Delay Performance of CSMA policies in Multihop Wireless Networks: A New Perspective

Mahdi Lotfinezhad and Peter Marbach
Department of Computer Science, University of Toronto
E-mail: {mahdi,marbach}@cs.toronto.edu

***Abstract*—In this paper, we study the delay performance of CSMA policies in wireless networks, where the delay is defined as the average time that a silent wireless link needs to wait until it accesses the channel for packet transmission. It is well-known that CSMA policies can incur an access delay that may be correlated over time and may grow exponentially with the network size. This discourages practical implementation of CSMA policies in even mid-sized networks. In this paper, we provide a new perspective on the delay performance of CSMA policies. We present recently developed results for two important interference models and show how CSMA policies can be used to ensure an access delay that is memoryless over time or that does not grow with the network size. The two interference models that we consider are primary interference and the "lattice interference graph". Our results suggest that CSMA policies can achieve a delay performance, as well as a delay-throughput trade-off, that makes them viable to be used in practice.**

I. Introduction

CSMA scheduling policies are examples of simple distributed policies that are throughput optimal [1]. Despite their attractive features, these policies may incur a delay that is heavily correlated over time and that may exponentially increase with the network size. This is perhaps the main hurdle towards practical implementation of CSMA policies in multihop wireless networks.

In this paper, we study the access delay of CSMA policies, i.e., the average time that a silent link needs to wait until its next packet transmission. Ideally, in a network with stochastic arrivals, it is desirable to have an access delay that is independent of the system history or the network size. We are, therefore, interested in the following questions:

- For which topologies, CSMA policies become *memoryless* in the sense that links' access delays become independent of the system history?
- Do CSMA polices fundamentally fail in topologies in which CSMA access delay exponentially increases with the network size? Or, there exists a simple fix to the operation of CSMA policies that turns the exponential delay to a delay that does not depend on the network size?

To answer the first question, we first define an *ideal network* which provides sufficient conditions under which CSMA policies become memoryless. We then consider the primary interference model [2], [3]. Using this model, we focus on network graphs that can be represented as bipartite graphs, where any of N sender nodes transmits data to any of another set of N receiver nodes. We show that under proper limiting regime, large bipartite network graphs provide examples of the defined ideal network. Therefore, these networks serve as interesting examples where CSMA access delay becomes independent of system history.

To answer the second question, we base our analysis on an idealized CSMA policy as considered in [1], [4]. We assume a continuous distribution of the random back-off time of CSMA policies to avoid packet collisions. In addition, in our analysis, we focus on networks where interference can be characterized by a lattice, or a grid-like, interference graph. The reason for assuming a lattice interference graph is that 1) under this graph, the delay of CSMA policies grows exponentially with the network size [5], and 2) it allows a formal analysis.

For the lattice interference graph, we formulate a model that allows to characterize the transient behavior of CSMA policies as how data throughput changes over time. This model is based on two simplifying assumptions that allow a formal analysis of the transient behavior. Using the model, we show that CSMA policies converge to a maximum schedule that achieves a high throughput at a rate that is independent of the network size. This insight is one of the key results of the paper. Using this insight, we propose a slight variant of the classical CSMA policy and show that its delay performance does not depend on the network size, i.e., the delay stays bounded even when the networks size approaches infinity. This result is rather surprising and at first counter-intuitive. We also characterize the throughput-delay trade-off of the proposed CSMA policy and show that this trade-off is also independent of the network size.

While our analysis, to answer the second question, has been developed for the special case of an interference graph with a lattice topology, through numerical case studies, we have verified that the intuition obtained from the analysis carries over to more general network topologies. We have also observed a similar behavior where the proposed CSMA policy is combined with a rate control mechanism.

The above results suggest that CSMA policies while throughput optimal, simple, and distributed, can be used to ensure a good delay performance. In addition, the analysis shows that CSMA policies can obtain good delay performance even for very large networks. To the best of our knowledge, this is the first time that such a result has been obtained. This result opens up the possibility of having multihop wireless networks that can be deployed in a simple distributed manner and have a very good performance in terms of both throughput and delay.

