_{\max} \geq i^*$. □

The result in Theorem 2 in fact holds for any stationary traffic process as long as the service is infinitely divisible, with the ρ_i values representing the mean amount of traffic generated per time unit as measured in units of transmission time. It is worth observing that the form of the stability conditions is rather reminiscent of those for polling systems with k_i-limited or Weighted Fair Queueing (WFQ) service disciplines [11] and Generalized Processor Sharing (GPS) queues [6, 20].

Noting that $i_{\max} = N$ if and only if $\rho_i < \hat{\tau}_i$ for all $i = 1, \ldots, N$, Theorem 2 in particular gives the following necessary and sufficient condition for all nodes to be stable.

Corollary 1. *All nodes are stable if and only if $\rho_i < \hat{\tau}_i$ for all $i = 1, \ldots, N$.*

The explicit and relatively simple form of the stability condition established in Corollary 1 is highly remarkable as it starkly contrasts with those for slotted Aloha systems, which even for a complete interference graph with three or more nodes have remained largely elusive, see for instance [1, 29, 30] for bounds and partial results.

The next result shows that for full interference graphs, the dual property of Proposition 1 does hold (which in general is not the case; see Section 3.1).

Corollary 2. *All nodes are stable if $\lambda_i < \theta_i^*$ for all $i = 1, \ldots, N$.*

Proof. For full interference graphs we have that $\tau_i^* = \sigma_i/(1 + \sum_{i=1}^N \sigma_i)$. Note that $\lambda_i < \theta_i^*$ for all $i = 1, \ldots, N$ implies $\rho_N < \tau_N^*$ and $\tau_N^* = (1 - \tau_1^* - \cdots - \tau_{N-1}^*)\frac{\sigma_N}{1+\sigma_N} < (1 - \rho_1 - \cdots - \rho_{N-1})\frac{\sigma_N}{1+\sigma_N} = \hat{\tau}_N$. This yields $i_{\max} = N$, which completes the proof. □

We conclude this section by two further consequences of Theorem 2 that are helpful when only the total load $\sum_{i=1}^N \rho_i$ is known. Denote $\sigma_{\min} = \min_{i=1,\ldots,N} \sigma_i$ and $\sigma_{\max} = \max_{i=1,\ldots,N} \sigma_i$.

Corollary 3. *All nodes are stable if $\sum_{i=1}^N \rho_i < \frac{\sigma_{\min}}{1+\sigma_{\min}}$. This condition is sharp in the sense that if $\rho_1 = \cdots = \rho_{N-1} = 0$ and $\sigma_N = \sigma_{\min}$, then $\rho_N < \frac{\sigma_{\min}}{1+\sigma_{\min}}$ is necessary for node N to be stable.*

Corollary 4. *At least one node is unstable if $\sum_{i=1}^N \rho_i > \frac{\sigma_{\max}}{1/N+\sigma_{\max}}$. This condition is sharp in the sense that if $\rho_1 = \cdots = \rho_N = \rho$ and $\sigma_1 = \cdots = \sigma_N = \sigma$, then $\rho < \frac{\sigma}{1+N\sigma}$ is sufficient for all nodes to be stable.*

5 Partial interference graphs

In Section 4 we focused on the case of a full interference graph, and derived explicit necessary and sufficient conditions for stability. In this section we allow for partial interference graphs, and will show that in general the stability conditions cannot be represented in such an explicit form. In particular, we illustrate the difficulties that arise in star topologies, and then argue that all (non-complete) graphs contain such a star topology as an induced subgraph.

5.1 Star topologies

Consider a star topology, where the leaf nodes $1, \ldots, N-1$ all interfere with the root node N, but not with each other, i.e., $E = \{1, \ldots, N-1\} \times \{N\}$. The stability region may then be characterized by:

$$\rho_i < (1 - \rho_N)\frac{\sigma_i}{1 + \sigma_i}, \qquad i = 1, \ldots, N-1, \tag{5}$$

and

$$\rho_N < \hat{\tau}_N, \tag{6}$$

with $\hat{\tau}_N$ representing the fraction of time that node N would be active if it were saturated.

