

Delay performance of CSMA in networks with bounded degree conflict graphs

Vijay G. Subramanian
EECS Dept., Northwestern University

Murat Alanyali
ECE Dept., Boston University

Abstract—We analyze packet delay in CSMA-based random access schemes in networks under the protocol interference model. Using a stochastic coupling argument we identify a subset of the throughput-region where queue lengths can be bounded uniformly for all network sizes. This conclusion provides a throughput-region of interest for delay sensitive applications and suggests that delay bounds based on mixing time analyses may be loose.

I. INTRODUCTION

It was recently shown that CSMA-based random access schemes can achieve the full throughput region under the protocol model of interference [1], [2]. Here a collection of n wireless links are considered under a conflict graph G which stipulates that no two neighbors in G may transmit simultaneously. In turn transmitting links form an independent set of the conflict graph, and the full throughput region is the convex hull of all such independent sets. The papers [1], [2] consider the transmission patterns that arise under CSMA using Markovian models and establish that the persistence of links in accessing the medium can be tuned to achieve any point in the interior of the throughput region.

It is also widely recognized that the throughput optimality of CSMA comes at the expense of short-term unfairness. Namely, instantaneous transmission patterns may tend to spend excessive time around maximal independent sets and thereby starving remaining links. This effect is typically more pronounced for larger throughputs and for larger systems. It appears fundamentally related to the mixing time of the Markov process of transmission patterns. In fact by upper-bounding the mixing time Jiang et al. [3] identify a subset of the throughput region in which mean queue lengths in a related queuing system grow as $O(\log n)$ where n is the number of links, or the number of nodes in G . Interestingly, there is also a matching lower-bound for the mixing time in general graphs [4]. Hence if the mixing time is an indicator of dependence between short-term fairness and the system size n , then one would expect that the dependence would always be present, though possibly weak for some throughputs and strong for others.

In this paper we establish that queue lengths under models akin to that of [3] are uniformly bounded in an explicitly characterized subset of the throughput region. This conclusion

has two implications. Firstly it identifies a throughput region that would be of interest for delay sensitive applications. Secondly it suggests that mixing time may not be a suitable instrument for sharp characterization of such regions.

The alluded bound on queue lengths relies on a stochastic coupling argument. Specifically, we stochastically bound each individual queue length with an auxiliary process obtained by omitting all other links except the neighbors. The auxiliary process provides a lower bound on throughput and an upper bound on the queue length, where the upper bound depends only on local contention faced by each link.

Sections II and III establish the coupling and resulting throughput bounds in discrete and continuous models, respectively. Some numerical results are provided in Section IV and the paper concludes with final remarks in Section V.

II. DISCRETE-TIME CSMA NETWORK

Following the notation of [2], [3] we represent the independent sets of G with \mathcal{I} and denote by \mathcal{N}_i the neighbors of node i in G . Since each node of G represents a wireless link in the present context we use the terms node and link interchangeably. Define

$$\sigma_i(t) = \begin{cases} 1 & \text{if link } i \text{ transmits in slot } t \\ 0 & \text{otherwise.} \end{cases}$$

The conflict graph implies that if $\sigma_i(t) = 1$ then $\sigma_j(t) = 0$ for all $j \in \mathcal{N}_i$.

a) *Parallel Glauber Dynamics*: Let $q = \{q_m : m \in \mathcal{I}\}$ be an arbitrary but fixed probability distribution on \mathcal{I} such that $\pi_i := \sum_{m \ni i} q_m > 0$ for each link i . Construct the process

$$\sigma(t) = (\sigma_1(t), \sigma_2(t), \dots, \sigma_n(t)) : t = 0, 1, 2, \dots$$

from an arbitrary initial value as follows: For each $t > 0$

- 1) Randomly choose a decision schedule $m(t) \in \mathcal{I}$ with probability distribution q ;
- 2) For every link $i \in m(t)$ take the following steps. With probability $p_i = \frac{\lambda_i}{1+\lambda_i}$ link i samples the (local) medium. If $\sum_{j \in \mathcal{N}_i} \sigma_j(t-1) = 0$, then set $\sigma_i(t) = 1$ and otherwise set $\sigma_i(t) = 0$. Also, with probability $\bar{p}_i = 1 - p_i$ link i decides on not sampling the medium and sets $\sigma_i(t) = 0$;
- 3) For every link $i \notin m(t)$, set $\sigma_i(t) = \sigma_i(t-1)$.

