Delay performance and mixing times in random-access networks

Document status and date:
Published: 01/01/2013

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
• The final author version and the galley proof are versions of the publication after peer review.
• The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

Download date: 05. Oct. 2023

Delay Performance and Mixing Times in Random-Access Networks[*]

Niek Bouman, Sem C. Borst, Johan S.H. van Leeuwaarden
Eindhoven University of Technology
Den Dolech 2, 5612 AZ Eindhoven, The Netherlands
{n.bouman, s.c.borst, j.s.h.v.leeuwaarden}@tue.nl

Abstract

We explore the achievable delay performance in wireless random-access networks. While relatively simple and inherently distributed in nature, suitably designed queue-based random-access schemes provide the striking capability to match the optimal throughput performance of centralized scheduling mechanisms in a wide range of scenarios. The specific type of activation rules for which throughput optimality has been established, may however yield excessive queues and delays.

Motivated by that issue, we examine whether the poor delay performance is inherent to the basic operation of these schemes, or caused by the specific kind of activation rules. We derive delay lower bounds for queue-based activation rules, which offer fundamental insight in the cause of the excessive delays. For fixed activation rates we obtain lower bounds indicating that delays and mixing times can grow dramatically with the load in certain topologies as well.

1   Introduction

Emerging wireless mesh networks typically lack any centralized access control entity, and instead vitally rely on the individual nodes to operate autonomously and fairly share the medium in a distributed fashion. A particularly popular mechanism for distributed medium access control is provided by the so-called Carrier-Sense Multiple-Access (CSMA) protocol. In the CSMA protocol each node attempts to access the medium after a certain random 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 were originally developed by Boorstyn *et al.* [2], and further pursued by Wang & Kar [30], Durvy *et al.* [6, 7] and Garetto *et al.* [10]. Although the representation of the CSMA back-off mechanism in the above-mentioned models is less detailed than in the landmark work of Bianchi [1], they accommodate a general interference graph and thus cover a broad range of topologies. Experimental results of Liew *et al.* [17] demonstrate that these models, while idealized, provide throughput estimates that match remarkably well with measurements in actual real-life networks.

1

Despite their asynchronous and distributed nature, CSMA-like algorithms have been shown to offer the capability of achieving the full capacity region and thus match the optimal throughput performance of centralized scheduling mechanisms operating in slotted time, see for instance Jiang & Walrand [15], Liu *et al.* [18] and Tassiulas & Ephremides [28]. Based on this observation, various clever algorithms have been developed for finding the back-off rates that yield a particular target throughput vector or that optimize a certain concave throughput utility function in scenarios with saturated buffers, see for instance Jiang *et al.* [14, 15] and Marbach & Eryilmaz [20].

In the same spirit, several powerful approaches have been devised for adapting the transmission periods based on the queue lengths in non-saturated scenarios, see for instance Rajagopalan *et al.* [22], Shah & Shin [24] and Shah *et al.* [25]. Roughly speaking, the latter algorithms provide maximum-stability guarantees under the condition that the transmission durations of the various nodes behave as logarithmic functions of the queue lengths.

Unfortunately, however, simulation experiments demonstrate that such activation rules can induce excessive queues and delays, which has sparked a strong interest in developing approaches for improving the delay performance, see for instance Ghaderi & Srikant [12], Lotfinezhad & Marbach [19], Ni *et al.* [21] and Shah & Shin [23]. In particular, it has been shown that more aggressive schemes, where the transmission durations grow faster as function of the queue lengths, can reduce the delays, see for instance Bouman *et al.* [3].

In order to gain insight in the root cause for the poor delay performance, we establish in the present paper lower bounds for the average steady-state delay. To the best of our knowledge, the derivation of lower bounds for the average steady-state delay in random-access networks has received hardly any attention so far. An interesting paper by Shah *et al.* [26] showed that low-complexity schemes cannot be expected to achieve low delay in arbitrary topologies (unless P equals NP), since that would imply that certain NP-hard problems could be solved efficiently. However, the notion of delay in [26] is a transient one, and it is not exactly clear what the implications are for the average steady-state delay in specific networks, if any.

Jiang *et al.* [13, 16] derived *upper* bounds for the average steady-state delay based on mixing time results for Glauber dynamics, where the mixing time represents the amount of time required for the process to come close to its equilibrium distribution. The bounds show that for sufficiently low load the delay only grows polynomially with the number of nodes in bounded-degree interference graphs. Subramanian & Alanyali [27] presented similar upper bounds for bounded-degree interference graphs with low load based on analysis of neighbor sets and stochastic coupling arguments. While some of the conceptual notions in the present paper are similar (cliques, mixing times), we focus on *lower* rather than upper bounds, and exploit quite different techniques.

The lower bounds that we derive for queue-based activation schemes provide fundamental insight why the kind of rules that guarantee maximum stability yield excessive delays. We further obtain lower bounds for the delay and mixing time in case fixed back-off rates are used. In both cases, the bounds bring to light that the delay and mixing time can grow dramatically with the load of the system. Specifically, we establish that the expected delay grows as $F(1/(1-\rho))$ as $\rho \uparrow 1$, where ρ is the load and $F(\cdot)$ is a superlinear function, implying that the growth rate may be polynomially or even exponentially faster than is typically the case in queueing systems at high load. The specific form and growth rate of the function $F(\cdot)$ depends on the activation rule as well as the topology of the network, as we will show for several scenarios of interest. Various partial versions of the results presented here appeared in Bouman *et al.* [4, 5].

The remainder of the paper is organized as follows. In Section 2 we present a detailed model

description, followed by some preliminary results in Section 3. In Section 4 we derive delay lower bounds for queue-based activation schemes. We establish generic lower bounds for the delay and mixing time in case of fixed back-off rates in Section 5. In Sections 6 and 7 we apply these generic bounds to a canonical class of partite interference graphs, which includes several specific cases of interest such as grid topologies. Simulation experiments are conducted in Section 8 to support the analytical results. In Section 9, we make some concluding remarks and identify topics for further research.

2 Model description

Network, interference graph, and traffic model. We consider a network of several nodes sharing a wireless medium according to a random-access mechanism. The network is represented by an undirected graph $G = (V, E)$ where the set of vertices $V = \{1, \ldots, N\}$ correspond to the various nodes and the set of edges $E \subseteq V \times V$ indicate which pairs of nodes interfere. Nodes that are neighbors in the interference graph are prevented from simultaneous activity, and thus the independent sets of G correspond to the feasible joint activity states of the network. A node is said to be blocked whenever the node itself or any of its neighbors is active, and unblocked otherwise. Define $\Omega \subseteq \{0, 1\}^N$ as the set of all feasible joint activity states of the network.

Packets arrive at node i as a Poisson process of rate λ_i. The packet transmission times at node i are independent and exponentially distributed with mean $1/\mu_i$. Denote by $\rho_i = \lambda_i/\mu_i$ the traffic intensity of node i.

Let $U(t) \in \Omega$ represent the joint activity state of the network at time t, with $U_i(t)$ indicating whether node i is active at time t or not. Denote by $L_i(t)$ the number of packets at node i at time t (including any packet that may be in the process of being transmitted).

Random-access mechanism. The nodes share the medium according to a random-access mechanism. When a node ends an activity period (consisting of possibly several back-to-back packet transmissions), it starts a back-off period. The back-off times of node i are independent and exponentially distributed with mean $1/\nu_i$. The back-off period of a node is suspended whenever it becomes blocked by activity of any of its neighbors, and is only resumed once the node becomes unblocked again. Thus the back-off period of a node can only end when none of its neighbors are active. Now suppose a back-off period of node i ends at time t. Then the node starts a transmission with probability $\phi_i(L_i(t))$, and begins a next back-off period otherwise. When a transmission of node i ends at time t, it releases the medium and begins a back-off period with probability $\psi_i(L_i(t^+))$, or starts the next transmission otherwise. We allow for $\phi_i(0) > 0$ and $\psi_i(0) < 1$, so a node may be active even when its buffer is empty, and transmit dummy packets. A dummy transmission is terminated when a new packet arrives and the transmission of this packet is started immediately. Equivalently, node i may be thought of as activating at an exponential rate $f_i(L_i(t)) = \nu_i \phi_i(L_i(t))$, whenever it is unblocked at time t, and de-activating at rate $g_i(L_i(t)) = \mu_i \psi_i(L_i(t) - 1)$, whenever it is active at time t. For conciseness, the functions $f_i(\cdot)$ and $g_i(\cdot)$ will be referred to as activation and de-activation functions, respectively, and we define $h_i(\cdot) = f_i(\cdot)/g_i(\cdot)$ as the nominal activation function.

Network dynamics. Under the above-described queue-based schemes, the process $\{(U(t), L(t))\}_{t \geq 0}$ evolves as a Markov process with state space $\Omega \times \mathbb{N}_0^N$. Transitions (due to arrivals) from a state

3

(U, L) to $(U, L + e_i)$ occur at rate λ_i, transitions (due to activations) from a state (U, L) with $U_i = 0$ and $U_j = 0$ for all neighbors j of node i to $(U + e_i, L)$ occur at rate $f_i(L_i)$, transitions (due to transmission completions followed back-to-back by a subsequent transmission) from a state (U, L) with $U_i = 1$ to $(U, L - e_i I_{\{L_i > 0\}})$ occur at rate $\mu_i - g_i(L_i)$, transitions (due to transmission completions followed by a back-off period) from a state (U, L) with $U_i = 1$ to $(U - e_i, L - e_i I_{\{L_i > 0\}})$ occur at rate $g_i(L_i)$.

Product-form distribution. We now proceed with some additional notation and preliminary results. For any $u \in \Omega$, define $\pi(u) = \lim_{t \to \infty} \mathbb{P}\{U(t) = u\}$ as the steady-state probability that the activity process resides in state u. Further define $\theta_i = \sum_{u \in \Omega} \pi(u) u_i$ as the steady-state fraction of time that node i is active. Note that for fixed activation and de-activation rates, i.e., $\phi_i(\cdot) \equiv \phi_i$ and $\psi_i(\cdot) \equiv \psi$, the activity process $\{U(t)\}$ does not depend on the process $\{L(t)\}$, and in fact constitutes a reversible Markov process with product-form stationary distribution [2]

$$\pi(u) = Z^{-1} \prod_{i=1}^{N} \sigma_i^{u_i}, \qquad u \in \Omega, \tag{1}$$

and normalization constant

$$Z = \sum_{u \in \Omega} \prod_{i=1}^{N} \sigma_i^{u_i},$$

with $\sigma_i = \nu_i \phi_i / (\mu_i \psi_i)$ representing a nominal activity factor.

Stability. In general it is difficult to establish under what conditions the system is stable, i.e., when the process $\{(U(t), L(t))\}_{t \geq 0}$ is positive recurrent. Denoting by conv$(\cdot)$ the convex hull operator and by int(conv$(\cdot)$) its interior, it is easily seen that $(\rho_1, \ldots, \rho_N) \in$ int(conv(Ω)) is a necessary condition for stability.