By definition, inequality (6) is necessary and sufficient for the root node N to be stable, and given that node N is stable, the inequalities (5) are necessary and sufficient for all the leaf nodes $1, \ldots, N-1$ to be stable as well. The boundary of the stability region consists of a total of N segments, with $N-1$ linear segments defined by the inequalities (5), where the corresponding leaf node is critically loaded, and 1 segment which is not likely to be linear in general, described by the inequality (6), where the root node is critically loaded.

There does not seem to be a closed-form expression available for $\hat{\tau}_N$ in general, in fact not even for $N = 3$, so that the inequality (6) is not so explicit. The next lemma however provides a useful closed-form lower bound for $\hat{\tau}_N$.

Corollary 5. *Assuming exponential back-off times and transmission durations, we have $\hat{\tau}_N \geq \tau_N^*$, with*

$$\tau_N^* = \frac{\sigma_N}{\sigma_N + \prod_{i=1}^{N-1}(1 + \sigma_i)}.$$

Proof. Noting that the star topology is a bipartite graph, the statement of the corollary follows directly from Proposition 2 (presented in Section 7), with $V_1 = \{N\}$ and $V_2 = \{1, \ldots, N-1\}$. □

Corollary 5 implies that as long as $\rho_N < \tau_N^*$, the root node N is guaranteed to be stable, and thus the conditions (5) are necessary and sufficient for all the leaf nodes $1, \ldots, N-1$ to be stable as well. Noting that $\tau_i^* = (1 - \tau_N^*)\frac{\sigma_i}{1+\sigma_i}$, the inequalities (5) may be expressed as

$$\frac{\rho_i}{\tau_i^*} < \frac{1 - \rho_N}{1 - \tau_N^*}, \qquad i = 1, \ldots, N-1,$$

or

$$\max_{i=1,\ldots,N-1} \frac{\rho_i}{\tau_i^*} < \frac{1 - \rho_N}{1 - \tau_N^*},$$

or

$$\max_{i=1,\ldots,N-1} \frac{\rho_i(1+\sigma_i)}{\sigma_i} < 1 - \rho_N.$$

Thus simple necessary and sufficient conditions for stability arise when the traffic load of the root node N is sufficiently low, in particular when the ratio of the traffic load ρ_N to the saturated throughput τ_N^* is relatively low compared to that of the leaf nodes $1, \ldots, N-1$. Specifically, suppose that $\frac{\rho_N}{\tau_N^*} \leq \max_{i=1,\ldots,N-1} \frac{\rho_i}{\tau_i^*}$, and let $j_{\max} = \arg\max_{j=1,\ldots,N-1} \frac{\rho_j}{\tau_j^*}$. Then $\frac{\rho_{j_{\max}}}{\tau_{j_{\max}}^*} < 1 - \rho_N$ is a simple necessary and sufficient condition for stability of all nodes. In order to prove that, it suffices to show that the latter condition implies that $\rho_N < \tau_N^*$. Suppose that is not the case, then $\max_{i=1,\ldots,N-1} \frac{\rho_i}{\tau_i^*} \geq 1$, so the condition $\max_{i=1,\ldots,N-1} \frac{\rho_i}{\tau_i^*} < \frac{1-\rho_N}{1-\tau_N^*}$ implies $\frac{1-\rho_N}{1-\tau_N^*} > 1$, i.e., $\rho_N < \tau_N^*$. In particular, if $\frac{\rho_N}{\tau_N^*} = \max_{i=1,\ldots,N-1} \frac{\rho_i}{\tau_i^*}$, then $\rho_N < \tau_N^*$ is a simple necessary and sufficient condition for the stability of all nodes.

A further observation is that $\rho_i < \tau_i^*$, $i = 1, \ldots, N$, is a sufficient condition for all nodes $1, \ldots, N$ to be stable. It might seem that this is a trivial fact, which in fact should hold for any interference graph, but that is not the case as was illustrated in the counterexample in Section 3.1. However, an explicit formulation of the *necessary* and sufficient stability condition for star networks stays beyond our reach, as $\hat{\tau}_N$ in (6) is seemingly untractable.

5.2 Stability conditions for general graphs

The fact that an explicit condition for stability in the star network appears elusive, illustrates the difficulty of obtaining the stability condition for general networks. Indeed, an explicit characterization of the stability region is difficult for any network that is not a complete graph, for which the stability condition was explicitly obtained in Theorem 2.