Given the link activity process $\sigma(t)$, the queue-length for the flow sending traffic along link i has the following dynamics:

$$Q_i(t+1) = \max(Q_i(t) + a_i(t+1) - \sigma_i(t+1), 0). \quad (1)$$

V. Subramanian was supported in part by SFI via a Short-Term Travel Fellowship and PI grant 07/IN.1/1901, and NSF under grant CCF-0905224. M. Alanyali was funded in part by NSF through grant CNS-1018154.

Here $a_i(t)$ is the number of packets that join the backlog of link i at the beginning of slot t , and it takes one slot to transmit each packet. For simplicity we shall assume that $a_i(t) : t = 1, 2, 3, \dots$ are iid, with $\nu_i = E[a_i(t)] < \infty$ and $E[a_i^2(t)] < \infty$.

b) *Stochastic Coupling*: We will upper bound each queue length $Q_i(t)$ by comparing it with a related queue-length process whose link activity is dominated by $\sigma_i(t)$. To obtain such a process consider the subgraph \tilde{G}_i of G that is obtained by omitting all edges that are incident on neighbors of i , except those edges that connect the neighbors with i . Hence \tilde{G}_i has a disconnected component that is a star with i in the center. Let the process $\tilde{\sigma}(t)$ indicate the Parallel Glauber Dynamics on the new conflict graph \tilde{G}_i . Furthermore assume that $\sigma(t)$ and $\tilde{\sigma}(t)$ are coupled by sharing the same decision schedules and medium-sampling decisions in steps 1) and 2) of the description of the dynamics.

Since neighbors of i compete only with i to access the medium under the conflict graph \tilde{G}_i it is arguable that they are more advantageous to transmit, and in turn i transmits less frequently, relative to the setting under G . The next theorem formalizes this intuition.

Theorem 1. *If $\sigma(0) = \tilde{\sigma}(0)$ then $\sigma_i(t) \geq \tilde{\sigma}_i(t)$ for all t .*

Proof: We establish the stronger condition that if $\sigma(0) = \tilde{\sigma}(0)$ then $\sigma_i(t) \geq \tilde{\sigma}_i(t)$ and $\sigma_j(t) \leq \tilde{\sigma}_j(t)$ for all $j \in \mathcal{N}_i$ and for all t . Let $t^* > 0$ be the first time slot when this condition is violated. Then either

- a) $\tilde{\sigma}_i(t^*) = 1$ and $\sigma_i(t^*) = 0$, or
- b) $\tilde{\sigma}_j(t^*) = 0$ and $\sigma_j(t^*) = 1$ for some $j \in \mathcal{N}_i$.

Note that a) and b) are mutually exclusive since otherwise conditions set forth by the conflict graph are violated. If a) holds then there exists a neighbor $j \in \mathcal{N}_i$ such that $\sigma_j(t^* - 1) = 1$; but $\tilde{\sigma}_j(t^* - 1) = 0$ since $\tilde{\sigma}_i(t^*) = 1$. So b) should be true at time $t^* - 1$. Alternatively, if b) holds then $\sigma_i(t^* - 1) = 0$ and $\tilde{\sigma}_i(t^* - 1) = 1$ because i is the only neighbor of its neighbors in \tilde{G}_i . Therefore a) should hold at time $t^* - 1$. This conflicts with the definition of t^* ; so no such time exists. ■

Let $\tilde{Q}_i(t) : t = 1, 2, \dots$ be defined by:

$$\tilde{Q}_i(t+1) = \max(\tilde{Q}_i(t) + a_i(t+1) - \tilde{\sigma}_i(t+1), 0). \quad (2)$$

For two random variables X and Y , the notation $X \leq Y$ represents stochastic dominance [5] of Y over X . Then we have the following theorem relating Q_i and \tilde{Q}_i which is a corollary to Theorem 1.