In [12, 22, 24] it is shown that this condition is in fact also sufficient for activation and de-activation functions $f_i(l) = r_i(l)/(1 + r_i(l))$ and $g_i(l) = 1/(1 + r_i(l))$ with suitably chosen $r_i(\cdot)$, e.g., $r_i(l) = \log(l + 1)$. For more aggressive queue-based activation functions, [11] shows that the necessary condition is not always sufficient though.

In the case of fixed activation and de-activation rates, a simple necessary and sufficient condition for stability is $\rho_i < \theta_i$, for all $i = 1, \ldots, N$. Furthermore, there exists a unique vector $(\sigma_1, \ldots, \sigma_N)$ that yields $(\theta_1, \ldots, \theta_N) \in$ int(conv(Ω)) [15, 29]. Hence, for any traffic intensity vector obeying the necessary stability condition, $(\rho_1, \ldots, \rho_N) \in$ int(conv(Ω)), there exists a vector $(\sigma_1, \ldots, \sigma_N)$ such that $(\rho_1, \ldots, \rho_N) < (\theta_1, \ldots, \theta_N) \in$ int(conv(Ω)), though determining the right vector $(\sigma_1, \ldots, \sigma_N)$ is non-trivial in general.

3 Preliminary results

In this section we state some preliminary results in preparation for the derivation of delay lower bounds in the next sections. Throughout we assume that the system under consideration is stable, because otherwise such lower bounds are not particularly meaningful. More specifically, we derive lower bounds for the expected aggregate stationary queue length in subsets of nodes $\mathcal{A} \subseteq V$. Note that, using Little's law, this also provides a lower bound for the expected aggregate stationary

delay. That is, $\sum_{i \in \mathcal{A}} \mathbb{E}\{L_i\} \geq \alpha$ implies that $\sum_{i \in \mathcal{A}} \lambda_i \mathbb{E}\{W_i\} \geq \alpha$, with W_i a random variable representing the delay (waiting time plus service time) of an arbitrary packet at node i.

The notion of a *clique* will play a pivotal role in the derivation of the lower bounds. A clique is a subset $\mathcal{C} \subseteq V$ of vertices in the interference graph G such that the subgraph induced by $\mathcal{C}$ is complete. Note that in a clique at most one node can be active at a time. The aggregate load in a clique should therefore be less than one if the system is to be stable. For compactness, we use the notation $\lambda_{\mathcal{C}} = \sum_{j \in \mathcal{C}} \lambda_j$ and $\rho_{\mathcal{C}} = \sum_{j \in \mathcal{C}} \rho_j$. We say that a clique $\mathcal{C}$ is in heavy traffic when $\rho_{\mathcal{C}}$ is close to one. Further we denote by $L_{i,\mathcal{C}}$ the number of packets at node i at an arbitrary epoch during a non-serving interval for the clique $\mathcal{C}$, i.e., a time interval during which none of the nodes in $\mathcal{C}$ is transmitting a packet.

Observe that the total number of packets in the clique $\mathcal{C}$ is bounded from below by that in a single node carrying the aggregate traffic, yielding the simple lower bound

$$\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\} \geq \frac{\lambda_{\mathcal{C}} \sum_{i \in \mathcal{C}} \lambda_i / \mu_i^2}{1 - \rho_{\mathcal{C}}} + \rho_{\mathcal{C}}. \tag{2}$$

Thus, the total expected number of packets in any system grows at least linearly in $1/(1 - \rho)$ as ρ increases to one, with $\rho = \max_{\mathcal{C}} \rho_{\mathcal{C}}$ the maximum traffic intensity in any clique.

The lower bound in (2) is only based on sheer load considerations and does not account for the effect of the back-off mechanism. In the next sections we will derive lower bounds for queue-based strategies as well as fixed-rate strategies that do capture the effect of the back-off mechanism, and turn out to be considerably tighter and exhibit superlinear growth in $1/(1 - \rho_{\mathcal{C}})$.

The derivation of the lower bounds starts from the observation that stability of the system requires the non-serving intervals for a clique in heavy traffic to be short or happen infrequently. That is, in each clique, most of the time, one of the nodes should be active, since otherwise the average rate of arriving packets would exceed the average rate of departing packets. For this to be the case, the activity factors should be big at high load. The next lemma quantifies this statement.

Lemma 3.1. *Assume that the system is stable. Then, for any clique $\mathcal{C} \subseteq V$ containing node i,*

$$\frac{\mathbb{E}\{f_i(L_{i,\mathcal{C}})\}}{\mathbb{E}\{g_i(L_i)\}} \geq \frac{\rho_i}{1 - \rho_{\mathcal{C}}}. \tag{3}$$

Proof Observing that the mean number of activations at node i equals the mean number of de-activations at node i per unit of time, we obtain

$$\mathbb{E}\{f_i(L_i) \mathrm{I}_{\{U_j = 0 \text{ for all } j \in \mathcal{N}_i^+\}}\} = \rho_i \mathbb{E}\{g_i(L_i^d)\}, \tag{4}$$

where $\mathcal{N}_i^+$ denotes the set of neighbors of node i in the graph G, along with i itself, and L_i^d denotes the number of packets waiting for transmission at node i at a departure epoch. Note that L_i^d is, in distribution, equal to L_i.

Further,

$$\mathbb{E}\{f_i(L_i) \mathrm{I}_{\{U_j = 0 \text{ for all } j \in \mathcal{N}_i^+\}}\} \leq \mathbb{E}\{f_i(L_i) \mathrm{I}_{\{U_j = 0 \text{ for all } j \in \mathcal{C}\}}\}$$
$$= \mathbb{E}\{f_i(L_{i,\mathcal{C}})\} \mathbb{P}\{U_j = 0 \text{ for all } j \in \mathcal{C}\}. \tag{5}$$

5

Since the events $\{U_j = 1\}$ are mutually exclusive for all $j \in \mathcal{C}$, it follows that

$$\mathbb{P}\{U_j = 0 \text{ for all } j \in \mathcal{C}\} = 1 - \mathbb{P}\{U_j = 1 \text{ for some } j \in \mathcal{C}\}$$
$$= 1 - \sum_{j \in \mathcal{C}} \mathbb{P}\{U_j = 1\} = 1 - \sum_{j \in \mathcal{C}} \theta_j. \tag{6}$$

Thus, we find (3) using (4), (5), and the fact that $\rho_i \leq \theta_i$ for all $i \in V$ is a necessary condition for stability of the system. $\square$

In particular, for fixed-rate strategies it follows that stability entails

$$\sigma_i \geq \frac{\rho_i}{1 - \rho_{\mathcal{C}}}, \tag{7}$$

which in fact could also have been established using the product-form distribution (1).

Lemma 3.1 shows that the activity factors in each clique should be big at high load. In the next sections we will demonstrate that this also causes the delay and mixing time to grow dramatically in heavy traffic.

Queue-based strategies. For queue-based strategies we examine in Section 4 activation functions that are such that a node becomes increasingly more aggressive when the total number of packets at that node increases. For that natural class of activations functions, we exploit the result of Lemma 3.1 to find a lower bound of the form $h^{-1}(1/(|\mathcal{C}|(1 - \rho_{\mathcal{C}})))$ for the aggregate number of packets in the clique $\mathcal{C}$, where $h^{-1}(\cdot)$ is the inverse function of $h(\cdot)$.

A prominent example is $f(l) = r(l)/(1 + r(l))$ and $g(l) = 1/(1 + r(l))$ with $r(l) = \log(l + 1)$, so that $h^{-1}(l) = \exp(l) - 1$, the class of backlog-based strategies for which maximum stability is guaranteed as mentioned earlier. In this case we find that the queue length scales at least exponentially in $1/(1 - \rho_{\mathcal{C}})$.

Fixed-rate strategies. In the case of fixed-rate strategies the delay lower bounds revolve around two simple observations: (i) high activation rates cause long mixing times, in particular slow transitions between dominant activity states; (ii) slow transitions between dominant states imply long starvation periods for some nodes, and hence huge queue lengths and delays. In Section 5 we formalize (ii), and establish lower bounds for the expected aggregate weighted queue length and delay in terms of the expected return times of the process $\{U(t)\}$.

In order to lower bound these return times, we will build in Sections 6 and 7 on insight (i) for a canonical class of partite interference graphs. That is, we examine topologies where the nodes belong to one of K different components such that nodes in the same component do not interfere with each other and every node belongs to a clique of size K (of which the other $K - 1$ nodes necessarily belong to $K - 1$ different components). This class of K-partite interference graphs covers a wide range of network topologies with nearest-neighbor interference, e.g., linear topologies, ring networks with an even number of nodes, two-dimensional grid networks, tori (two-dimensional grid networks with a wrap-around boundary), and *complete* K-partite graphs, where all nodes are connected except those that belong to the same component, with star topologies as a prime example.

4 Queue-based strategies

In this section we derive delay lower bounds for queue-based strategies that use a concave activation function or a convex de-activation function.

Theorem 4.1. *Assume λ_i, ν_i, μ_i, $f_i(\cdot)$ and $g_i(\cdot)$, $i = 1, \ldots, N$, are such that the system is stable. Then, for any clique $\mathcal{C} \subseteq V$,*

(i) *If $f_i(\cdot) \equiv f(\cdot)$ for $i \in \mathcal{C}$ is an increasing concave function and $g_i(\cdot) \geq \xi_i > 0$ for $i \in \mathcal{C}$, then*

$$\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\} \geq \frac{\lambda_{\mathcal{C}} \sum_{i \in \mathcal{C}} \lambda_i / \mu_i^2}{1 - \rho_{\mathcal{C}}} + |\mathcal{C}| f^{-1}\Big(\frac{1}{|\mathcal{C}|} \frac{\sum_{i \in \mathcal{C}} \rho_i \xi_i}{1 - \rho_{\mathcal{C}}}\Big) + \rho_{\mathcal{C}}. \tag{8}$$

(ii) *If $f_i(\cdot) \leq \xi_i$ for $i \in \mathcal{C}$ and $g_i(\cdot) \equiv g(\cdot)$ for $i \in \mathcal{C}$ is a decreasing convex function, then*

$$\sum_{i \in \mathcal{C}} \rho_i \mathbb{E}\{L_i\} \geq \rho_{\mathcal{C}} g^{-1}\Big(\frac{(1 - \rho_{\mathcal{C}}) \sum_{i \in \mathcal{C}} \xi_i}{\rho_{\mathcal{C}}}\Big). \tag{9}$$

(iii) *If $f_i(\cdot) \equiv f(\cdot)$ for $i \in \mathcal{C}$ is an increasing concave function and $g_i(\cdot) \equiv g(\cdot)$ for $i \in \mathcal{C}$ is a decreasing convex function, then*

$$\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\} \geq h^{-1}\Big(\frac{1}{|\mathcal{C}|} \frac{\rho_{\mathcal{C}}}{1 - \rho_{\mathcal{C}}}\Big). \tag{10}$$