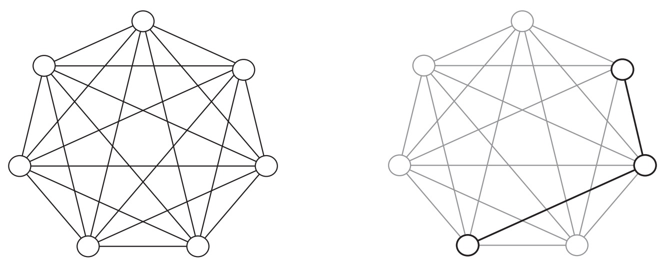

(a) A 7-node complete graph

(b) A 7-node graph containing a 3-node star subgraph

Figure 2: A complete graph and a 3-node star subgraph.

One way to argue this is to show that any network is either a complete graph or contains a 3-node star network as a subgraph. In order to see that, we may focus on a connected graph. Consider an arbitrary node in the graph, say node 1, as well as the set C_1 of all of its neighbors. If one of the nodes in the set C_1 has a neighbor that is not node 1 or in the set C_1, then this induces a 3-node star network. Otherwise, node 1 along with the nodes in the set C_1 make up the entire graph (since the graph is connected). If the nodes in the set C_1 are not fully connected, then this induces a 3-node star network again. Otherwise, node 1 along with the nodes in the set C_1 are fully connected, and the graph is complete. This argument is illustrated in Figure 2.

Hence, the fact that the 3-node star network (for which the stability condition does not seem to admit an explicit characterization) is contained in every non-complete network, provides strong indication of the hardness of the problem of determining the stability region for general networks. That is, characterizing the set of traffic vectors $(\rho_1, \dots, \rho_N)$ for which the system is stable is challenging for given activation rates $\nu_1, \dots, \nu_N$.

In contrast, determining whether there exist activation rates $\nu_1, \dots, \nu_N$ for which the system is stable for a given traffic vector $(\rho_1, \dots, \rho_N)$ is relatively easy. Specifically, in case of a star topology, such activation rates exist if and only if $\rho_N + \max_{i=1,\dots,N-1} \rho_i < 1$. In general, such activation rates exist if and only if the traffic vector $(\rho_1, \dots, \rho_N)$ belongs to the interior of $\text{conv}(\Omega)$. The latter property in fact serves as the basis for various adaptive strategies that achieve stability whenever feasible at all [13, 14, 22, 23, 26, 27].

6 Conclusion

We examined the insensitivity and stability of CSMA networks. We proved that when all nodes are saturated, the stationary distribution of these networks is insensitive to the distribution of both the transmission durations and the back-off times. The insensitivity holds irrespective of whether or not an active node freezes the back-off process of neighboring nodes.

We then turned to a situation where nodes are subject to packet dynamics, and established a simple necessary condition for stability for general interference graphs. Explicit necessary and sufficient stability conditions were derived for the case of complete interference graphs. Moreover, we illustrated the difficulty of deriving similar conditions for partial interference graphs.

In the above analysis, we have used a continuous-time model to capture CSMA dynamics. In practice however, time is slotted and back-off periods are expressed in numbers of slots. As a consequence, neighboring nodes may start transmitting simultaneously, leading to a collision. Extending our results to the case where collisions may occur is difficult. In the saturated scenario however, if we assume that when two nodes are involved in a collision, they start and end the corresponding transmissions at the same epochs, then insensitivity can be shown using the Engset network representation (see the discussion following Theorem 1) only in the case without back-off freezing. It remains unclear whether insensitivity holds when time is slotted in the case with back-off freezing. Regarding stability issues, adding collisions greatly complicates the analysis even in the case of networks with full interference (similar to the classical problem of deriving network stability conditions under Aloha protocols). In Section 4, the absence of collisions simplifies the stability analysis, because in this case, packets are successfully being transmitted whenever the medium is busy.

7 Auxiliary results and remaining proofs

7.1 A stochastic comparison result

Consider a bipartite graph such that $V = V_1 \cup V_2$, with $V_1 \cap V_2 = \emptyset$ and $E \subseteq V_1 \times V_2$. We will show that in the situation where all the nodes are saturated, the thoughputs of the nodes in V_1 and V_2 are lower and higher respectively, then in case only the nodes in V_1 are saturated. Let $V_1 \subseteq \{1, \ldots, N\}$ and define $\theta_i(V_1)$ as the throughput of node i in the modified version of the network where the nodes in V_1 are saturated.