Theorem 2. *If $Q_i(0) = \tilde{Q}_i(0)$ then $Q_i(t) \leq \tilde{Q}_i(t)$ for all t . Furthermore, the relationship also holds for the stationary queue-lengths \mathbf{Q}_i and $\tilde{\mathbf{Q}}_i$.*

Proof: We compare the two queue-lengths by considering the same sequence of arrivals to link i for both systems. Since for $t > 0$ we have

$$Q_i(t) = \max\left(0, Q_i(0) + \sum_{j=1}^t a_i(j) - \sigma_i(j),\right.$$

$$\left. \begin{aligned} & \max_{1 \leq k \leq t} \sum_{j=k}^m a_i(j) - \sigma_i(j) \Big) \\ \tilde{Q}_i(t) &= \max\left(0, \tilde{Q}_i(0) + \sum_{j=1}^t a_i(j) - \tilde{\sigma}_i(j), \right. \\ & \left. \max_{1 \leq k \leq t} \sum_{j=k}^m a_i(j) - \tilde{\sigma}_i(j) \Big) \end{aligned}$$

with $\tilde{\sigma}_i(t) \leq \sigma_i(t)$ for all $t > 0$, which implies that $Q_i(t) \leq \tilde{Q}_i(t)$. Thus, it follows [5] that we have the required stochastic dominance. From [6] we know that the stationary queue-length random variables are given by (with empty sum being 0)

$$\mathbf{Q}_i = \sup_{t \geq 0} \sum_{j=1}^t a_i(j) - \sigma_i(j), \quad \tilde{\mathbf{Q}}_i = \sup_{t \geq 0} \sum_{j=1}^t a_i(j) - \tilde{\sigma}_i(j),$$

the coupling argument again yields the result. ■

Remark 1. *The equilibrium distribution of Parallel Glauber Dynamics has a well-known product form (see for example [3, Theorem 1]). In the special case of \tilde{G}_i the equilibrium expectation of $\tilde{\sigma}_i(t)$ can be obtained by consideration of a star topology and can be expressed as*

$$\tilde{s}_i := \lim_{t \rightarrow \infty} P(\tilde{\sigma}_i(t) = 1) = \frac{\lambda_i}{\lambda_i + \prod_{j \in \mathcal{N}_i} (1 + \lambda_j)}. \quad (3)$$

This is the long-term rate at which the hypothetical queue $\tilde{Q}_i(t)$ that dominates $Q_i(t)$ is serviced; in turn the dominating process is positive recurrent if $\nu_i < \tilde{s}_i$.

c) *Scale-free Rate Region:* The achievable rate region Λ is the collection of arrival rates $\nu = \{\nu_i : i = 1, 2, \dots, n\}$ for which there exists a schedule of transmissions that stabilizes all queue lengths in the system. It is well-known that Λ is the convex hull of independent sets of G and that any interior point of Λ can be stabilized via CSMA by proper choice of fugacity vector $\lambda = \{\lambda_i : i = 1, 2, \dots, n\}$. The following theorems identify subsets of Λ for which each queue length can be bounded using local parameters. It turns, in such regions, queue lengths scale as $O(1)$ as the size of G increases.

Theorem 3. *If*

$$\nu_i < \frac{\lambda_i}{\lambda_i + \prod_{j \in \mathcal{N}_i} (1 + \lambda_j)} \quad \text{for all } i$$

then the queue length $Q_i(t)$ at each link i is dominated by a stable process uniformly for $G \setminus \mathcal{N}_i$.

Proof: By Theorem 2 $Q_i \leq \tilde{Q}_i$ where the dominating process does not depend on the topology $G \setminus \mathcal{N}_i$ beyond the immediate neighborhood of each link i , and it is stable if $\nu_i < \tilde{s}_i$. ■

By considering uniform fugacity one can obtain an explicit subset of the throughput region A in which the scale-free property holds for conflict graphs with bounded-degree:

Theorem 4. *If*

$$\nu_i < \frac{1}{1 + \frac{d^d}{(d-1)^{(d-1)}}$$

and $\lambda_i = (d-1)^{-1}$ for all i then each queue length $Q_i(t)$ is dominated by a stable process uniformly for all G with maximum degree d .

Proof: Let d_i be the degree of node i in the original conflict graph G . Since $d_i \leq d$,

$$\frac{u}{u + \prod_{j \in \mathcal{N}_i} (1+u)} \geq \frac{u}{u + (1+u)^d}$$

for any positive u . The desired conclusion follows by choosing $u = (d-1)^{-1}$ and by referring to Theorem 3 with the understanding that $\lambda_i = u$. ■

Remark 2. *Since $\Lambda \subseteq [0, 1]^n$ Theorem 4 implies that for*

$$\nu \in \frac{1}{1 + \frac{d^d}{(d-1)^{(d-1)}}} \Lambda^o$$

all queue lengths scale as $O(1)$ as $n \rightarrow \infty$ by setting $\lambda_i = (d-1)^{-1}$ for all i .