Proof The Fuhrmann-Cooper decomposition property [9] (applied to the total number of packets in the clique $\mathcal{C}$) implies

$$\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\} = \frac{\lambda_{\mathcal{C}} \sum_{i \in \mathcal{C}} \lambda_i / \mu_i^2}{1 - \rho_{\mathcal{C}}} + \sum_{i \in \mathcal{C}} \mathbb{E}\{L_{i,\mathcal{C}}\} + \rho_{\mathcal{C}}. \tag{11}$$

This corroborates (2) since the second term in (11) is non-negative, but in case (i) that term might in fact be dominant as we now proceed to show. From (3) we know that, in case (i),

$$(1 - \rho_{\mathcal{C}}) \sum_{i \in \mathcal{C}} \mathbb{E}\{f_i(L_{i,\mathcal{C}})\} \geq \sum_{i \in \mathcal{C}} \rho_i \xi_i.$$

Since $f(\cdot)$ is concave, it follows from Jensen's inequality that

$$\sum_{i \in \mathcal{C}} \mathbb{E}\{f(L_{i,\mathcal{C}})\} \leq |\mathcal{C}| f\Big(\frac{1}{|\mathcal{C}|} \sum_{i \in \mathcal{C}} \mathbb{E}\{L_{i,\mathcal{C}}\}\Big). \tag{12}$$

Because $f(\cdot)$ is increasing we thus get

$$\sum_{i \in \mathcal{C}} \mathbb{E}\{L_{i,\mathcal{C}}\} \geq |\mathcal{C}| f^{-1}\Big(\frac{1}{|\mathcal{C}|} \frac{\sum_{i \in \mathcal{C}} \rho_i \xi_i}{1 - \rho_{\mathcal{C}}}\Big),$$

which completes the proof for case (i).

The proof for case (ii) proceeds along similar lines. From (3) we obtain

$$\sum_{i \in \mathcal{C}} \rho_i \mathbb{E}\{g(L_i)\} \leq (1 - \rho_{\mathcal{C}}) \sum_{i \in \mathcal{C}} \xi_i.$$

7

Since $g(\cdot)$ is convex, it follows from Jensen's inequality that

$$\sum_{i \in \mathcal{C}} \rho_i \mathbb{E}\{g(L_i)\} \geq \rho_{\mathcal{C}} g\Big(\frac{1}{\rho_{\mathcal{C}}} \sum_{i \in \mathcal{C}} \rho_i \mathbb{E}\{L_i\}\Big). \tag{13}$$

Since $g(\cdot)$ is decreasing we thus get

$$\sum_{i \in \mathcal{C}} \rho_i \mathbb{E}\{L_i\} \geq \rho_{\mathcal{C}} g^{-1}\Big(\frac{(1 - \rho_{\mathcal{C}}) \sum_{i \in \mathcal{C}} \xi_i}{\rho_{\mathcal{C}}}\Big),$$

yielding (9).

To prove case (iii), note that combining (11) and (12) gives

$$\sum_{i \in \mathcal{C}} \mathbb{E}\{f(L_{i,\mathcal{C}})\} \leq |\mathcal{C}| f\Big(\frac{1}{|\mathcal{C}|}\Big(\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\} - \frac{\lambda_{\mathcal{C}} \sum_{i \in \mathcal{C}} \lambda_i/\mu_i^2}{1 - \rho_{\mathcal{C}}} - \rho_{\mathcal{C}}\Big)\Big)$$

and hence, because $f(\cdot)$ is increasing,

$$\sum_{i \in \mathcal{C}} \mathbb{E}\{f(L_{i,\mathcal{C}})\} \leq |\mathcal{C}| f\Big(\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\}\Big).$$

Further, since $\rho_i \leq \rho_{\mathcal{C}}$ for $i \in \mathcal{C}$ and because $g(\cdot)$ is decreasing we obtain from (13) that

$$\sum_{i \in \mathcal{C}} \rho_i \mathbb{E}\{g(L_i)\} \geq \rho_{\mathcal{C}} g\Big(\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\}\Big).$$

From (3) we then find

$$|\mathcal{C}| f\Big(\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\}\Big)(1 - \rho_{\mathcal{C}}) \geq \rho_{\mathcal{C}} g\Big(\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\}\Big),$$

or

$$h\Big(\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\}\Big) \geq \frac{1}{|\mathcal{C}|} \frac{\rho_{\mathcal{C}}}{1 - \rho_{\mathcal{C}}}.$$

Thus as $h(\cdot) = f(\cdot)/g(\cdot)$ is increasing because $f(\cdot)$ is increasing and $g(\cdot)$ is decreasing, we get (10). $\qed$

The three cases covered in Theorem 4.1 all reveal the same effect, namely that the mean number of packets in a clique is at least of the order of $h^{-1}(1/(|\mathcal{C}|(1 - \rho_{\mathcal{C}})))$, where $h^{-1}(\cdot)$ is the inverse function of $h(\cdot)$. In case (ii) this effect is observed because the argument of $g^{-1}(\cdot)$ is reciprocal. Further, noting that $f(l) = \log(l + 1)/(1 + \log(l + 1))$ is an increasing concave function and $g(l) = 1/(1 + \log(l + 1))$ is a decreasing convex function, we have

$$\sum_{i \in \mathcal{C}} \mathbb{E}\{L_i\} \geq \mathrm{Exp}\Big(\frac{1}{|\mathcal{C}|} \frac{\rho_{\mathcal{C}}}{1 - \rho_{\mathcal{C}}}\Big) - 1$$

for the class of functions for which maximum stability is guaranteed.

The results of Theorem 4.1 suggest that in order to improve the delay performance one should use more aggressive access schemes. In fact, if $h(\cdot)$ is a superlinear function, i.e., if $h(\cdot)$ grows faster than linear, we find a lower bound that is loose in heavy traffic and (2) provides a better lower bound in that case. Remember however that maximum stability is not guaranteed in case a superlinear function $h(\cdot)$ is used, hence the delay performance might actually deteriorate, and even instability could occur as shown in [11].

5 Fixed-rate strategies

In the previous section we derived delay bounds for queue-based activation rules and we saw that the type of activation rules for which throughput optimality has been established yield excessive delays and queues. We now proceed to construct lower bounds for the expected aggregate weighted queue length and delay in the case of fixed activation and de-activation rates, i.e., we take $\phi_i(\cdot) \equiv \phi_i$ and $\psi_i(\cdot) \equiv \psi_i$.

We first introduce some useful notation. Define $Q(S)$ as the transition rate out of the subset $S \subseteq \Omega$, i.e.,

$$Q(S) = \sum_{u \in S} \sum_{u' \in \Omega \setminus S} \pi(u) q(u, u') = \sum_{u \in \Omega \setminus S} \sum_{u' \in S} \pi(u) q(u, u'),$$

with $q(u, u')$ denoting the transition rate from state u to state u' of the component $\{U(t)\}$ of the Markov process as specified in Section 2, i.e., $q(u, u+e_i) = \nu_i \phi_i$ and $q(u+e_i, u) = \mu_i \psi_i$, $u, u+e_i \in \Omega$. With minor abuse of notation, denote by $\pi(S) = \sum_{u \in S} \pi(u)$ the fraction of time that the system resides in one of the activity states in the subset S. The bottleneck ratio of the subset S is defined as

$$\Phi(S) = \frac{Q(S)}{\pi(S)}.$$

Further define for arbitrary weights $w \in \mathbb{R}_+^N$ and for any $\mathcal{A} \subseteq V$, $S \subseteq \Omega$,

$$Y(w, \mathcal{A}, S) = \max_{u \in S} \sum_{i \in \mathcal{A}} w_i \mu_i u_i,$$

and denote

$$D(w, \mathcal{A}, S) = \sum_{i \in \mathcal{A}} w_i \lambda_i - Y(w, \mathcal{A}, S).$$

The coefficient $Y(w, \mathcal{A}, S)$ represents the maximum aggregate weighted service rate of the nodes in $\mathcal{A}$ when the system resides in one of the activity states in the subset S. Noting that $\sum_{i \in \mathcal{A}} w_i \lambda_i$ is the weighted arrival rate of the nodes in $\mathcal{A}$, the coefficient $D(w, \mathcal{A}, S)$ may thus be interpreted as the minimum drift in the aggregate weighted queue length of the nodes in $\mathcal{A}$ when the system resides in one of the activity states in the subset S.

Proposition 5.1. *For any $w \in \mathbb{R}_+^N$, $\mathcal{A} \subseteq V$,*

$$\sum_{i \in \mathcal{A}} w_i \mathbb{E}\{L_i\} \geq \frac{1}{2} \max_{S \subseteq \Omega} D(w, \mathcal{A}, S) \pi(S) \frac{1}{\Phi(S)}. \tag{14}$$

Proof Denote by T_S a random variable representing the equilibrium return time to the subset of activity states $\Omega \setminus S$ and denote by T_S^e a random variable representing the elapsed equilibrium lifetime of T_S, i.e.,

$$\mathbb{P}\{T_S^e < t\} = \frac{1}{\mathbb{E}\{T_S\}} \int_{s=0}^{t} \mathbb{P}\{T_S > s\} \mathrm{d}s. \tag{15}$$

Now observe that when the system resides in one of the activity states in S, which is the case with probability $\pi(S)$, the aggregate weighted queue length of the nodes in $\mathcal{A}$ have experienced a drift no less than $D(w, \mathcal{A}, S)$ for an expected amount of time $\mathbb{E}\{T_S^e\}$. This observation indicates

9

that the expected aggregate weighted queue length of the nodes in $\mathcal{A}$ is bounded from below by $\pi(S)D(w, \mathcal{A}, S)\mathbb{E}\{T_S^e\}$ for any choice of S and hence

$$\sum_{i \in \mathcal{A}} w_i \mathbb{E}\{L_i\} \geq \max_{S \subseteq \Omega} D(w, \mathcal{A}, S)\pi(S)\mathbb{E}\{T_S^e\}. \tag{16}$$

Using (15) we obtain

$$\mathbb{E}\{T_S^e\} = \frac{1}{\mathbb{E}\{T_S\}} \int_{t=0}^{\infty} \int_{s=t}^{\infty} \mathbb{P}\{T_S > s\}\mathrm{d}s\mathrm{d}t = \frac{\mathbb{E}\{T_S^2\}}{2\mathbb{E}\{T_S\}} \geq \frac{1}{2}\mathbb{E}\{T_S\}.$$

Finally, because $Q(S)$ is the expected number of times the process enters S per unit of time and $\mathbb{E}\{T_S\}$ is the expected amount of time the process stays in S after entering, the expected fraction of the time the process resides in S, $\pi(S)$, is given by $\pi(S) = Q(S)\mathbb{E}\{T_S\}$. Thus, $\mathbb{E}\{T_S\} = \frac{1}{\Phi(S)}$, and (14) follows. $\square$