The rest of this paper is organized as follows. In the next section, we provide the network model and the details of CSMA policies. In Section III, we provide an example that motivates our study in this paper. In section IV, we focus on the first question posed above and the topologies in which CSMA policies become memoryless. In Section V, we focus on the second question posed above and introduce our modified CSMA policy. Finally, in Section VI, we conclude the paper.

II. Model

In this section, we introduce the network and CSMA models that we will use in this paper.

A. Network Model

We consider a fixed wireless network composed of a set $\mathcal{N}$ of nodes with cardinality N, and a set $\mathcal{L}$ of undirected links with cardinality L. An undirected link $l = (n, m) \in \mathcal{L}$ indicates that node n and m are within transmission range of each other and can exchange data packets.

We model the contention between links by an interference graph $G(\mathcal{L}, \mathcal{E})$, where $\mathcal{L}$ is the set of links and $\mathcal{E}$ is the set of edges. An edge $e = (l, l') \in \mathcal{E}$ in the graph $G(\mathcal{L}, \mathcal{E})$ indicates that the two links l and l', $l, l' \in \mathcal{L}$ interfere with each other. In the following, we will refer to $\mathcal{L}$ as the node set of the interference graph, and to the set $\mathcal{E}$ as its edge set.

Throughout the paper, we assume that all traffic is one-hop. We let $\lambda_l = \lambda_{(i,j)}$ indicate the total packet arrival rate to link $l = (i, j)$, and let $\boldsymbol{\lambda} = (\lambda_l)_{l \in \mathcal{L}}$ be the arrival rate vector for a given network. We assume that the rate of transmission is the same for all links and equal to unity. Moreover, if a link is scheduled for transmission but does not have any packets to transmit, dummy packets are transmitted.

B. CSMA Policies

A CSMA policy can be characterized by 1) how links make decisions as when to transmit, and 2) by a parameter β defined as the sensing delay or the idle period. When $\beta = 0$, we have the *idealized CSMA* model, which is very similar to the one presented in [1], [4]. Given a wireless network with interference graph $G(\mathcal{L}, \mathcal{E})$, every link $l \in \mathcal{L}$ independently of others senses transmissions of any conflicting link in the interference graph $G(\mathcal{L}, \mathcal{E})$, i.e. of any link l' such that the edge $e = (l, l')$ is contained in the edge set $\mathcal{E}$. If link l senses that any of its interfering links is transmitting, then it waits until all its interfering links become silent. Once this happens, link l sets a backoff timer with a value that is exponentially distributed with mean $1/z_l$, and starts to reduce the backoff timer. If the timer reaches zero before any of its interfering links start a transmission, then link l starts a transmiission. Otherwise, link l simply waits until all its interfering links become silent again, and repeats the above process. For the case of $\beta = 0$, we assume that all transmission times (packet lengths) are independently (across links and over time) and exponentially distributed with unit mean.

The above models an idealized CSMA protocol [1], [4]. Since the distribution of the backoff timers is continuous, the probability that two interfering links start to transmit a packet exactly at the same time is equal to zero, and we assume that we can ignore the collision of the transmissions of any two interfering links. In addition, we assume that links can always sense transmissions of their interfering links, and there is no "hidden-terminal" problem that can create packet collisions.

When $\beta > 0$, we consider a synchronized time-slotted system where each timeslot has duration β. For this case, CSMA policy works as follows: every link $l \in \mathcal{L}$ independently of others senses transmissions of any conflicting link. We define the channel to be sensed as idle by link l, if link l senses no transmission of its conflicting links. If the channel is sensed as idle by link l for the duration of one given timeslot, then in the following timeslot, link l starts transmission of one single packet with probability p_l, independent of all other events in the network. For the case of $\beta > 0$, we assume packet transmission times (packet lengths) are geometrically distributed with unit mean, independently across links and over time. If link l does not start a packet transmission, then link l has to sense the channel as idle for the duration of another timeslot before it again has the chance to start a packet transmission as explained above.