d) Mean Delay Bounds: Since the decision schedule process is independent of everything else, we will analyze the link activity in the auxiliary system by considering the time between two consecutive epochs when link i releases the medium, i.e., $\tilde{\sigma}_i$ transitions from 1 to 0. Between these two times there is exactly one block of time that $\tilde{\sigma}_1$ equals 1. Suppose we start by considering $\tilde{\sigma}_i(t-1) = 1$, then the only way that we can have $\tilde{\sigma}_i(t) = 0$ is if $m(t)$ is such that $i \in m(t)$ and link i chooses not to sample the medium. This happens with probability $\frac{\pi_i}{1+\lambda_i}$. Thus, it is easy to argue that $\tilde{\sigma}_i$ remains at value 1 for a duration that is geometrically distributed with parameter $\frac{\pi_i}{1+\lambda_i}$. When the $\tilde{\sigma}_i(t) = 0$, then the medium (at least the local neighbourhood of link i) idles for at least one slot and for more slots if either none of the links in $\{i\} \cup \mathcal{N}_i$ get picked in the decision schedule which happens with probability $1 - \pi_i - \sum_{j \in \mathcal{N}_i} \pi_j$, or if the picked neighbour does not choose to sample the medium (and switch $\tilde{\sigma}$ to 1) which happens with probability $\sum_{j \in \mathcal{N}_i \cup \{i\}} \frac{\pi_j}{1+\lambda_j}$. Thus, the idle slot is also geometrically distributed but with parameter $\frac{\pi_i \lambda_i}{1+\lambda_i} + \sum_{j \in \mathcal{N}_i} \frac{\pi_j \lambda_j}{1+\lambda_j}$.

Once the (local) medium stops being idle, either link i regains the medium or the neighbours transmit. If a neighbour grabs the (local) medium, then a period of outage of the channel occurs to link i and after this is finished an idle period ensues, and then this procedure repeats. Note that number of such episodes is geometrically distributed (starting from 0) with parameter $\frac{\pi_i \lambda_i / (1+\lambda_i)}{\pi_i \lambda_i / (1+\lambda_i) + \sum_{j \in \mathcal{N}_i} \pi_j \lambda_j / (1+\lambda_j)}$. With probability $\frac{\sum_{j \in \mathcal{N}_i} \pi_j \lambda_j / (1+\lambda_j)}{\pi_i \lambda_i / (1+\lambda_i) + \sum_{j \in \mathcal{N}_i} \pi_j \lambda_j / (1+\lambda_j)}$ one of the neighbours grabs the medium first. The key point of our analysis is that the duration of time when one of the neighbours disbars the activity of link i is of discrete phase-type [7, Chapter 2], the parameters of which can be determined solely from the parameters of link i and its neighbours. The state-