The question arises how to choose S such that the maximum and thus the tightest possible lower bound in (14) is obtained. Evidently, the more S includes states with some of the nodes in $\mathcal{A}$ active, the larger the potential aggregate weighted service rate of the nodes in $\mathcal{A}$, i.e., the larger $Y(w, \mathcal{A}, S)$, and the smaller $D(w, \mathcal{A}, S)$. In other words, we need to ensure that S excludes some of the states with nodes in $\mathcal{A}$ active. Indeed, if S includes all states with maximal subsets of the nodes in $\mathcal{A}$ active, then $Y(w, \mathcal{A}, S) = \max_{u \in \Omega} \sum_{i=1}^{N} \hat{w}_i \mu_i u_i$, with $\hat{w}_i = w_i$ if $i \in \mathcal{A}$ and $\hat{w}_i = 0$ otherwise. The fact that $(\rho_1, \ldots, \rho_N) \in \text{int}(\text{conv}(\Omega))$ then implies that $Y(w, \mathcal{A}, S) \geq \sum_{i=1}^{N} \hat{w}_i \mu_i \rho_i = \sum_{i=1}^{N} \hat{w}_i \lambda_i = \sum_{i \in \mathcal{A}} w_i \lambda_i$, so that $D(w, \mathcal{A}, S) \leq 0$, yielding an irrelevant lower bound. However, observe that the expected equilibrium return time to $\Omega \setminus S$, denoted $\mathbb{E}\{T_S\}$, may be small when S includes very few states. Hence, to obtain the sharpest possible lower bound, it may not necessarily be optimal to exclude all the states with nodes in $\mathcal{A}$ active from S. For high values of ν, which are necessary for stability at high load as Lemma 3.1 showed, the above argument suggests that we should choose S so that it contains a state with many active nodes, while the boundary of S only contains states with few active nodes.

Define

$$\partial S = \{u \in S : \sum_{u' \notin S} q(u, u') > 0\}$$

as the 'boundary' of S and $K(S', \mathcal{A}') = \max_{u \in S'} \sum_{i \in \mathcal{A}'} u_i$. In order to get a tight lower bound in (14) we thus need to find a subset S such that $K(S, V)$ is large, $K(\partial S, V)$ is small, and $K(S, \mathcal{A})$ is small.

We will now first give an example to illustrate the use of Proposition 5.1.

Example 5.1. Suppose that S is such that $u + e_i \notin \Omega \setminus S$ for all $u \in \partial S$. In case $\phi_i \equiv 1$, $\psi_i \equiv 1$, $\mu_i \equiv 1$ and $\nu_i \equiv \nu \geq 1$, we then have $Q(S) \leq N\pi(\partial S)$, and thus using (1),

$$\frac{1}{\Phi(S)} = \frac{\pi(S)}{Q(S)} \geq \frac{\sum_{u \in S} \pi(u)}{N \sum_{u \in \partial S} \pi(u)} = \frac{\sum_{u \in S} \nu^{\sum_{i=1}^{N} u_i}}{N \sum_{u \in \partial S} \nu^{\sum_{i=1}^{N} u_i}} \geq \frac{1}{N} \nu^{K(S,V)-K(\partial S,V)}.$$

We thus see that in this example we indeed need to choose S such that $K(S, V) - K(\partial S, V)$ is maximized.

Now suppose the interference graph is a symmetric complete bipartite graph. That is, the nodes in $V_1 = \{1, \ldots, N/2\}$ interfere with, and only with, the nodes in $V_2 = \{N/2 + 1, \ldots, N\}$. In this case we have $K(S, V) \leq N/2$. Further, as S is such that $u + e_i \notin \Omega \setminus S$, we have $S = \Omega$ if and only if $K(\partial S, V) = 0$. Thus, because $S = \Omega$ yields an irrelevant lower bound, we have $K(\partial S, V) \geq 1$.

Assuming that $\mathcal{A} \subseteq V_1$ it is clear that $K(S, \mathcal{A}) = 0$ if S only contains states where nodes in V_2 are active. Hence in this case we should choose $S = \{u \in \Omega : \sum_{i \in V_2} u_i \geq 1\}$, the set of activity states where at least one of the nodes in V_2 is active, as this gives $K(S, V) = N/2$, $K(\partial S, V) = 1$ and $K(S, \mathcal{A}) = 0$. We thus see that the delay grows at least as fast as $\nu^{N/2-1}$.

As mixing times are typically long when transitions between dominant activity states are slow, it is likely that we can construct a lower bound for the mixing time that is similar to (14). The mixing time of a process represents the amount of time required for the process to come close to its equilibrium distribution, and is formally defined as

$$t_{\mathrm{mix}}(\epsilon) = \inf\{t : d(t) \leq \epsilon\},$$

where $d(t)$ denotes the maximal distance (in total variation) between $U(t)$ and π, i.e.,

$$d(t) = \max_{U(0) \in \Omega} \frac{1}{2} \sum_{u \in \Omega} |\mathbb{P}\{U(t) = u\} - \pi(u)|.$$

As the next proposition shows, the bottleneck ratio $\Phi(\cdot)$ provides a lower bound on the mixing time of the activity process $\{U(t)\}$.

Proposition 5.2. *The mixing time of $\{U(t)\}$ satisfies*

$$t_{\mathrm{mix}}(\epsilon) \geq \max_{S \subseteq \Omega} (1 - 2\epsilon - \pi(S)) \frac{1}{\Phi(S)}. \tag{17}$$

Proof Zocca *et al.* [31] show that

$$t_{\mathrm{mix}}(\epsilon) \geq (1 - 2\epsilon - r) \frac{1}{\Phi_r^*},$$

with

$$\Phi_r^* = \min_{\{S \subseteq \Omega : \pi(S) \leq r\}} \Phi(S).$$

Therefore,

$$
\begin{aligned}
t_{\mathrm{mix}}(\epsilon) &\geq \max_r \max_{\{S \subseteq \Omega : \pi(S) \leq r\}} (1 - 2\epsilon - r) \frac{1}{\Phi(S)} \\
&= \max_r \max_{\{S \subseteq \Omega : \pi(S) = r\}} (1 - 2\epsilon - r) \frac{1}{\Phi(S)},
\end{aligned}
$$

and (17) follows. $\qquad\qquad\square$

We thus found a lower bound for the mixing time that has a similar form as the bound we found in Proposition 5.1 for the aggregate weighted queue length. Note, however, that to find a tight lower bound for the mixing time, for sufficiently small ϵ, we only need $K(S, V)$ to be large and $K(\partial S, V)$ to be small.

6 Complete partite graphs

In the previous section we derived generic lower bounds for the expected aggregate weighted queue length and delays in terms of the bottleneck ratio of any subset $S \subseteq \Omega$, an approach that is also used to find a lower bound for the mixing time of the activity process $\{U(t)\}$. In this section and the next we describe how to find a subset $S \subseteq \Omega$ with the desired properties discussed in the previous section, for a broad class of K-partite interference graphs. We additionally assume that each of the nodes belongs to at least one clique of size K (of which the other $K-1$ nodes necessarily belong to $K-1$ different components).

We first introduce some further notation and state a few preparatory lemmas. Denote by $V_k \subseteq V$ the subset of nodes that belong to the k-th component and $M_k = |V_k|$, $k = 1, \ldots, K$. For compactness, define

$$\Upsilon_k = \prod_{i \in V_k} (1 + \sigma_i) - 1 = \sum_{I \subseteq V_k} \prod_{i \in I} \sigma_i - 1 = \sum_{\emptyset \neq I \subseteq V_k} \prod_{i \in I} \sigma_i.$$

In particular when $\sigma_i \equiv \hat{\sigma}_k$ for all $i \in V_k$, we have $\Upsilon_k = (1 + \hat{\sigma}_k)^{M_k} - 1$.

Throughout we assume that $\rho_i = \hat{\rho}_k$ for all $i \in V_k$, and denote $\rho = \sum_{k=1}^{K} \hat{\rho}_k$, and $\rho_{\min} = \min_{k=1,\ldots,K} \hat{\rho}_k$. For convenience, we also assume $\phi_i \equiv 1$, $\psi_i \equiv 1$, $\mu_i \equiv 1$, so that $\sigma_i = \nu_i$ for all $i = 1, \ldots, N$. Define $M = \max_{k=1,\ldots,K} M_k$ as the maximum component size.

In order to gain some useful intuition, we focus first on *complete K-partite graphs*, where all nodes are connected except those that belong to the same component. In other words, the complement of the graph consists of K fully connected components. Thus, transmission activity is mutually exclusive across the various components.

In this case, the normalization constant in (1) satisfies

$$Z = 1 + \sum_{k=1}^{K} \sum_{\emptyset \neq I \subseteq V_k} \prod_{i \in I} \sigma_i = 1 + \sum_{k=1}^{K} \Upsilon_k.$$

For any $k = 1, \ldots, K$, define $S_k = \{u \in \Omega : \sum_{i \in V_k} u_i \geq 1\}$ as the set of activity states where at least one of the nodes in V_k is active. We will use these sets to find a lower bound for the delay and mixing time. As discussed in Example 5.1, these sets are likely to provide a tight lower bound.