When $\beta = 0$, we can characterize a CSMA policy by the vector $\mathbf{z} = (z_l)_{l \in \mathcal{L}}$ where $z_l > 0$, and refer to z_l as the *transmission attempt-rate* of link l. When $\beta > 0$, we can characterize a CSMA policy by transmission attempt probability vector $\mathbf{p} = (p_l)_{l \in \mathcal{L}} \in [0, 1]^L$. Given a CSMA policy characterized by $\mathbf{z}$ or $\mathbf{p}$, let μ_l, $l \in \mathcal{L}$, be the service rate of link l, i.e., the fraction of time link l is transmitting under the given CSMA policy.

For a given network, we say that the given CSMA policy stabilizes the network for a given rate vector $\boldsymbol{\lambda}$ if

$$\lambda_l < \mu_l, \qquad l \in \mathcal{L}.$$

Given a fixed network, we then define the achievable rate region $\mathcal{C}$ of the idealized CSMA policy with $\beta = 0$ as the set of all rate vectors $\boldsymbol{\lambda}$ for which there exists a vector $\mathbf{z}$ that stabilizes the network for $\boldsymbol{\lambda}$, i.e. we have that $\lambda_l < \mu_l(\mathbf{z})$, $l \in \mathcal{L}$.

It is well-known that the idealized CSMA policy with $\beta = 0$ is throughput optimal, i.e., the set $\mathcal{C}$ contains all arrival rate vectors $\boldsymbol{\lambda}$ that are inside the capacity region Γ, where Γ is the set of all $\boldsymbol{\lambda}$'s that can be stabilized by any other policy [1].

C. Primary Interference Model

In this model, a packet transmission over link $l = (i, j)$ is successful if only if within the transmission duration, there exists no other activity over any other link (m, n) which shares a node with (i, j) [2], [3]. The primary interference model applies, for example, to wireless systems where multiple frequencies/codes are available (using FDMA or CDMA) to avoid interference, but each node has only a single transceiver

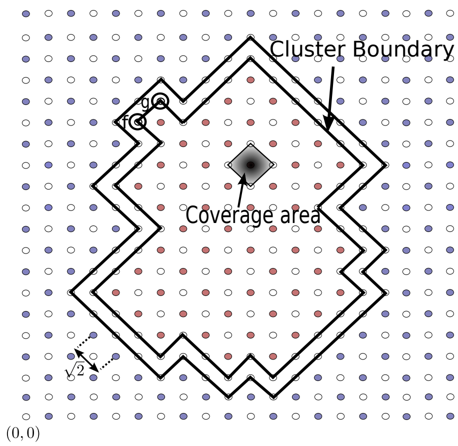

Fig. 1. Demonstration of the lattice graph G_L in $\mathbb{R}^2$; active links and their coverage area; and clusters and their boundary. Active links are colored.

and hence can only send to or receive from one other node at any time.

D. Lattice Interference Graph

In the lattice interference graph, the interference graph $G(\mathcal{L}, \mathcal{E})$ can be represented by a two-dimensional grid or lattice (see Fig. 1 for an illustration). More precisely, we assume that the interference graph $G(\mathcal{L}, \mathcal{E})$ is such that all links $l \in \mathcal{L}$ of the interference graph can be represented by *coordinates* (i, j), $i, j \in \{0, ..., n\}$, and there is an edge $e \in \mathcal{E}$ between any two links $l = (i, j)$ and $l' = (i', j')$, $l, l' \in \mathcal{L}$, if l and l' differ in exactly one coordinate and we have that

$$|i - i'| + |j - j'| = 1.$$

Throughout the paper, we use the notation $G_L = G_L(\mathcal{L}, \mathcal{E})$ to indicate that the interference graph can be represented by a lattice as described above.

Given a lattice interference graph G_L with links (i, j), $i, j \in \{0, ..., n\}$, we define $\mathcal{B}(G_L)$ as the boundary of G_L, i.e., $\mathcal{B}(G_L)$ is the set of all links for which at least one coordinate is equal to 0 or n. For the purpose of illustration, we assume that each link l represented by coordinates (i, j) can be interpreted and mapped to the point (i, j) in $\mathbb{R}^2$. With such an extension, we have mapped the vertex set $\mathcal{L}$ of G_L to a subset of points in $\mathbb{R}^2$.