Proposition 2. *Assuming exponential back-off times and transmission durations, we have $\theta_i(V_1) \geq \theta_i^*$ for all $i \in V_1$ and $\theta_i(V_1) \leq \theta_i^*$ for all $i \in V_2$.*

Proof. The proof relies on stochastic coupling [25]. Let $N_i^\lambda(t)$, $i \in V_2$, $N_i^\mu(t)$, $i = 1, \ldots, N$, and $N_i^\nu(t)$, $i = 1, \ldots, N$, be independent Poisson processes of rates λ_i, μ_i, and ν_i, respectively. We will use these Poisson processes to construct processes $X(t) = (X_1(t), \ldots, X_N(t))$ and $Y(t) = (Y_i(t))_{i \in V_2}$, representing the activity process and the queue length process in the scenario where the nodes $i \in V_1$ are saturated, and $X^*(t) = (X_1^*(t), \ldots, X_N^*(t))$ representing the activity process in case all the nodes are saturated. We assume that $X_i(0) \geq X_i^*(0)$ for all $i \in V_1$ and $X_i(0) \leq X_i^*(0)$ for all $i \in V_2$, and allow $Y_i(0)$, $i \in V_2$, to be arbitrary. We will prove that $X_i(t) \geq X_i^*(t)$ for all $i \in V_1$ and $X_i(t) \leq X_i^*(t)$ for all $i \in V_2$. Since the stationary distribution of the processes $X(t)$ and $X^*(t)$ does not depend on the initial state, and $\theta_i(V_1) = \mu_i \mathbb{E}\{X_i\} = \mu_i \mathbb{P}\{X_i = 1\}$ and $\theta_i^* = \mu_i \mathbb{E}\{X_i^*\} = \mu_i \mathbb{P}\{X_i^* = 1\}$, the statement of the proposition then follows.

In order to prove the above inequalities, we will use induction. Let t be a time epoch at which an event occurs in one of the Poisson processes. We will show that if the inequalities hold at time t^-, that they then continue to hold at time t^+. We distinguish three cases, depending on in which of the various Poisson processes the event occurs.

We first consider an event in the process $N_i^\lambda(t)$, reflecting a packet arrival at one of the non-saturated nodes $i \in V_2$. In that case, we set $Y_i(t^+) = Y_i(t^-) + 1$. Note that the values of $X_i(t)$ and $X_i^*(t)$ are not affected, and hence the inequalities trivially continue to be valid. Second, we consider an event in the process $N_i^\mu(t)$, corresponding to a potential transmission completion at node i. In that case, we set $X_i(t^+) = X_i^*(t^+) = 0$, and in case $i \in V_2$,

$$Y_i(t^+) = Y_i(t^-) - X_i(t^-), \tag{7}$$

reflecting a potential packet departure. Since $X_i(t^+) = X_i^*(t^+) = 0$, the inequalities remain trivially satisfied. Third, we consider an event in the process $N_i^\nu(t)$, corresponding to a potential activation at node i. In that case we set $X_i^*(t^+) = 1$ if $X_j^*(t^-) = 0$ for all $j \in C_i$, with C_i representing the set of neighbors of node i. Moreover, in case $i \in V_1$, we set $X_i(t^+) = 1$ if $X_j(t^-) = 0$ for all $j \in C_i$, while in case $i \in V_2$, we set $X_i(t^+) = 1$ if $Y_i(t^-) \geq 1$ and $X_j(t^-) = 0$ for all $j \in C_i$.

The fact that for $i \in V_1$, $C_i \subseteq V_2$, and $X_j(t^-) \leq X_j^*(t^-)$ for all $j \in V_2$ implies that $X_i^*(t^+) = 1$ forces $X_i(t^+) = 1$. Likewise, the fact that for $i \in V_2$, $C_i \subseteq V_1$, and $X_j(t^-) \geq X_j^*(t^-)$ for all $j \in V_1$ implies that $X_i(t^+) = 1$ forces $X_i^*(t^+) = 1$. Hence, the inequalities continue to hold. Also, note that $X_i(t) = 0$ whenever $Y_i(t) = 0$, or equivalently, $X_i(t) = 1$ can only occur when $Y_i(t) \geq 1$, so that (7) leaves $Y_i(t) \geq 0$ for all $t \geq 0$.