Lemma 6.1. *For any activation rate vector $(\nu_1, \ldots, \nu_N)$ such that the system is stable, for any $k = 1, \ldots, K$,*

$$Q(S_k) = Q(\Omega \setminus S_k) < M_k(1 - \sum_{l \neq k} \hat{\rho}_l) \left(\frac{1-\rho}{\hat{\rho}_k}\right)^{M_k - 1}, \tag{18}$$

$$\hat{\rho}_k < \pi(S_k) < 1 - \sum_{l \neq k} \hat{\rho}_l, \tag{19}$$

$$\sum_{l \neq k} \hat{\rho}_l < \pi(\Omega \setminus S_k) < 1 - \hat{\rho}_k. \tag{20}$$

Proof Using (1) we obtain

$$\theta_i = Z^{-1} \sum_{u \in \Omega, u_i = 1} \prod_{j=1}^{N} \sigma_j^{u_j} \geq Z^{-1} \sigma_i \prod_{l \in V_k \setminus \{i\}} (1 + \sigma_l) = Z^{-1} \frac{\sigma_i}{1 + \sigma_i} (\Upsilon_k + 1).$$

12

Also, from (1) we know
$$\theta_i = \sigma_i \mathbb{P}\{U_j = 0 \text{ for all } j \in \mathcal{N}_i^+\}.$$

Furthermore,
$$\mathbb{P}\{U_j = 0 \text{ for all } j \in \mathcal{N}_i^+\} \le \mathbb{P}\{U_j = 0 \text{ for all } j \in \mathcal{C}\},$$

and hence we get
$$\theta_i \le \sigma_i[1 - \sum_{j \in \mathcal{C}} \theta_j],$$

from (6), which may be rewritten as
$$\theta_i \le \frac{\sigma_i}{1 + \sigma_i}[1 - \sum_{j \in \mathcal{C}\backslash\{i\}} \theta_j],$$

so that
$$\Upsilon_k \le Z[1 - \sum_{j \in \mathcal{C}\backslash\{i\}} \theta_j] - 1,$$

and thus, using the fact that $\rho_i < \theta_i$ for all $i \in V$ is a necessary condition for stability,
$$\Upsilon_k < Z[1 - \sum_{j \in \mathcal{C}\backslash\{i\}} \rho_j] - 1. \tag{21}$$

Next, note that $Q(\Omega \backslash S_k) = \pi(0) \sum_{i \in V_k} \sigma_i$, and similarly,
$$Q(S_k) = \pi(0) \sum_{i \in V_k} \sigma_i = \frac{1}{Z} \sum_{i \in V_k} \sigma_i.$$

Using this we get,
$$Q(S_k) \le \frac{M_k}{Z} \max_{i \in V_k} \sigma_i = M_k \max_{i \in V_k} \frac{\frac{\sigma_i}{1+\sigma_i}(\Upsilon_k + 1)}{\frac{1}{1+\sigma_i}(\Upsilon_k + 1)} = \frac{M_k}{Z} \max_{i \in V_k} \frac{\frac{\sigma_i}{1+\sigma_i}(\Upsilon_k + 1)}{\prod_{l \in V_k\backslash\{i\}}(1 + \sigma_l)}$$
$$< \frac{M_k(\Upsilon_k + 1)}{Z(1 + \min_{i \in V_k} \sigma_i)^{M_k - 1}}.$$

Invoking Lemma 3.1 and (21) gives (18).

Also, because $\hat{\rho}_k < \pi(S_k) = \Upsilon_k/Z$ is needed for stability, (21) gives,
$$\pi(S_k) < 1 - \sum_{l \neq k} \hat{\rho}_l - \frac{1}{Z},$$

which proves (19). Noting that $\pi(S_k) + \pi(\Omega \backslash S_k) = 1$ gives (20). $\square$

Using Lemma 6.1 we can find a lower bound for the expected aggregate weighted queue length at some subset of nodes in $\mathcal{A} \subseteq V_k$.

Theorem 6.2. *For any activation rate vector $(\nu_1, \ldots, \nu_N)$ such that the system is stable and for any $w \in \mathbb{R}_+^N$, $\mathcal{A} \subseteq V_k$,*
$$\sum_{i \in \mathcal{A}} w_i \mathbb{E}\{L_i\} > \frac{1}{2M}(\rho_{\min})^{M+1} \sum_{i \in \mathcal{A}} w_i \lambda_i \left(\frac{1}{1 - \rho}\right)^{M-1}.$$

For the symmetric scenario $M_k \equiv M$ and $\hat{\rho}_k \equiv \rho/K$ for all $k = 1, \ldots, K$,

$$\mathbb{E}\{L_i\} > \frac{(K-1)^2 \rho^{M+2}}{2MK^{M+1}(K-(K-1)\rho)} \left(\frac{1}{1-\rho}\right)^{M-1}.$$

Proof The proof relies on applying Proposition 5.1, taking S to be (i) $\Omega \setminus S_k$ and (ii) S_l, $l \neq k$. In either case, $u_i = 0$ for all $i \in \mathcal{A}$, $u \in S$, so that $Y(w, \mathcal{A}, S) = 0$, i.e.,

$$D(w, \mathcal{A}, S) = \sum_{i \in \mathcal{A}} \lambda_i w_i.$$

First consider case (i). In this case we obtain the lower bound

$$\sum_{i \in \mathcal{A}} w_i \mathbb{E}\{L_i\} > \frac{\left(\sum_{l \neq k} \hat{\rho}_l\right)^2}{2M_k} \sum_{i \in \mathcal{A}} w_i \lambda_i \left(\frac{\hat{\rho}_k}{1-\rho}\right)^{M_k - 1}$$

from Proposition 5.1 and Lemma 6.1. Taking $\mathcal{A} = V_k$ yields the second statement of the lemma for a symmetric scenario.

In order to complete the proof of the first part of the lemma, we now turn to case (ii). Using Proposition 5.1 and Lemma 6.1, we arrive at the lower bound

$$\sum_{i \in \mathcal{A}} w_i \mathbb{E}\{L_i\} > \frac{\hat{\rho}_l^2}{2M_l} \sum_{i \in \mathcal{A}} w_i \lambda_i \left(\frac{\hat{\rho}_l}{1-\rho}\right)^{M_l - 1}.$$

Combining the above two lower bounds yields the first part of the lemma. $\square$

Theorem 6.2 states that in a complete K-partite graph the expected queue length grows at least as fast as $1/(1-\rho)^{M-1}$, with M the size of the largest component. Based on the observations in Section 3, this may be heuristically explained as follows. In order for the system to be stable, each node must at least have an activation rate of the order $1/(1-\rho)$, see Lemma 3.1. In turn, the transition times between the various activity states as governed by the maximum-size component occur on a time scale of the order ν^{M-1}, when each node has a fixed activation rate ν.

For $M = 1$ (full interference graph), the lower bound established in Theorem 6.2 is loose, reflecting that it is not the slow transitions between the various components that cause the delays to be long in that case, but the sheer load. For $M = 2$, the lower bound could also have been obtained by treating cliques as single-resource systems and is in fact similar to (2). For $M \geq 3$, the lower bound is particularly relevant, and reflects that the slow transitions between the various components cause the delays to be exponentially larger than can be explained from sheer load considerations alone.

Lemma 6.1 also provides a corresponding lower bound for the mixing time of the activity process $\{U(t)\}$ as established in the next theorem.

Theorem 6.3. *For any activation rate vector $(\nu_1, \ldots, \nu_N)$ such that the system is stable,*

$$t_{\mathrm{mix}}(\epsilon) > ((K-1)\rho_{\min} - 2\epsilon) \frac{(\rho_{\min})^M}{M} \left(\frac{1}{1-\rho}\right)^{M-1}.$$

Proof Applying Proposition 5.2 for $S = S_k$ and using Lemma 6.1 gives

$$t_{\text{mix}}(\epsilon) > (1 - 2\epsilon - (1 - \sum_{l \neq k} \hat{\rho}_l)) \frac{\hat{\rho}_k}{M_k(1 - \sum_{l \neq k} \hat{\rho}_l)} \left(\frac{\hat{\rho}_k}{1 - \rho}\right)^{M_k - 1}$$

$$> ((K - 1)\rho_{\min} - 2\epsilon) \frac{\rho_{\min}}{M_k} \left(\frac{\rho_{\min}}{1 - \rho}\right)^{M_k - 1},$$

for $k = 1, \ldots, K$, and the result follows. $\qquad\square$

Note that the bound derived in Theorem 6.3 is not necessarily the tightest bound we can find using the results of Lemma 6.1. In fact, the bound is irrelevant if $\epsilon > (K - 1)\rho_{\min}/2$. However, if $\rho_{\min} > 0$ and ϵ is small enough, we can conclude that the mixing time grows at least like $1/(1-\rho)^{M-1}$ as ρ increases to 1.

For equal activation rates, i.e., for the activation rate vector $(\nu, \ldots, \nu)$, it is shown in [31] that $t_{\text{mix}}(\epsilon) \sim \nu^{M^*-1}$ as $\nu \to \infty$, with M^* the size of the second largest component, so that the heavy-traffic behavior is governed by M^* instead of M. Note however that this activation rate vector does not provide a stable system, unless $\rho_{\min} = 0$ or $M = M^*$, as follows from Lemma 6.1 and the fact that $\pi(S_k) \to 0$ as $\nu \to \infty$ for all k such that $M_k < M$. Hence the activity process mixes slower in heavy traffic if the system is stable as compared to a system with equal activation rates.

7 Extensions

In this section we turn attention to the broader class of (not necessarily complete) K-partite graphs. Thus, transmission activity is no longer mutually exclusive across the various components. However, we make the next assumption implying that joint activity across various components is relatively inefficient. Denote by $v^{(k)} = 1_{V_k}$ the incidence vector of V_k, i.e., $v_i^{(k)} = 1$ if $i \in V_k$ and $v_i^{(k)} = 0$ otherwise, and define $\Omega^* = \{v^{(1)}, \ldots, v^{(K)}\}$.

Assumption 7.1. For any $u \in \Omega$,

$$H(u) = \sum_{k=1}^{K} \sum_{i \in V_k} \frac{u_i}{M_k} \leq 1,$$

with strict inequality for any $u \notin \Omega^*$.

Based on the above assumption, we define

$$\zeta = 1 - \max_{u \in \Omega \setminus \Omega^*} H(u) > 0.$$

An illustrative example is provided by a $2B \times 2B$ grid with nodes labeled as $\{(i,j)\}$, $i, j = 1, \ldots, 2B$, and nearest-neighbor interference. The two components are $V_1 = \{(i,j) : (i+j) \bmod 2 = 1\}$ and $V_2 = \{(i,j) : (i+j) \bmod 2 = 0\}$, with $M_1 = M_2 = 2B^2$. In order for $m \geq 1$ nodes in V_1 to be active, at least $m + 1$ or $m + 3$ nodes in V_2 must be inactive (depending on whether or not we assume a wrap-around boundary). Thus $\sum_{i=1}^{2B} \sum_{j=1}^{2B} u_{(i,j)} \leq 2B^2 - 1$ (or $2B^2 - 3$) for all $u \in \Omega \setminus \Omega^*$, and $\zeta = \frac{1}{2B^2}$ (or $\frac{3}{2B^2}$).

The next lemma shows that in order for the system to be stable, joint activity across the various components can only occur a negligible fraction of the time at high load.

Lemma 7.1. *In order for the system to be stable, it should hold that*

$$\sum_{u \in \Omega \setminus \Omega^*} \pi(u) < \frac{1-\rho}{\zeta},$$

and

$$\pi(v^{(k)}) > \hat{\rho}_k - \frac{1-\rho}{\zeta}.$$

Proof In order for the system to be stable, we must have $\rho_i < \theta_i$ for all $i = 1, \ldots, N$. Thus,

$$
\begin{aligned}
\rho &= \sum_{k=1}^{K} \hat{\rho}_k = \sum_{k=1}^{K} \frac{1}{M_k} \sum_{i \in V_k} \rho_i \\
&< \sum_{k=1}^{K} \frac{1}{M_k} \sum_{i \in V_k} \sum_{u \in \Omega} \pi(u) u_i = \sum_{u \in \Omega} \pi(u) \sum_{k=1}^{K} \frac{1}{M_k} \sum_{i \in V_k} u_i \\
&= \sum_{u \in \Omega^*} \pi(u) \sum_{k=1}^{K} \frac{1}{M_k} \sum_{i \in V_k} u_i + \sum_{u \in \Omega \setminus \Omega^*} \pi(u) \sum_{k=1}^{K} \frac{1}{M_k} \sum_{i \in V_k} u_i \\
&\le \sum_{u \in \Omega^*} \pi(u) + (1 - \zeta) \sum_{u \in \Omega \setminus \Omega^*} \pi(u) \\
&= 1 - \zeta \sum_{u \in \Omega \setminus \Omega^*} \pi(u),
\end{aligned}
$$

where the last inequality follows from Assumption 7.1. The first part of the lemma follows.