Given a lattice interference graph G_L, we define a link $l = (i, j) \in \mathcal{L}$ as an *even link* iff $i + j$ is an even number. We define $\mathcal{L}^{(e)}$ as the set of all such even links. Similarly, we define a link $l = (i, j) \in \mathcal{L}$ as an *odd link* iff $i + j$ is an odd number, and define $\mathcal{L}^{(o)}$ as the set of all odd links. Note that all even links can transmit simultaneously without causing an interference. The same holds for all odd links.

For the lattice interference graph G_L, in the following, we focus on the idealized CSMA policies with $\beta = 0$ and uniform transmission attempt-rates so that

$$z_l = z, \qquad z > 0, \ l \in \mathcal{L}.$$

In addition, we also focus on the case of uniform packet arrival rates, i.e., we have

$$\lambda_l = \lambda, \qquad 0 < \lambda < 0.5, \ l \in \mathcal{L}. \tag{1}$$

III. Motivating Example

Consider the lattice interference graph model as defined in Section II-D. An important question for this model is to determine the maximal packet arrival rate λ that can be stabilized by a CSMA policy $\mathbf{z}$. By now it is well-known that for any packet arrival rate $\lambda < 0.5$, there exists a CSMA policy $\mathbf{z}$ that can stabilize the network [4]–[6]. Roughly, this can be achieved by choosing a CSMA policy $\mathbf{z}$ with uniform tranmission attempt-rates $z_l = z$, $l \in \mathcal{L}$, and by letting z become larger and larger. In the case where the attempt rate z becomes large, the network state will mainly alternate between two transmission patterns where either mostly links in the set of even links $\mathcal{L}^{(e)}$, or links in the set of odd links $\mathcal{L}^{(o)}$, are transmitting. However, transitions between these two transmission patters happen very infrequently, i.e., the CSMA policy tends to "lock into" one of the two transmission patterns for a very long time before it switches to the other pattern.

This "locking-in" behavior has two important consequences. First, the time epochs at which a link l accesses the channel are heavily correlated over time. In other words, CSMA has scheduling memory. Second, the average delay a silent link needs to wait until it accesses the channel can be prohibitively large. It is well-known that this "locking-in" behavior of the CSMA policy is necessary in order to achieve a high throughput, i.e., in order to support an arrival rate close to 0.5. However, while locking into a transmission pattern benefits the throughput, it dramatically hurts the delay performance. An exact delay analysis of this case is difficult; the currently best characterization of the delay performance (see, e.g., [7]) provide an upper-bound for the delay of the form z^L, i.e. a delay bound which increases exponentially in the network size L. To support an arrival rate of $\lambda = 0.5 - \epsilon$, $\epsilon > 0$, it has been shown that an attempt rate z roughly of the order $z = 1/\epsilon$ is needed [8]. As a result, in order to support an arrival rate of $\lambda = 0.5 - \epsilon$, $\epsilon > 0$, the upper-bound on the delay can grow as fast as $(1/\epsilon)^L$. Furthermore, for the case where ϵ is sufficiently small, a lower-bound on the delay has been derived that grows exponentially with exponent $\sqrt{L}/(\log L)^2$ [9]. These bounds imply that the delay performance of CSMA becomes very poor for large networks, i.e., as L becomes large.

The fact that CSMA delay can be heavily correlated over time and exponentially large discourages practical implementation of CSMA protocols, even for mid-sized wireless networks. This motivates to ask the following questions. First, is a correlated access delay a fundamental property of CSMA policies, or otherwise, there exist non-trivial topologies for which CSMA access delay exhibits memory-less properties? Second, for topologies such as the grid interference graph, is it true that any CSMA-like policy incurs large delay, or otherwise, it is possible to dramatically improve the delay by only slightly modifying the classical CSMA protocol? Of