space of the phases is $\{0, 1\}^{\mathcal{N}_i}$ encoding the activity state of each of the neighbours with $\mathbf{0}$ the state corresponding to all neighbours with $\tilde{\sigma}$ being 0, being the absorbing state. The phases are denoted by the identity of the neighbours that have $\tilde{\sigma}$ being 1. A discrete phase-type distribution is parameterized by the entry distribution $\boldsymbol{\tau}$, the non-negative exit vector \mathbf{T}_0 and the (internal) non-negative transition matrix \mathbf{T} where $\mathbf{T}_0 + \mathbf{T}\mathbf{1} = \mathbf{1}$ with $\mathbf{1}$ being the all ones vector. The entry distribution $\boldsymbol{\tau}$ is easily determined as entry (from the absorbing state) is only to states $\{j\}$ where $j \in \mathcal{N}_i$, i.e., where exactly one of the neighbours j has $\tilde{\sigma}_j = 1$. This occurs with probability $\frac{\pi_j \lambda_j / (1+\lambda_j)}{\sum_{k \in \mathcal{N}_i} \pi_k \lambda_k / (1+\lambda_k)}$. The exit vector \mathbf{T}_0 is also easily determined since at the most one of the neighbours of link i is chosen at any given time by the decision schedule, the exit to the absorbing state is only through one of the states $\{j\}$ where $j \in \mathcal{N}_i$, i.e., from a state where only one neighbour has $\tilde{\sigma}$ equalling 1. From state $\{j\}$ the absorbing state is reached with probability $\pi_j / (1+\lambda_j)$, i.e., neighbour j is chosen by the decision schedule and it chooses to change the value of $\tilde{\sigma}_j$. For the transition matrix \mathbf{T} we distinguish between three types of states. Single active states $\{j\}$ where $j \in \mathcal{N}_i$. As mentioned earlier, there is a probability of transitioning to the absorbing state $\mathbf{0}$. One can also transition to state $\{j, k\}$ where $k \in \mathcal{N}_i \setminus \{j\}$ which occurs with probability $\frac{\pi_k \lambda_k}{1+\lambda_k}$. Finally, the system can remain in state $\{j\}$ otherwise, which occurs with probability $1 - \frac{\pi_j}{1+\lambda_j} - \sum_{k \in \mathcal{N}_i \setminus \{j\}} \frac{\pi_k \lambda_k}{1+\lambda_k}$. All active state is \mathcal{N}_i . One can transition to state $\mathcal{N}_i \setminus \{j\}$ when neighbour j switches state, which occurs with probability $\frac{\pi_j}{1+\lambda_j}$. The other possibility is for the system to remain in state \mathcal{N}_i , which occurs with probability $1 - \sum_{j \in \mathcal{N}_i} \frac{\pi_j}{1+\lambda_j}$. Intermediate active state are states $\{j_1, \dots, j_l\}$ where $1 < l < |\mathcal{N}_i|$. Here one of the active neighbours can switch its $\tilde{\sigma}$ to 0, which for neighbour $j \in \{j_1, \dots, j_l\}$ occurs with probability $\frac{\pi_j}{1+\lambda_j}$ with the new state becoming $\{j_1, \dots, j_l\} \setminus \{j\}$. A neighbour $k \in \mathcal{N}_i \setminus \{j_1, \dots, j_l\}$ can set its $\tilde{\sigma}$ to 1, which leads to a transition to state $\{j_1, \dots, j_l, k\}$ with probability $\frac{\pi_k \lambda_k}{1+\lambda_k}$. Finally, with probability $1 - \sum_{j \in \{j_1, \dots, j_l\}} \frac{\pi_j}{1+\lambda_j} - \sum_{k \in \mathcal{N}_i \setminus \{j_1, \dots, j_l\}} \frac{\pi_k \lambda_k}{1+\lambda_k}$ the system remains in state $\{j_1, \dots, j_l\}$. The mean of a discrete phase-type distribution with parameters $(\boldsymbol{\tau}, \mathbf{T}_0, \mathbf{T})$ is given by $\boldsymbol{\tau}(\mathbf{I} - \mathbf{T})^{-1}\mathbf{1}$ where \mathbf{I} is the identity matrix of the same size as \mathbf{T} , which for our parameters above we will denote by $\tilde{\mathbf{T}}_i$. The probability generating function of the discrete phase-type distribution above is given by $z\boldsymbol{\tau}(\mathbf{I} - z\mathbf{T})^{-1}\mathbf{T}_0 =: \tau(z)$.

Thus, from the perspective of link i , the medium is available for a duration that is geometrically distributed after which the medium is unavailable for a random duration (whose distribution can be computed) that is independent of how long the medium was available for link i and the queue-length at link i . The probability generating function of the outage time of link i , T_i^{outage} is given by

$$E[z^{T_i^{\text{outage}}}] = \frac{\frac{\pi_i \lambda_i}{1+\lambda_i}}{\sum_{j \in \mathcal{N}_i} \frac{\pi_j \lambda_j}{1+\lambda_j}} \frac{\frac{\pi_i}{1+\lambda_i} z}{1 - z + \frac{\pi_i}{1+\lambda_i} z} + 1 - \frac{\pi_i}{1+\lambda_i} z \tau(z) \quad (4)$$

An exact analysis for the delay can be carried out by assuming that the medium outage is owing to a higher priority service and that the server always serves one packet in a given slot; in fact, the results in [12] exactly cover the scenario that we are interested in.