Also, for any $i \in V_k$,

$$\hat{\rho}_k = \rho_i < \sum_{u \in \Omega} \pi(u) u_i \le \pi(v^{(k)}) + \sum_{u \in \Omega \setminus \Omega^*} \pi(u),$$

which combined with the first statement yields the second part of the lemma. $\square$

In the next lemma we show that the fraction of time the activity process $\{U(t)\}$ spends in any component V_k relative to the traffic intensity of the nodes in that component, is almost equal for all components if ρ is large enough.

Lemma 7.2. *Assume the system is stable and $\rho \ge \rho_\gamma = 1 - \gamma \zeta \rho_{\min}^2$, $\gamma > 0$. Then*

$$\min_{k=1,\ldots,K} \frac{1}{\hat{\rho}_k} \prod_{i \in V_k} \sigma_i \ge (1 - 3\gamma) \max_{k=1,\ldots,K} \frac{1}{\hat{\rho}_k} \prod_{i \in V_k} \sigma_i.$$

Proof For compactness, we denote $\Pi_k = \prod_{i \in V_k} \sigma_i$ and $R_k = \Pi_k / \hat{\rho}_k$, and define $k_{\min} = \arg\min_{k=1,\ldots,K} R_k$ and $k_{\max} = \arg\max_{k=1,\ldots,K} R_k$.

Lemma 7.1 implies

$$
\begin{aligned}
\hat{\rho}_{k_{\min}} - \gamma \rho_{\min}^2 &\le \hat{\rho}_{k_{\min}} - (1-\rho)/\zeta \le \pi(v^{(k_{\min})}) = Z^{-1} \Pi_{k_{\min}} \\
&\le \frac{\Pi_{k_{\min}}}{\sum_{k=1}^{K} \Pi_k} = \frac{\Pi_{k_{\min}}}{\Pi_{k_{\min}} + \Pi_{k_{\max}} + \sum_{k \ne k_{\min}, k_{\max}} \Pi_k} \\
&\le \frac{\Pi_{k_{\min}}}{\Pi_{k_{\min}} + \Pi_{k_{\max}} + \sum_{k \ne k_{\min}, k_{\max}} \frac{\hat{\rho}_k}{\hat{\rho}_{k_{\min}}} \Pi_{k_{\min}}},
\end{aligned}
$$

16

yielding

$$(1 - (\hat{\rho}_{k_{\min}} + \sum_{k \neq k_{\min}, k_{\max}} \hat{\rho}_k)(1 - \frac{\gamma \rho_{\min}^2}{\hat{\rho}_{k_{\min}}}))\frac{\Pi_{k_{\min}}}{\Pi_{k_{\max}}} \geq \hat{\rho}_{k_{\min}} - \gamma \rho_{\min}^2,$$

or equivalently,

$$(1 - (\rho - \hat{\rho}_{k_{\max}})(1 - \gamma \rho_{\min}^2/\hat{\rho}_{k_{\min}}))\Pi_{k_{\min}} \geq (\hat{\rho}_{k_{\min}} - \gamma \rho_{\min}^2)\Pi_{k_{\max}}.$$

Using $\rho \geq 1 - \gamma \zeta \rho_{\min}^2$ and $\rho_{\min} \leq \hat{\rho}_{k_{\min}}$, it follows that

$$(1 - (1 - \gamma \zeta \rho_{\min}^2 - \hat{\rho}_{k_{\max}})(1 - \gamma \rho_{\min}))\Pi_{k_{\min}} \geq (1 - \gamma)\hat{\rho}_{k_{\min}}\Pi_{k_{\max}}.$$

This yields

$$(1 + 2\gamma)\hat{\rho}_{k_{\max}}\Pi_{k_{\min}} \geq (1 - \gamma)\hat{\rho}_{k_{\min}}\Pi_{k_{\max}},$$

and thus,

$$R_{k_{\min}} \geq \frac{(1 - \gamma)R_{k_{\max}}}{1 + 2\gamma} \geq (1 - 3\gamma)R_{k_{\max}}.$$

$\square$

In order to state a lower bound for the expected aggregate weighted queue length at some subset of nodes $\mathcal{A} \subseteq V_k$, we now first introduce some further notation and concepts.

A sequence of states $(u^{(0)}, u^{(1)}, \ldots, u^{(l)})$, with $u^{(k)} \in \Omega$, $k = 0, \ldots, l$, is called a *path* from $u^{(0)}$ to $u^{(l)}$ if $(u^{(k)}, u^{(k+1)})$ are feasible transitions, i.e., $q(u^{(k)}, u^{(k+1)}) > 0$ for all $k = 0, \ldots, l - 1$. For a given path $p = (u^{(0)}, u^{(1)}, \ldots, u^{(l)})$, denote by $m(p) = \min_{k=0,1,\ldots,l} H(u^{(k)})$ the minimum value of the function $H(\cdot)$, as defined in Assumption 7.1, along the path. For given states $u, v \in \Omega$, denote by $P(u, v)$ the collection of all paths from u to v. Define $M(u, v) = \max_{p \in P(u,v)} m(p)$ as the maximum of the minimum value of the function $H(\cdot)$ along any path from state u to state v, with the convention that $M(u, u) = \infty$.

For all $\mathcal{A} \subseteq V$ such that $\mathcal{A} \subseteq V_k$ for some $k \in \{1, \ldots, K\}$, denote by $\Delta(\mathcal{A})$ the set of states in which the expected drift of the aggregate weighted queue length in $\mathcal{A}$ is non-positive, i.e.,

$$\Delta(\mathcal{A}) = \{u \in \Omega : \sum_{i \in \mathcal{A}} w_i \lambda_i \leq \sum_{i \in \mathcal{A}} w_i \mu_i u_i\}.$$

Further define $\delta(\mathcal{A})$ as the minimal expected drift of the aggregate weighted queue length in $\mathcal{A}$ if the system does not reside in of one of the states in $\Delta(\mathcal{A})$, i.e.,

$$\delta(\mathcal{A}) = \sum_{i \in \mathcal{A}} w_i \lambda_i - \max_{u \in \Omega \setminus \Delta(\mathcal{A})} \sum_{i \in \mathcal{A}} w_i \mu_i u_i.$$

Note that $\delta(\mathcal{A}) > 0$ by construction. For all $l \neq k$, define $m_l(\mathcal{A}) = \max_{u \in \Delta(\mathcal{A})} M(v^{(l)}, u)$, and

$$S_l(\mathcal{A}) = \{u \in \Omega : M(v^{(l)}, u) > m_l(\mathcal{A})\}$$

as the set of states that can be reached from $v^{(l)}$ via a path p with $m(p) > m_l(\mathcal{A})$. Also, define $m_k(\mathcal{A}) = \max_{l \neq k} m_l(\mathcal{A})$, and

$$S_k(\mathcal{A}) = \{v \in \Omega : \max_{u \in \Delta(\mathcal{A})} M(u, v) > m_k(\mathcal{A})\}$$

17

as the set of states that can be reached from $\Delta(\mathcal{A})$ via a path p with $m(p) > m_k(\mathcal{A})$.

Finally, define $H_l(\mathcal{A}) = \max_{u \in \partial S_l(\mathcal{A})} H(u)$, $H^*(\mathcal{A}) = \min_{l=1,\ldots,K} H_l(\mathcal{A})$ and $H^*_{\min} = \min_{\mathcal{A} \subseteq V : \exists k : \mathcal{A} \subseteq V_k} H^*(\mathcal{A})$.

In the remainder of this subsection we will assume that the activation rates of nodes in the same component are equal, i.e. $\sigma_i = \hat{\sigma}_k$ if $i \in V_k$ for all $k = 1, \ldots, K$. Denote $\sigma^* = \min_{k=1,\ldots,K} \hat{\sigma}_k^{M_k}$ and $k^* = \operatorname{argmin}_{k=1,\ldots,K} \hat{\sigma}_k^{M_k}$.

Remark. It is not clear when there exists an activation rate vector $(\nu_1, \ldots, \nu_N)$ with $\nu_i = \nu_j$ if $i, j \in V_k$ that stabilizes the system. For symmetric topologies, e.g. ring networks with an even number of nodes or tori with an even number of nodes in both directions, it seems plausible that such an activation rate vector can stabilize the system for any $\rho < 1$. For asymmetric typologies, e.g. linear topologies and two-dimensional grid networks, this is not clear.

In the next lemma we derive an upper bound for the fraction of the time the system spends in the boundary of $S_l(\mathcal{A})$ for any $l = 1, \ldots, K$.

Lemma 7.3. *Assume the system is stable and $\rho \geq \rho_\gamma = 1 - \gamma \zeta \rho_{\min}^2$, $\gamma > 0$. Then*

$$\max_{u \in \partial S_l(\mathcal{A})} \prod_{j=1}^N \sigma_j^{u_j} \leq \left(\frac{\sigma^*}{(1 - 3\gamma)\rho_{\min}} \right)^{H_l(\mathcal{A})}.$$

Proof Since $\sigma_i = \hat{\sigma}_k$ for all $k = 1, \ldots, K$, we obtain

$$\max_{u \in \partial S_l(\mathcal{A})} \prod_{j=1}^N \sigma_j^{u_j} = \max_{u \in \partial S_l(\mathcal{A})} \prod_{k=1}^K \prod_{i \in V_k} \sigma_i^{u_i} = \max_{u \in \partial S_l(\mathcal{A})} \prod_{k=1}^K \hat{\sigma}_k^{\sum_{i \in V_k} u_i}$$

$$= \max_{u \in \partial S_l(\mathcal{A})} \prod_{k=1}^K (\hat{\sigma}_k^{M_k})^{\frac{1}{M_k} \sum_{i \in V_k} u_i}.$$

Lemma 7.2 gives

$$\hat{\sigma}_k^{M_k} \leq \frac{\hat{\rho}_k}{1 - 3\gamma} \min_{l=1,\ldots,K} \frac{1}{\hat{\rho}_l} \hat{\sigma}_l^{M_l} \leq \frac{\sigma^*}{(1 - 3\gamma)\rho_{\min}},$$

and thus,

$$\max_{u \in \partial S_l(\mathcal{A})} \prod_{j=1}^N \sigma_j^{u_j} \leq \max_{u \in \partial S_l(\mathcal{A})} \prod_{k=1}^K \left(\frac{\sigma^*}{(1 - 3\gamma)\rho_{\min}} \right)^{\frac{1}{M_k} \sum_{i \in V_k} u_i}$$

$$\leq \max_{u \in \partial S_l(\mathcal{A})} \left(\frac{\sigma^*}{(1 - 3\gamma)\rho_{\min}} \right)^{\sum_{k=1}^K \frac{1}{M_k} \sum_{i \in V_k} u_i}$$

$$= \left(\frac{\sigma^*}{(1 - 3\gamma)\rho_{\min}} \right)^{\max_{u \in \partial S_l(\mathcal{A})} H(u)}$$

$$= \left(\frac{\sigma^*}{(1 - 3\gamma)\rho_{\min}} \right)^{H_l(\mathcal{A})}.$$

$\square$

We are now in the position to derive bounds for $Q(S_l(\mathcal{A}))$, $\pi(S_l(\mathcal{A}))$ and $\pi(\Omega \setminus S_l(\mathcal{A}))$ that are qualitatively similar to the bounds in Lemma 6.1.