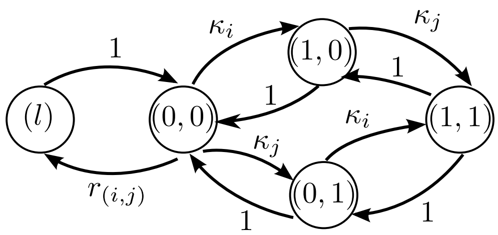

Fig. 2. Markov chain from the view point of link $l = (i, j)$. We have shown *transition rates* in the limit of $\beta \to 0$.

course, any modification of the CSMA protocol should be done in a such a way that it does not negatively affect the attractive feature of CSMA policies, namely, the fact that these policies can be easily implemented in a distributed manner. In the following, we provide answers to the above questions.

IV. Topologies with Memoryless CSMA Delay

To get intuition into the cases where CSMA policies may have a memoryless delay, i.e., an access delay that is not correlated over time and does not depend on the history of the system, we first introduce the notion of an ideal network, as follows. Throughout this section, we consider CSMA policies with $\beta > 0$, and hence, a timeslotted system, as described in Section II-B.

A. An Ideal Picture

In this section, we define an *ideal network* and discuss its access delay properties. Consider link $l = (i, j)$, from node i to node j. We define an ideal network to be the one in which for any link $l = (i, j)$ the following hold:

- When link l is not active, and at least one of the nodes i or j are active, the idle and active periods of node i (node j) are geometrically distributed with mean $\frac{1}{\kappa_i}$ and one, respectively, independent of activities of nodes j (node i).
- When link l is idle, in the next timeslot, with probability $p_l = p_{(i,j)} = r_{(i,j)}\beta$, the link becomes active, otherwise nodes i or j independently of each other become active with probability $\kappa_i\beta$ or $\kappa_j\beta$, respectively.

Under the above ideal hypothesis, the idle and active periods of link $l = (i, j)$ can be characterized by a simple Markov chain with state $S(n)$ at timeslot n. The state space of this Markov chain can be defined as $\mathcal{S} = \{S_k, k = 1, \cdots, 5\} = \{(s_i, s_j)\} \cup \{(l)\}$, where $s_i = 0$ means node i is idle, and $s_i = 1$ means it is active. The state (l) means that link l is active (transmitting). Fig. 2 shows how the ideal network will look like from the view point of link l, in terms of a Markov chain with *transition rates* in the limit of $\beta \to 0$.

We see that by choosing

$$r_{(i,j)} = \frac{\lambda_{(i,j)}}{1 - \lambda_{(i,j)}}(1 + \kappa_i)(1 + \kappa_j), \tag{2}$$

asymptotically as $\beta \to 0$, we have $\frac{\pi_l}{\lambda_{(i,j)}} = 1$, where π_l is the steady state probability that link l is transmitting. Hence, if the ideal hypothesis holds, it becomes fairly easy to choose the correct transmission probability to support any supportable throughput for link l.

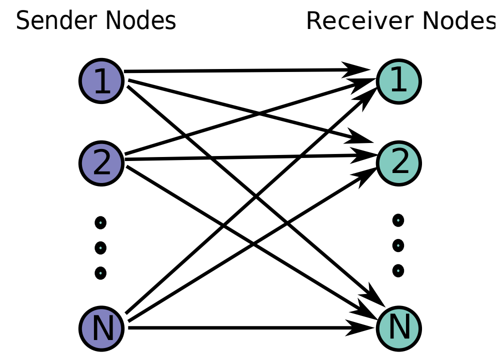

Fig. 3. Illustration of bipartite network graph.

Define τ_l, as the access delay, to be the time until the next packet transmission by link l given that link l is not currently active. We have the following theorem [10]

Theorem 1. *Let $\tau_l^s = \lambda_l \tau_l$ be the scaled version of the access delay τ_l. Under the ideal network assumption, independent of the past*

$$\lim_{\lambda_l \to 0} \lim_{\beta \to 0} \tau_l^s \xrightarrow{\mathcal{D}} X,$$

where X is an exponential r.v. with unit mean, and convergence is in distribution.