It is easily verified that viewed in isolation, the processes $X(t)$ and $Y(t)$ as constructed above obey the same statistical laws as the activity process and the queue length process in the scenario where the nodes $i \in V_1$ are saturated, while the process $X^*(t)$ is governed by the same statistical laws as the activity process in case all the nodes are saturated. □

7.2 Proof Lemma 1

Proof. We have that

$$\beta_u\Big(q_{u,m_i+1}\Big(1-\frac{\alpha_u}{1-\alpha_{m_i+1}}\Big)+\sum_{\substack{l=1\\l\neq u}}^{m_i} q_{u,l}\Big) = \sum_{\substack{l=1\\l\neq u}}^{m_i} \beta_l\Big(q_{l,u}+q_{l,m_i+1}\frac{\alpha_u}{1-\alpha_{m_i+1}}\Big), \quad (8)$$

with $\sum_{u=1}^{m_i} \beta_u = 1$, and

$$\eta_u\Big(r_{u,n_i+1}\Big(1-\frac{\gamma_u}{1-\gamma_{n_i+1}}\Big)+\sum_{\substack{l=1\\l\neq u}}^{n_i} r_{u,l}\Big) = \sum_{\substack{l=1\\l\neq u}}^{n_i} \eta_l\Big(r_{l,u}+r_{l,n_i+1}\frac{\gamma_u}{1-\gamma_{n_i+1}}\Big), \quad (9)$$

with $\sum_{u=1}^{n_i} \eta_u = 1$. It is readily seen that

$$\nu_i = (1-\alpha_{m_{i+1}})^{-1}\sum_{u=1}^{m_i} \beta_u q_{u,m_i+1}, \quad (10)$$

$$\mu_i = (1-\gamma_{n_{i+1}})^{-1}\sum_{u=1}^{n_i} \eta_u r_{u,n_i+1}. \quad (11)$$

In order to show that the π_{PH} in (1) is indeed the limiting distribution of the Markov process of interest, it suffices to show that π_{PH} satisfies the global balance equations of this process. However, rather than doing this directly, we study for each node i the partial balance equations that equate the rate into and out of a state by changes to this node only. As the global balance equations can be obtained by summing these partial balance equations over all nodes, it is sufficient to show that π_{PH} satisfies all partial balance equations.

Let C_i denote the set of neighbors of node i, and define $C_i^+ = C_i \cup \{i\}$. Let $T_k^i(\omega)$ denote the operator that changes the i-th component of ω to k, while leaving intact the other components. When node i is inactive and unblocked we see the following transitions to node i (irrespective of freezing)

$$\pi_{\mathrm{PH}}(\omega)\left(q_{\omega_i,m_i+1}\Big(1-\frac{\alpha_{\omega_i}\gamma_{n_i+1}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}\Big)+\sum_{\substack{k=1\\k\neq\omega_i}}^{m_i} q_{\omega_i,k}\right)=\sum_{k=1}^{n_i}\pi_{\mathrm{PH}}(T_k^i(\omega))r_{k,n_i+1}\frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}$$
$$+\sum_{\substack{k=1\\k\neq\omega_i}}^{m_i}\pi_{\mathrm{PH}}(T_{-k}^i(\omega))\Big(q_{k,\omega_i}+q_{k,m_i+1}\frac{\alpha_{\omega_i}\gamma_{n_i+1}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}\Big),\quad \forall\omega \text{ s.t. } \omega_j\leq -1\ \forall j\in C_i^+. \tag{12}$$

If node i is active we have (irrespective of freezing)

$$\pi_{\mathrm{PH}}(\omega)\left(r_{\omega_i,n_i+1}\Big(1-\frac{\alpha_{m_i+1}\gamma_{\omega_i}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}\Big)+\sum_{\substack{k=1\\k\neq\omega_i}}^{n_i} r_{\omega_i,k}\right)=\sum_{k=1}^{m_i}\pi_{\mathrm{PH}}(T_{-k}^i(\omega))q_{k,m_i+1}\frac{\gamma_{\omega_i}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}$$
$$+\sum_{\substack{k=1\\k\neq\omega_i}}^{n_i}\pi_{\mathrm{PH}}(T_k^i(\omega))\Big(r_{k,\omega_i}+r_{k,n_i+1}\frac{\alpha_{m_i+1}\gamma_{\omega_i}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}\Big),\quad \forall\omega \text{ s.t. } \omega_i\geq 1. \tag{13}$$