Lemma 7.4. *Assume $\rho \geq \rho_\gamma = 1 - \gamma\zeta\rho_{\min}^2$, $\gamma > 0$. For any activation rate vector $(\nu_1, \ldots, \nu_N)$ such that the system is stable and with $\nu_i = \nu_j$ if $i, j \in V_k$ for some k, for any $l = 1, \ldots, K$,*

$$Q(S_l(\mathcal{A})) = Q(\Omega \setminus S_l(\mathcal{A}))$$

$$< \frac{2^N}{(1 - 3\gamma)\rho_{\min}} \left(\frac{\rho_{k^*}}{1 - \rho}\right)^{M(H_l(\mathcal{A}) - 1)}, \tag{22}$$

$$(1 - \gamma)\rho_{\min} < \pi(S_l(\mathcal{A})) < 1 - (1 - \gamma)\rho_{\min}, \tag{23}$$

$$(1 - \gamma)\rho_{\min} < \pi(\Omega \setminus S_l(\mathcal{A})) < 1 - (1 - \gamma)\rho_{\min}. \tag{24}$$

Proof First note that

$$Q(S_l(\mathcal{A})) = \sum_{u \in \partial S_l(\mathcal{A})} \pi(u)$$

and

$$Q(\Omega \setminus S_l(\mathcal{A})) = \sum_{u \in \partial S_l(\mathcal{A})} \pi(u).$$

Further,

$$Q(S_l(\mathcal{A})) = \sum_{u \in \partial S_l(\mathcal{A})} \pi(u) = Z^{-1} \sum_{u \in \partial S_l(\mathcal{A})} \prod_{j=1}^N \sigma_j^{u_j}$$

$$\leq Z^{-1} |\partial S_l(\mathcal{A})| \max_{u \in \partial S_l(\mathcal{A})} \prod_{j=1}^N \sigma_j^{u_j}.$$

Noting that $Z \geq \sigma^*$ and $|\partial S_l(\mathcal{A})| \leq 2^N$ yields, using Lemma 7.3,

$$Q(S_l(\mathcal{A})) \leq \frac{2^N \left(\frac{\sigma^*}{(1 - 3\gamma)\rho_{\min}}\right)^{H_l(\mathcal{A})}}{\sigma^*} \leq \frac{2^N (\sigma^*)^{H_l(\mathcal{A}) - 1}}{(1 - 3\gamma)\rho_{\min}},$$

and (22) follows from Lemma 3.1.

Further, using Lemma 7.1,

$$\pi(S_l(\mathcal{A})) = \pi(S_l(\mathcal{A})) \geq \pi(v^{(l)}) > \hat{\rho}_l - \frac{1 - \rho}{\zeta} \geq (1 - \gamma)\rho_{\min}.$$

Now note that by definition $S_l \cap \Delta(\mathcal{A}) = \emptyset$ for $l \neq k$ and $\Delta(\mathcal{A}) \subseteq S_k$ if $\mathcal{A} \subseteq V_k$. Hence, for $l \neq k$,

$$\pi(S_l(\mathcal{A})) \leq 1 - \pi(v^{(k)}) < 1 - (1 - \gamma)\rho_{\min},$$

and

$$\pi(S_k(\mathcal{A})) \leq 1 - \sum_{l \neq k} \pi(v^{(l)}) < 1 - (K - 1)(1 - \gamma)\rho_{\min},$$

which gives (23). Noting that $\pi(S_l(\mathcal{A})) + \pi(\Omega \setminus S_l(\mathcal{A})) = 1$ gives (24). □

Using a similar approach as in Section 6, the bounds in Lemma 7.4 can be utilized to establish a lower bound for the expected aggregate weighted queue length in some subset of nodes and for the mixing time of the activity process.

Theorem 7.5. *Assume $\rho \geq \rho_\gamma = 1 - \gamma \zeta \rho_{\min}^2$, $\gamma > 0$. For any activation rate vector $(\nu_1, \ldots, \nu_N)$, with $\nu_i = \nu_j$ if $i, j \in V_k$ for some k, such that the system is stable and for any $w \in \mathbb{R}_+^N$, $\mathcal{A} \subseteq V_k$,*

$$\sum_{i \in \mathcal{A}} w_i \mathbb{E}\{L_i\} > \frac{\delta(\mathcal{A})(1 - 4\gamma)\rho_{\min}^{M+3}}{2^{N+1}} \left(\frac{1}{1-\rho}\right)^{M(1-H^*(\mathcal{A}))},$$

Proof The proof of this theorem proceeds along similar lines as the proof of Theorem 6.2 and relies on applying Proposition 5.1, taking S to be (i) $\Omega \setminus S_k(\mathcal{A})$ and (ii) $S = S_l(\mathcal{A})$, $l \neq k$. First note that by definition $S_l(\mathcal{A}) \cap \Delta(\mathcal{A}) = \emptyset$, and thus $D(w, \mathcal{A}, S_l(\mathcal{A})) \geq \delta(\mathcal{A})$, $l \neq k$. Also note that $\Delta(\mathcal{A}) \subseteq S_k(\mathcal{A})$, so that $D(w, \mathcal{A}, \Omega \setminus S_k(\mathcal{A})) \geq \delta(\mathcal{A})$.

Further, using Lemma 7.4 we obtain the lower bound

$$\sum_{i \in \mathcal{A}} w_i \mathbb{E}\{L_i\} > \frac{\delta(\mathcal{A})(1 - 4\gamma)\rho_{\min}^{M+3}}{2^{N+1}} \left(\frac{1}{1-\rho}\right)^{M(1-H_l(\mathcal{A}))},$$

for $l = 1, \ldots, K$, and the result follows. $\qquad\square$

Theorem 7.5 states that in a general K-partite interference graph the expected queue length grows at least as fast as $1/(1-\rho)^{M(1-H^*)}$, where the coefficient H^* depends on the specific topology and is in general hard to calculate. We however know that $\frac{1}{M} \leq H^* \leq 1$ and for some specific topologies we can explicitly determine H^*.

The next theorem provides a corresponding lower bound for the mixing time of the activity process $\{U(t)\}$.

Theorem 7.6. *Assume $\rho \geq \rho_\gamma = 1 - \gamma \zeta \rho_{\min}^2$, $\gamma > 0$. For any activation rate vector $(\nu_1, \ldots, \nu_N)$ such that the system is stable and with $\nu_i = \nu_j$ if $i, j \in V_k$ for some k,*

$$t_{\mathrm{mix}}(\epsilon) > ((1-\gamma)\rho_{\min} - 2\epsilon)\frac{(1-4\gamma)\rho_{\min}^{M+2}}{2^N} \left(\frac{1}{1-\rho}\right)^{M(1-H_{\min}^*)}.$$

Proof Take $\mathcal{A} \subseteq V$ such that there exists a k such that $\mathcal{A} \subseteq V_k$. Using Lemma 6.1 we then find for any $l \in \{1, \ldots, K\}$,

$$\frac{1}{\Phi(S_l(\mathcal{A}))} > \frac{(1-\gamma)\rho_{\min}(1-3\gamma)\rho_{\min}}{2^N} \left(\frac{\rho_{k^*}}{1-\rho}\right)^{M(1-H_l(\mathcal{A}))}$$

$$\geq \frac{(1-4\gamma)\rho_{\min}^{M+2}}{2^N} \left(\frac{1}{1-\rho}\right)^{M(1-H_l(\mathcal{A}))}.$$

Hence, using Proposition 5.2,

$$t_{\mathrm{mix}}(\epsilon) > ((1-\gamma)\rho_{\min} - 2\epsilon)\frac{(1-4\gamma)\rho_{\min}^{M+2}}{2^N} \left(\frac{1}{1-\rho}\right)^{M(1-H_l(\mathcal{A}))},$$

and the result follows by optimizing over l and $\mathcal{A}$. $\qquad\square$

The value of the coefficient $H^*(\mathcal{A})$ depends strongly on the specific properties of the interference graph G. For a complete partite graph, the sets $S_l(\mathcal{A})$ coincide with those in the previous subsection, and we have $\partial S_l(\mathcal{A}) = \bigcup_{i \in V_l}\{e_i\}$, so that $H_l(\mathcal{A}) = 1/M_l$, and $H^*(\mathcal{A}) = 1/M$, recovering the result of Theorem 6.2. On the other hand, when the graph consists of N/K fully connected
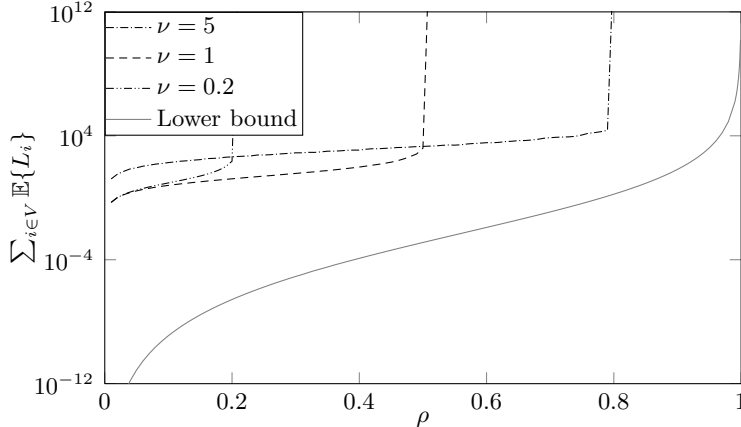
20

Figure 1: Average total number of packets for several fixed activation rates.

components, we have $H_l(\mathcal{A}) \equiv 1$, and the result trivializes. An interesting intermediate situation is the $2B \times 2B$ grid mentioned earlier with $M = M_1 = M_2 = 2B^2$, for which we conjecture that $H^*(\mathcal{A}) = H_1 = H_2 = 1 - 1/B$ or $1 - 1/(2B)$ if $B \geq 2$, depending on whether or not we assume a wrap-around boundary, suggesting that the mean queue lengths would grow as $1/(1-\rho)^B$ or $1/(1-\rho)^{2B}$.