This theorem essentially states that, due to the Markovian nature of the system, the scaled access delays are i.i.d. and exponentially distributed with unit mean. Hence, CSMA is memoryless on the ideal networks with small flows (since $\lambda_l \to 0$). We next study an important network topology that in a proper limiting regime becomes an an ideal network with memoryless delay.

B. Large Bipartite Graphs as Ideal Networks

In this section, we show that ideal networks as defined in Section IV-A are indeed feasible by considering the limit of large bipartite graphs. Consider a sequence of networks $\{(\mathcal{N}^{(N)}, \mathcal{L}^{(N)}), N \in \mathbb{N}\}$, where each network can be represented as a bipartite graph with a set of N sender nodes $\mathcal{N}_S = \{1, \cdots, N\}$ and a set of N receiver nodes $\mathcal{N}_R = \{N + 1, \cdots, 2N\}$ (see Fig. 3). For the Nth network, we assume the sensing period is $\beta^{(N)}$ such that

$$\lim_{N \to \infty} N^2 \beta^{(N)} = 0. \tag{3}$$

We also assume a symmetric CSMA policy $\mathbf{p}^{(N)} = \{p^{(N)}_{l=(i,j)}\}$ on each network where the transmission probability of link l is given by

$$p^{(N)}_{l=(i,j)} = \frac{\kappa^2}{N}\beta^{(N)}, \ i \in \mathcal{N}_S, \ j \in \mathcal{N}_R. \tag{4}$$

The above choice for transmission probabilities is motivated by the results in [11]. Using the results of [11], it is not hard to see that for large κ, the service rate or throughput of any link is close to $\frac{1}{N}(1 - \frac{1}{\kappa})$.

Let $Y_S^{(N)}(n) = (Y_{(i)}^{(N)}(n), i \in \mathcal{N}_S)$ be the vector of sender nodes' states at timeslot n, where $Y_{(i)}^{(N)}(n) = 1$, $i \in \mathcal{N}_S$, indicates that node i is transmitting, otherwise $Y_{(i)}^{(N)}(n) = 0$. Similarly, let $Y_R^{(N)}(n) = (Y_{(j)}^{(N)}(n), j \in \mathcal{N}_R)$ be the vector of receiver nodes' states at timeslot n, where $Y_{(j)}^{(N)}(n) = 1$, $j \in \mathcal{N}_R$, indicates that node j is receiving, otherwise $Y_{(j)}^{(N)}(n) = 0$. Let $M_{S,0}^{(N)}(t)$ be the occupancy measure for idle sender nodes, i.e.,

$$M_{S,0}^{(N)}(t) = \frac{1}{N} \sum_{i=1}^{N} \mathbf{1}_{Y_{(i)}^{(N)}(n)=0}.$$

Similarly, let the occupancy measure for idle receiver nodes be

$$M_{R,0}^{(N)}(t) = \frac{1}{N} \sum_{j=N+1}^{2N} \mathbf{1}_{Y_{(j)}^{(N)}(n)=0}.$$

Finally, let Φ^∞ be the positive root of

$$1 - \kappa^2(\Phi^\infty)^2 - \Phi^\infty = 0.$$

We have the following result [10]

Theorem 2. *Suppose* $\lim_{N\to\infty} N^2\beta^{(N)} = 0$. *For any* $\epsilon > 0$,

$$\lim_{N\to\infty} \lim_{\tau\to\infty} \frac{1}{\tau} \int_0^\tau \mathbf{1}_{|M_{S,0}^{(N)}(t) - \Phi^\infty| < \epsilon}\, dt = 1, \text{ a.s.} \tag{5}$$

The above holds if $M_{S,0}^{(N)}(t)$ *is replaced with* $M_{R,0}^{(N)}(t)$.