The final balance equation concerns all states where node i is inactive, but at least one of its neighbors is active, so node i is blocked. The case of back-off freezing yields a trivial partial balance equation, as in this case no such state can be entered or exited due to changes in the state of node i. On the other hand, when the back-off process of blocked nodes is not frozen, the state can change due to a transition within the back-off process of node i:

$$\pi_{\mathrm{PH}}(\omega)\left(q_{\omega_i,m_i+1}\Big(1-\frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}}\Big)+\sum_{\substack{k=1\\k\neq\omega_i}}^{m_i} q_{\omega_i,k}\right)$$
$$=\sum_{\substack{k=1\\k\neq\omega_i}}^{m_i}\pi_{\mathrm{PH}}(T_{-k}^i(\omega))\Big(q_{k,\omega_i}+q_{k,m_i+1}\frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}}\Big),\quad \forall\omega \text{ s.t. } \omega_i\leq -1,\ \exists j\in C_i:\omega_j\geq 1. \tag{14}$$

We now proceed to show that π_{PH} from (1) indeed satisfies (12)-(14). Substituting π_{PH} into (12), and canceling common terms yields

$$\beta_{\omega_i}\Big(q_{\omega_i,m_i+1}\Big(1-\frac{\alpha_{\omega_i}\gamma_{n_i+1}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}\Big)+\sum_{\substack{k=1\\k\neq\omega_i}}^{m_i} q_{\omega_i,k}\Big)$$

$$= \sum_{\substack{k=1\\k\neq\omega_i}}^{m_i} \beta_k\Big(q_{k,\omega_i}+q_{k,m_i+1}\frac{\alpha_{\omega_i}\gamma_{n_i+1}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}\Big)+\sigma_i\sum_{k=1}^{n_i}\eta_k r_{k,n_i+1}\frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}.$$

By (11), and adding $\beta_{\omega_i}q_{\omega_i,m_i+1}\Big(\frac{\alpha_{\omega_i}\gamma_{n_i+1}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}-\frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}}\Big)$ on both sides, we have

$$\beta_{\omega_i}\Big(q_{\omega_i,m_i+1}\Big(1-\frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}}\Big)+\sum_{\substack{k=1\\k\neq\omega_i}}^{m_i} q_{\omega_i,k}\Big)=\sum_{\substack{k=1\\k\neq\omega_i}}^{m_i}\beta_k q_{k,\omega_i}$$

$$+\sum_{k=1}^{m_i}\beta_k q_{k,m_i+1}\frac{\alpha_{\omega_i}\gamma_{n_i+1}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}+\nu_i\frac{\alpha_{\omega_i}(1-\gamma_{n_i+1})}{1-\alpha_{m_i+1}\gamma_{n_i+1}}-\beta_{\omega_i}q_{\omega_i,m_i+1}\frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}}.$$

By (10) and rearranging on both sides we get

$$\beta_{\omega_i}\Big(q_{\omega_i,m_i+1}\Big(1-\frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}}\Big)+\sum_{\substack{k=1\\k\neq\omega_i}}^{m_i} q_{\omega_i,k}\Big)=\sum_{\substack{k=1\\k\neq\omega_i}}^{m_i}\beta_k\Big(q_{k,\omega_i}+q_{k,m_i+1}\frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}}\Big)$$

$$+\nu_i\frac{\alpha_{\omega_i}\gamma_{n_i+1}(1-\alpha_{m_i+1})}{1-\alpha_{m_i+1}\gamma_{n_i+1}}+\nu_i\frac{\alpha_{\omega_i}(1-\gamma_{n_i+1})}{1-\alpha_{m_i+1}\gamma_{n_i+1}}-\nu_i\frac{\alpha_{\omega_i}}{1-\alpha_{m_i+1}}(1-\alpha_{m_i+1}).$$

Canceling the remaining terms we get (8), so π_{PH} indeed satisfies (12).