From the description above it is clear that we have an alternating renewal process for the medium availability for link i . The average time that link i holds onto the medium is given by $\bar{T}_i = \frac{1+\lambda_i}{\pi_i}$. The average time that the medium is unavailable for link i is given by

$$\bar{T}_i^{outage} = \frac{1+\lambda_i}{\pi_i \lambda_i} \left(1 + \tilde{T}_i \sum_{j \in \mathcal{N}_i} \frac{\pi_j \lambda_j}{1+\lambda_j} \right) \quad (5)$$

Thus, average fraction of time link i can transmit is given by

$$\tilde{s}_i = \frac{\bar{T}_i}{\bar{T}_i + \bar{T}_i^{outage}} = \frac{\lambda_i}{1+\lambda_i + \tilde{T}_i \sum_{j \in \mathcal{N}_i} \frac{\pi_j \lambda_j}{1+\lambda_j}} \quad (6)$$

After some algebraic manipulations it can be verified that the expression for \tilde{s}_i above coincides with that from (3). It is easy to see that the product of \tilde{T}_i and $\sum_{j \in \mathcal{N}_i} \frac{\pi_j \lambda_j}{1+\lambda_j}$ stays the same if the relative values of $\{\pi_j : j \in \mathcal{N}_i\}$ are kept fixed. Since the variance of the time link i gets the medium or has an outage is finite, it is easy to argue using Kingman's bounds [8], [9], [10] that the mean delay will be finite for all average arrival rates strictly less than \tilde{s}_i , if the variance of the arrival process is also finite. From Theorem 2 this is an upper-bound on the mean delay for the real system.

The bound on the mean delay only depends on the fugacities of link i and its neighbours and also the probabilities of choosing neighbours of link i as per the decision schedule. Thus, for bounded degree conflict graphs we have a subset rate-region where the mean delay is uniformly bounded. Note that this result is not in conflict with [11] which considers general graphs where an adversary can pick the worst conflict graph for a given algorithm.

III. CONTINUOUS-TIME CSMA NETWORK

There is an analogous continuous-time version of CSMA [13], [14], [15] that we analyze next. Again Glauber dynamics is followed but once a link is able to transmit, it does so for a duration of time that is exponentially distributed with parameter 1. In more detail, link i samples the activity of its local medium using inter-sampling times that are exponentially distributed with parameter λ_i . At a sampling instance if the (local) medium is occupied by any of its neighbors, then link i draws an exponentially distributed time (with parameter λ_i) as the inter-sampling time to determine the next sampling instance. If, instead, at the instance of sampling the (local) medium is not being used by any of the neighbours, then link i transmits for a duration of time that is exponentially distributed with parameter 1. As soon as link i releases the medium, it resumes the sampling process.

Again our approach will be to concentrate on a particular link i and consider an auxiliary system that operates on the reduced conflict graph \tilde{G}_i . Coupling here is trickier as one has

to use the same service times for link i whenever it is served. Our solution here is to use uniformization [16] to convert the process to a discrete-time process and then use the same procedure as in the discrete-time case. Equilibrium distribution of link activities admits the same expression with the Parallel Glauber Dynamics in terms of the fugacity vector λ ; so once counterparts of Theorems 1–2 are established Theorems 3–4 apply verbatim to the continuous time model. The details are omitted owing to space considerations but we should point out that uniformization is carried out using a Poisson process at a rate that equals the sum of the fugacities of all the links plus the number of links. Also note that regular single-site Glauber dynamics results instead of the more general parallel dynamics described in Section II.