8 Simulation experiments

In this section we will illustrate the theoretical results for the growth behavior of the aggregate queue length through simulation experiments. For cross comparison, we consider a system that can be represented by a symmetric complete bipartite ($K = 2$) interference graph with components of size $M = 5$. Because of space considerations, we do not report simulation results for other cases, but we observed qualitatively similar behavior in a broad range of scenarios.

To estimate the expected aggregate queue length for a given value of ρ, we set $t = 10^6$ and calculate the average total number of packets in the time intervals $[0, t]$ and $[t+1, 2t]$, starting from an initially empty system. We take the average of the two values to be our estimate if the values are less than 5% apart. Otherwise we set $t = 2t$ and repeat the procedure.

Figure 1 shows the average total number of packets in the system for various fixed activation rates. Note that we used a log-lin scale. We see that the simulated curves lie well above the lower bound of Theorem 6.2 for all chosen values of ν. Note that the system is not stable for all values of ρ, e.g. for $\nu = 1$ the system is unstable if $\rho \geq 2\theta_i = \frac{32}{63}$, explaining the jumps in the simulation result. Further note that the expected time between activation of nodes in the two components is smaller for small values of ν. This explains why small values of ν tend to perform better in case ρ is small, i.e., for large values of ν the nodes in one component will often be transmitting dummy packets while the nodes in the other component do have packets waiting to be transmitted.

Figure 2 shows the average total number of packets in the system for $f(l) = l$ and $g(l) = 1$. We see that the lower bound of Theorem 4.1 is remarkably close to the simulation result for small values of ρ. For larger values of ρ the bound and simulation result are farther apart. One explanation for this lies in the approximation made in (5). For small values of ρ this approximation is relatively
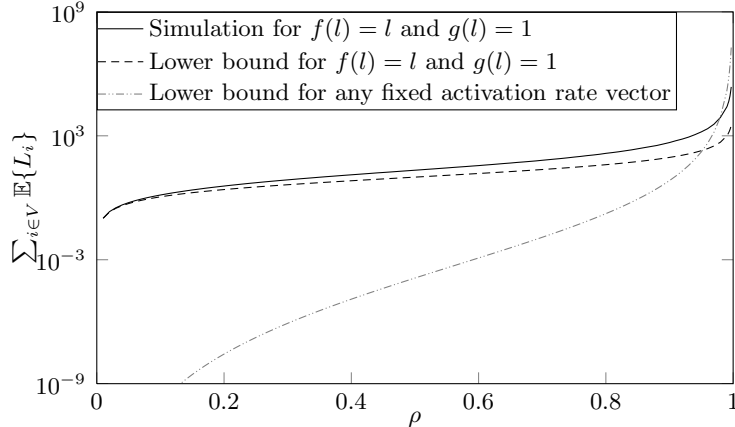
21

Figure 2: Average total number of packets for $f(l) = l$ and $g(l) = 1$.

good while for large values of ρ this approximation is off by a factor of about 2 in this case. While this does not explain the total discrepancy in this case, it does explain all discrepancies in case the rate of increase of the activation function is slow, e.g. $f(l) = \log(l + 1)$.

Finally note that the simulation result lies, for large values of ρ, below the lower bound for fixed activation rates established in Theorem 6.2. This suggests that the activation function $f_i(l) \equiv f(l) = l$ performs better in heavy traffic than $f_i(l) = \nu_i$ for any choice of the activation rate vector $(\nu_1, \ldots, \nu_N)$.

9   Conclusions

We have established lower bounds for the expected queue lengths and delays in wireless random-access networks. Both for queue-based strategies and fixed activation rates, the derivation of the bounds starts from the observation that stability of the system requires the activity factors to be big at high load. The specific subsequent arguments considerably differ however in both cases. Queue-based strategies for which maximum stability has been established, involve slow, logarithmic, activation functions, which require huge queue lengths at every node for the activity factors to be big enough, and cause the exponential delay scaling. In contrast, the delays for fixed activation rates are shown to result from excessive mixing times due to a bottleneck in the network topology together with the big activity factors required for stability. We also observe that the network topology plays a major role in case of fixed activation rates, while it only appears to matter somewhat implicitly in case of queue-based strategies as will be further discussed below.

For complete partite interference graphs, a comparison of both cases reveals that the expected delay for queue-based strategies grows faster than the lower bound $1/(1-\rho)^{M-1}$ for fixed activation and de-activation rates when $h(l)$ increases slower than $l^{1/(M-1)}$, with M denoting the maximum component size. This is for example the case if $f(l) = r(l)/(1 + r(l))$ and $g(l) = 1/(1 + r(l))$, with $r(l) = \log(l + 1)$. Conversely, when $h(l)$ increases faster than $l^{1/(M-1)}$, the lower bound for fixed activation and de-activation rates could potentially be beaten by sufficiently aggressive queue-based strategies. Simulation experiments demonstrate that the actual expected delays indeed exhibit the cross-over suggested by the lower bounds.

A challenging issue for further research is to examine whether more aggressive queue-based strategies can improve the delay performance in more general topologies as well. As noted earlier, maximum-stability guarantees in arbitrary topologies have only been established so far for nominal activation functions that grow logarithmically with the queue lengths [22, 24, 25]. Inspection of the proof arguments indicates that maximum stability will remain guaranteed as long as the fluid limits of the queue length process exhibit fast mixing behavior. This in turn means that the activity process for such queue-based strategies in fact behaves as if the activation rates are essentially fixed. Thus, in arbitrary topologies it is questionable whether queue-based strategies have the capability to outperform fixed-rate strategies.

In some specific topologies, however, maximum stability is maintained for highly aggressive queue-based strategies for which the fluid limits of the queue length process may exhibit slow mixing behavior [8, 11]. The complete partite interference graphs considered in the present paper are crucial examples of such topologies. In these scenarios there seems to be scope for more aggressive queue-based strategies to reduce the delays, as confirmed by the lower bounds and simulation results that we presented.

In conclusion, the question in what kind of scenarios more aggressive queue-based strategies can improve the delay performance appears to be inextricably linked to the question under what conditions such strategies provide maximum-stability guarantees. In both these questions, the mixing properties of the activity process seem to play a central role, and it would be interesting to explore this three-way connection further.

References

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

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

[3] N. Bouman, S.C. Borst, O.J. Boxma, and J.S.H. van Leeuwaarden. Queues with random back-offs. *Preprint*, 2012.

[4] N. Bouman, S.C. Borst, and J.S.H. van Leeuwaarden. Achievable delay performance in CSMA networks. In *Proc. 49th Allerton Conf.*, 2011.

[5] N. Bouman, S.C. Borst, J.S.H. van Leeuwaarden, and A. Proutière. Backlog-based random access in wireless networks: fluid limits and delay issues. In *Proc. ITC 23*, 2011.

[6] M. Durvy, O. Dousse, and P. Thiran. Modeling the 802.11 protocol under different capture and sensing capabilities. In *Proc. Infocom 2007 Conf.*, 2007.

[7] M. Durvy and P. Thiran. A packing approach to compare slotted and non-slotted medium access control. In *Proc. Infocom 2006 Conf.*, 2006.

[8] M. Feuillet, A. Proutière, and P. Robert. Random capture algorithms: Fluid limits and stability. In *Proc. ITA Workshop*, 2010.

[9] S.W. Fuhrmann and R.B. Cooper. Stochastic decompositions for the M/G/1 queue with generalized vacations. *Oper. Res.*, 33:1117–1129, 1985.

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

[11] J. Ghaderi, S.C. Borst, and P.A. Whiting. Backlog-based random-access algorithms: fluid limits and stability issues. In *Proc. WiOpt 2012 Conf.*, 2012.

[12] J. Ghaderi and R. Srikant. On the design of efficient CSMA algorithms for wireless networks. In *Proc. CDC 2010 Conf.*, 2010.

[13] L. Jiang, M. Leconte, J. Ni, R. Srikant, and J. Walrand. Fast mixing of parallel glauber dynamics and low-delay CSMA scheduling. In *Proc. Infocom 2011 Mini-Conf.*, 2011.

[14] L. Jiang, D. Shah, J. Shin, and J. Walrand. Distributed random access algorithm: scheduling and congestion control. *IEEE Trans. Inf. Theory*, 56(12):6182–6207, 2010.

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

[16] L. Jiang and J. Walrand. Stability and delay of distributed scheduling algorithms for networks of conflicting queues. *Queueing Syst.*, 72(1):161–187, 2012.

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

[18] J. Liu, Y. Yi, A. Proutière, M. Chiang, and H.V. Poor. Towards utility-optimal random access without message passing. *Wireless Commun. Mob. Comput.*, 10(1):115–128, 2010.

[19] M. Lotfinezhad and P. Marbach. Throughput-optimal random access with order-optimal delay. In *Proc. Infocom 2011 Conf.*, 2011.

[20] P. Marbach and A. Eryilmaz. A backlog-based CSMA mechanism to achieve fairness and throughput-optimality in multihop wireless networks. In *Proc. Allerton 2008 Conf.*, 2008.

[21] J. Ni, B. Tan, and R. Srikant. Q-CSMA: queue length based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks. In *Proc. Infocom 2010 Mini-Conf.*, 2010.

[22] S. Rajagopalan, D. Shah, and J. Shin. Network adiabatic theorem: an efficient randomized protocol for contention resolution. In *Proc. ACM SIGMETRICS/Performance 2009 Conf.*, 2009.

[23] D. Shah and J. Shin. Delay-optimal queue-based CSMA. In *Proc. ACM SIGMETRICS 2010 Conf.*, 2010.

[24] D. Shah and J. Shin. Randomized scheduling algorithm for queueing networks. *Ann. Appl. Prob.*, 22:128–171, 2012.

[25] D. Shah, J. Shin, and P. Tetali. Medium access using queues. In *Proc. FOCS 2011 Conf.*, 2011.

[26] D. Shah, D.N.C. Tse, and J.N. Tsitsiklis. Hardness of low delay network scheduling. *IEEE Trans. Inf. Theory*, 57(12):7810–7817, 2011.

[27] V. Subramanian and M. Alanyali. Delay performance of CSMA in networks with bounded degree conflict graphs. In *Proc. ISIT 2011 Conf.*, 2011.

[28] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Aut. Contr.*, 37:1936–1948, 1992.

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

[30] X. Wang and K. Kar. Throughput modeling and fairness issues in CSMA/CA based ad-hoc networks. In *Proc. Infocom 2005 Conf.*, 2005.

[31] A. Zocca, S.C. Borst, and J.S.H. van Leeuwaarden. Mixing properties of CSMA networks on partite graphs. In *Proc. ValueTools 2012 Conf.*, 2012.