Substituting π_{PH} into (13), and canceling common terms yields

$$\sigma_i\eta_{\omega_i}\Big(r_{\omega_i,n_i+1}\Big(1-\frac{\alpha_{m_i+1}\gamma_{\omega_i}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}\Big)+\sum_{\substack{k=1\\k\neq\omega_i}}^{n_i} r_{\omega_i,k}\Big)$$

$$=\sum_{k=1}^{m_i}\beta_k q_{k,m_i+1}\frac{\gamma_{\omega_i}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}+\sigma_i\sum_{\substack{k=1\\k\neq\omega_i}}^{n_i}\eta_k\Big(r_{k,\omega_i}+r_{k,n_i+1}\frac{\alpha_{m_i+1}\gamma_{\omega_i}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}\Big).$$

By (10), and adding $\sigma_i\eta_{\omega_i}r_{\omega_i,n_i+1}\Big(\frac{\alpha_{m_i+1}\gamma_{\omega_i}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}-\frac{\gamma_{\omega_i}}{1-\gamma_{n_i+1}}\Big)$ on both sides, we have

$$\sigma_i\eta_{\omega_i}\Big(r_{\omega_i,n_i+1}\Big(1-\frac{\gamma_{\omega_i}}{1-\gamma_{n_i+1}}\Big)+\sum_{\substack{k=1\\k\neq\omega_i}}^{n_i} r_{\omega_i,k}\Big)$$

$$=\nu_i\frac{\gamma_{\omega_i}(1-\alpha_{m_i+1})}{1-\alpha_{m_i+1}\gamma_{n_i+1}}+\sigma_i\sum_{\substack{k=1\\k\neq\omega_i}}^{n_i}\eta_k r_{k,\omega_i}+\sigma_i\frac{\alpha_{m_i+1}\gamma_{\omega_i}}{1-\alpha_{m_i+1}\gamma_{n_i+1}}\sum_{k=1}^{n_i}\eta_k r_{k,n_i+1}-\sigma_i\eta_{\omega_i}r_{k,n_i+1}\frac{\gamma_{\omega_i}}{1-\gamma_{n_i+1}}.$$

By (11) and rearranging both sides we get

$$\sigma_i \eta_{\omega_i} \left(r_{\omega_i, n_i+1} \Big(1 - \frac{\gamma_{\omega_i}}{1 - \gamma_{n_i+1}} \Big) + \sum_{\substack{k=1 \\ k \neq \omega_i}}^{n_i} r_{\omega_i, k} \right) = \sigma_i \sum_{\substack{k=1 \\ k \neq \omega_i}}^{n_i} \eta_k \Big(r_{k,\omega_i} + r_{k,n_i+1} \frac{\gamma_{\omega_i}}{1 - \gamma_{n_i+1}} \Big)$$

$$+ \nu_i \frac{\gamma_{\omega_i}(1 - \alpha_{m_i+1})}{1 - \alpha_{m_i+1}\gamma_{n_i+1}} + \nu_i \frac{\alpha_{m_i+1}\gamma_{\omega_i}(1 - \gamma_{n_i+1})}{1 - \alpha_{m_i+1}\gamma_{n_i+1}} - \nu_i \frac{\gamma_{\omega_i}}{1 - \gamma_{n_i+1}} (1 - \gamma_{n_i+1}).$$

Canceling the remaining terms we get (9), so π_{PH} indeed satisfies (13).

Thirdly, π_{PH} trivially satisfies (14), as substitution of π_{PH} into this equation immediately gives (8). $\square$

References

[1] V. Anantharam (1991). The stability region of the finite-user slotted ALOHA protocol. *IEEE Trans. Inform. Theory* **37 (3)**, 535-540.

[2] G. Bianchi (2000). Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE J. Sel. Areas Commun.* **18 (3)**, 535–547.

[3] T. Bonald (2006). The Erlang model with non-Poisson call arrivals. In: *Proc. ACM SIGMETRICS/Performance 2006*, Saint-Malo, France, June 26–30, 276–286.

[4] R.R. Boorstyn, A. Kershenbaum (1980). Throughput analysis of multihop packet radio. In: *Proc. ICC*, 1361–1366.

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

[6] S.C. Borst, O.J. Boxma, P.R. Jelenković (2003). Reduced-load equivalence and induced burstiness in GPS queues with long-tailed traffic flows. *Queueing Systems* **43 (4)**, 273–306.