In bounding mean delay we use the instances of link i releasing the (local) medium as renewal epochs. Once link i relinquishes the medium an idle period ensues which is exponentially distributed with parameter $\lambda_i + \sum_{j \in \mathcal{N}_i} \lambda_j$. If link i has the shortest inter-sampling duration for this round, then it holds the medium once again, and thus terminating the renewal cycle once the holding time expires. This happens with probability $\frac{\lambda_i}{\lambda_i + \sum_{j \in \mathcal{N}_i} \lambda_j}$. However, if a neighbour has the shortest inter-sampling duration, then there is period of outage for link i when someone in the neighbourhood of link i is transmitting followed by a medium idle period. The number of such episodes is once again geometrically distributed (starting with 0) with parameter $\frac{\lambda_i}{\lambda_i + \sum_{j \in \mathcal{N}_i} \lambda_j}$. As before we only need to determine the distribution of the duration of the outage for link i when someone in its neighbourhood is transmitting, and now it is a phase-type distribution [7]. Now the parameters¹ are $(\alpha, \mathbf{S}_0, \mathbf{S})$ with $\mathbf{S}_0 = -\mathbf{S}\mathbf{1}$. As before the phases correspond to the set of neighbours in a transmission state with $\mathbf{0}$ being the absorbing state. The initial entry distribution α is non-zero only for states $\{j\}$ where $j \in \mathcal{N}_i$ with the value being $\frac{\lambda_j}{\sum_{k \in \mathcal{N}_i} \lambda_k}$. The exit vector \mathbf{S}_0 is non-zero only from states $\{j\}$ where $j \in \mathcal{N}_i$ with the value being 1, corresponding to link j finishing its transmission. For the transition rate matrix \mathbf{S} we only need to specify the off-diagonal entries. Once again we have three cases. The single link active state which are states $\{j\}$ (with $j \in \mathcal{N}_i$) with the only transitions being to states $\{j, k\}$ with $k \in \mathcal{N}_i \setminus \{j\}$ at rate λ_k . The all active state which is the state \mathcal{N}_i with the only transitions being to states $\mathcal{N}_i \setminus \{j\}$ at rate 1. The intermediate active state which are states $\{j_1, \dots, j_l\}$ with $1 < l < |\mathcal{N}_i|$. Here transitions can take place to states $\{j_1, \dots, j_l\} \setminus \{j\}$ at rate 1 where $j \in \{j_1, \dots, j_l\}$. The only other transitions are to states $\{j_1, \dots, j_l, k\}$ at rate λ_k where $k \in \mathcal{N}_i \setminus \{j_1, \dots, j_l\}$. The mean duration for phase-type distributions is given by $-\alpha \mathbf{S}^{-1} \mathbf{1}$; for our parameters we will once again denote this by \tilde{T}_i . Following the same steps as before we now have the following:

$$\bar{T}_i = 1, \quad \bar{T}_i^{outage} = \frac{1}{\lambda_i} \left(1 + \tilde{T}_i \sum_{j \in \mathcal{N}_i} \lambda_j \right), \text{ and}$$

¹ \mathbf{S}_0 has non-negative entries and \mathbf{S} is a transition rate matrix, i.e., the diagonal entries are negative and the rest are non-negative.




Fig. 1. Conflict graphs of the two network topologies

$$\tilde{s}_i = \frac{\bar{T}_i}{\bar{T}_i + \bar{T}_i^{outage}} = \frac{\lambda_i}{1 + \lambda_i + \bar{T}_i \sum_{j \in \mathcal{N}_i} \lambda_j}, \quad (7)$$

and once again we can derive the same expression as (3). The moment generating function of outage of link i is given by

$$E[e^{sT_i^{outage}}] = \frac{\frac{\lambda_i}{1-s}}{\lambda_i + \left(\sum_{j \in \mathcal{N}_i} \lambda_j \right) \left(1 + \frac{\alpha(s\mathbf{I} + \mathbf{S})^{-1} \mathbf{S}_0}{1-s} \right)} \quad (8)$$

Now that we have the distributions for the amount of time link i can transmit and the amount of time that the link i is in outage (both for the related system), as a concrete example we can use results from [17] to upperbound the mean delay for Poisson arrivals.

IV. SIMULATIONS

To illustrate our results we present simulation results concentrating on two simple topologies for the discrete-time case. The conflict graph for each network is shown in Figures 1(a) and 1(b), respectively. In each case we will present results for link 3.

First we concentrate on network topology *A*. Initially we assume that $\lambda_i \equiv 1$ for all $i = 1, \dots, 7$. We also assume that the decision schedule is such that the following subsets of links, namely, $\{1, 4\}$, $\{1, 6\}$, $\{2, 5\}$, $\{5, 7\}$, $\{3\}$, are chosen with equal probability. This results in $\pi = [0.4 \ 0.2 \ 0.2 \ 0.2 \ 0.4 \ 0.2 \ 0.2]$. Numerically, one can determine that $\bar{T}_3 = 37.50$ and $\tilde{s}_3 = 1/17 \approx 0.0588$, in contrast for the real system one gets a value of 0.1362. Feeding link 3 with Bernoulli traffic with an average of 0.05 arrivals per unit time, one can get the delay (by simulation) for the real queue to be 114.983 units whereas the upper bound that the related system gives is 1153.03 units. Since many service opportunities present in the real system are ignored in the auxiliary system, the upper-bound is necessarily quite loose. For the same network, differentiating the throughput (in auxiliary system) one finds that setting $\lambda_i \equiv 1/(d_3 - 1) = 1/3$ yields a maximum throughput of $27/283 \approx 0.0954$.

Next we concentrate on network topology *B*. Again we start by assuming that $\lambda_i \equiv 1$ for all $i = 1, \dots, 4$. We also assume that the decision schedule is such that the following subsets of links, namely, $\{1\}$, $\{2\}$, $\{3\}$, $\{4\}$, are chosen with equal probability. This results in $\pi = [0.25 \ 0.25 \ 0.25 \ 0.25]$. Numerically, one can determine that $\bar{T}_3 = 12$ and $\tilde{s}_3 = 1/5$, in contrast for the real system one gets a value of 0.3422. Feeding link 3 with Bernoulli traffic with an average of 0.16 arrivals per unit time and where each arrival consists of one packet

one can get the delay (by simulation) for the real queue to be 24.71 units whereas the upper bound that the related system gives is 121.81 units. For this network, under the uniform fugacity assumption, one can easily verify that the maximum throughput (in auxiliary system) of link 3 is exactly 0.2.

V. CONCLUSIONS

For both a discrete-time and a continuous-time version of CSMA with fixed fugacities, we provided a lower bound on throughput and an upper bound on the queue length, where the upper bound depends only on local contention faced by each link. For networks with a conflict graph of bounded degree, this implies the existence of a subset of the capacity-region such that if the arrival rates are in this reduced rate-region, then the mean delay can be bounded independently of the network size. Therefore, if the arrival rates are small enough, then the delay bounds implied by the mixing time of Glauber dynamics are loose and do not reveal the correct scaling property of the delay with network size.

REFERENCES

- [1] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," *Proceedings 46th Annual Allerton Conference on Communication, Control and Computing*, September 2008.
- [2] J. Ni and R. Srikant, "Distributed CSMA/CD algorithms for achieving maximum throughput in wireless networks," in *Proc. of Information Theory and Applications Workshop*, San Diego, Feb 2009.
- [3] L. Jiang, M. Leconte, J. Ni, R. Srikant and J. Walrand, "Fast mixing of parallel Glauber dynamics and low-delay CSMA scheduling," *preprint*, <http://arxiv.org/abs/1008.0227>, 2010.
- [4] T. Hayes and A. Sinclair, "A general lower bound for mixing of single-site dynamics on graphs," *The Annals of Applied Probability*, vol. 17, no. 3, pp. 931–952, 2007.
- [5] A. Müller and D. Stoyan, "Comparison methods for stochastic models and risks," *Wiley Series in Probability and Statistics*, John Wiley & Sons Ltd., Chichester, 2002.
- [6] R. M. Loynes, "The stability of a queue with non-independent interarrival and service times," *Proc. Cambridge Philos. Soc.*, 58, 497–520, 1962.
- [7] M. F. Neuts, "Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach," Dover Publications Inc., 1981.
- [8] J. F. C. Kingman, "Some inequalities for the GI/G/1 queue," *Biometrika*, 49, 315–324, 1962.
- [9] K. T. Marshall, "Some inequalities in queueing," *Operations Research*, 16, 651–665, 1968.
- [10] J. F. C. Kingman, "Inequalities in the theory of queues," *J. R. Statist. Soc.*, B 32, 102–110, 1970.
- [11] D. Shah, D. N. C. Tse, and J. N. Tsitsiklis, "Hardness of Low Delay Network Scheduling," to appear *IEEE Transactions on Information Theory*, August 2009.
- [12] D. Fiems, B. Steyaert and H. Bruneel, "Discrete-time queues with generally distributed service time and renewal-type server interruptions," *Performance Evaluation*, 55, pp. 277–298, 2004.
- [13] R. Boorstyn, A. Kershenbaum, B. Maglaris and V. Sahin, "Throughput analysis in multihop CSMA packet networks," *IEEE Trans. on Comm.*, 35(3), pp. 267–274, March 1987.
- [14] X. Wang and K. Kar, "Throughput modeling and fairness issues in CSMA/CA based *ad-hoc* networks," in *Proc. of IEEE Infocom 2005*, Miami, March 2005.
- [15] S. C. Liew, C. Kai, J. Leung and B. Wong, "Back-of-the-envelope computation of throughput distributions in CSMA wireless networks," *preprint*, <http://arxiv.org/abs/0712.1854>, 2007.
- [16] S. M. Ross, "Introduction to probability models," Seventh edition, *Harcourt/Academic Press*, Burlington, MA, 2000.
- [17] B. Sengupta, "A queue with service interruptions in an alternating random environment," *Operations Research*, 38(2), pp. 308–318, March–April, 1990.