[7] M. Durvy, O. Dousse, P. Thiran (2007). Modeling the 802.11 protocol under different capture and sensing capabilities. In: *Proc. Infocom 2007*, Anchorage AK, May 6–12, 2356–2360.

[8] M. Durvy, O. Dousse, P. Thiran (2008). Border effects, fairness, and phase transitions in large wireless networks. In: *Proc. Infocom 2008*, Phoenix AZ, April 13–18, 601–609.

[9] M. Durvy, O. Dousse, P. Thiran (2009). On the fairness of large CSMA networks. *IEEE J. Sel. Areas Commun.* **27 (7)**, 1093–1104.

[10] M. Durvy, P. Thiran (2006). A packing approach to compare slotted and non-slotted medium access control. In: *Proc. Infocom 2006*, Barcelona, Spain, April 23–26.

[11] C. Fricker, M.R. Jaibi (1994). Monotonicity and stability of periodic polling models. *Queueing Systems* **15**, 211–238.

[12] M. Garetto, T. Salonidis, E.W. Knightly (2008). Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks. *IEEE/ACM Trans. Netw.* **16 (4)**, 864–877.

[13] L. Jiang, D. Shah, J. Shin, J. Walrand (2010). Distributed random access algorithm: scheduling and congestion control. *IEEE Trans. Inform. Theory*, to appear.

[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] F.P. Kelly (1979). *Reversibility and Stochastic Networks*. Wiley, Chichester.

[16] F.P. Kelly (1985). Stochastic models of computer communication systems. *J. Roy. Stat. Soc. B* **47 (3)**, 379–395.

[17] F.P. Kelly (1987). One-dimensional circuit-switched networks. *Ann. Prob.* **15 (3)**, 1166–1179.

[18] F.P. Kelly (1991). Loss networks. *Ann. Appl. Prob.* **1 (3)**, 319–378.

[19] A. Kershenbaum, R.R. Boorstyn, M. Chen (1987). An algorithm for evaluation of throughput in multihop packet radio networks with complex topologies. *IEEE J. Sel. Areas Commun.* **5 (6)**, 1003–1012.

[20] M. Lelarge (2009). Asymptotic behavior of generalized processor sharing queues under subexponential assumptions. *Queueing Systems* **62 (1–2)**, 51–73.

[21] S.C. Liew, C. 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.

[22] J. Liu, Y. Yi, A. Proutière, M. Chiang, H.V. Poor (2008). Maximizing utility via random access without message passing. Technical report MSR-TR-2008-128, Microsoft Research.

[23] P. Marbach, A. Eryilmaz (2008). A backlog-based CSMA mechanism to achieve fairness and throughput-optimality in multihop wireless networks. In: *Proc. 46th Annual Allerton Conf. Commun., Control, Comput.*

[24] K. Medepalli, F.A. Tobagi (2006). Towards performance modeling of IEEE 802.11 based wireless networks: A unified framework and its applications. In: *Proc. Infocom 2006*, Barcelona, Spain, April 23–26.

[25] A. Müller, D. Stoyan (2002). *Comparison Methods for Stochastic Models and Risks*. Wiley.

[26] S. Rajagopalan, D. Shah, J. Shin (2009). Network adiabatic theorem: an efficient randomized protocol for content resolution. In: *Proc. ACM SIGMETRICS/Performance 2009*, Seattle WA, June 15–19, 133–144.

[27] S. Rajagopalan, D. Shah, J. Shin (2009). Aloha that works. Preprint.

[28] Y. Suhov, U.A. Rozikov (2004). A hard-core model on a Cayley tree: An example of a loss network. *Queueing Systems* **46 (1-2)**, 197–212.

[29] W. Szpankowski (1988). Stability conditions for multidimensional queueing systems with computer applications. *Oper. Res.* **36 (6)**, 944–957.

[30] W. Szpankowski (1994). Stability conditions for some distributed systems: Buffered random access systems. *Adv. Appl. Prob.* **26**, 498–515.

[31] X. Wang, K. Kar (2005). Throughput modeling and fairness issues in CSMA/CA based ad-hoc networks. In: *Proc. Infocom 2005*, Miami FL, March 13–17.

[32] S. Zachary, I. Ziedins (1999). Loss networks and Markov random fields. *J. Appl. Prob.* **36 (2)**, 403–414.