Since

$$\lim_{\kappa\to\infty} \frac{\Phi^\infty}{\kappa^{-1}} = 1, \tag{6}$$

using Theorem 2, we observe that for large κ, the fraction of idle sender (or receiver) nodes should be close to κ^{-1} almost at all time instants[1]. Using (4), it is clear that if at all times the fraction of idle sending nodes and the fraction of idle receiving nodes are $\frac{1}{\kappa}$, then for instance each silent sender node attempts to access the channel with probability $\kappa\beta$. Hence, in limit of $\beta \to 0$, distribution of idle periods becomes independent of each other and exponentially distributed with mean $\frac{1}{\kappa}$, in which case we have an ideal network where CSMA access delay becomes memoryless. Theorem 2 shows that this becomes arbitrarily accurate in the limit of $N \to \infty$.

The bipartite network graphs provide an example where CSMA access delay becomes memoryless. We next show that for topologies such as the lattice interference graph, one can slightly change the operation of CSMA policies to convert an exponentially increasing delay to a delay that does not increase with the network size.

[1]In [10], we have shown that over any *finite* time interval, with probability approaching one as $N \to \infty$, the fraction of idle sender (or receiver) nodes quickly converges to Φ^∞ with rate $1 + 2\kappa^2\Phi^\infty$ and stays close to Φ^∞ for the rest of the interval.

V. CSMA Policies with Network-size Independent Delay

As stated in Section III, for the lattice interference graph, CSMA delay increases with the network size. To obtain a delay that is constant as a function of the network size, we modify the operation of CSMA policies. Throughout this section, we consider the idealized CSMA policies with $\beta = 0$, as described in Section II-B.

The basic idea behind our proposed CSMA policy is very simple. Periodically, i.e., at times $t = 0, T, 2T, 3T, ...$, we reset the transmission pattern of a CSMA policy by requiring all links to become silent and restart the CSMA protocol. Clearly, this approach will prevent a CSMA policy from locking into a particular transmission pattern for too long. As a result, this "resetting" or "unlocking" of the transmission pattern should help the delay performance as it prevents links from being locked-out from making any transmissions for long periods of time. However, on the other hand, it may dramatically hurt the throughput performance, i.e., if the unlocking period T is too small, then CSMA policy may fail to reach a good transmission pattern that is necessary to obtain a high throughput. Therefore, to understand the performance of this proposed CSMA policy we have to characterize both how the delay and throughput behave as a function of the "unlocking" period T.

In this paper, we characterize the delay and throughput performance of this modified CSMA policy with the above unlocking mechanism for a lattice interference graph and uniform transmission attempt and packet arrival rates, as described in Section II-D. While our analysis is carried out for a restricted model, through simulations, we have verified and illustrated the obtained results for more general network topologies, as well as more general transmission attempt-rates. In addition, to consider general packet arrival rates, we have studied how the proposed CSMA policy perform when combined with a rate control mechanism. The numerical results for these case studies will be provided in an extended version of this paper.

We note that the above alternative CSMA policy is an idealized scheme as it assumes that all links become immediately silent every T time units. One possible approach to implement our CSMA policy in a fully distributed manner is through propagating busy tones. Further discussion of this topic will be provided in an extended version of this paper.

A. Modeling Assumptions

Our analytical results presented in the next section are based on two simplifying assumptions. To formally state the assumptions, we first give several definitions.

Consider the lattice G_L, and for simplicity assume that links in G_L use the idealized CSMA policy with $\beta = 0$ for transmission. We define a link $l \in \mathcal{L}$ to be *active* at time t if it is transmitting at that time. At any time t, the set of active links can be considered as the union of clusters. Each cluster $\mathcal{C}_i$, $1 \le i \le i_{max}$, $i_{max} < \infty$, has the following properties:

1) C_i is a subset of active links in G_L.

2) For any two links l and l', where $l \neq l'$ and $l, l' \in \mathcal{C}_i$, there exists a path of n links $\{l_1, l_2, \cdots, l_n\}$ in $\mathcal{C}_i$ for some $n \geq 2$ where $l_1 = l$ and $l_n = l'$, such that $l_k = (i_k, j_k) \in \mathcal{C}_i$, $1 \leq k \leq n$, and $|i_k - i'_{k+1}| = |j_k - j'_{k+1}| = 1$. If $l = l' \in \mathcal{C}_i$, we define $\{l\}$ as a path.
3) $\mathcal{C}_i$ is maximal is the sense that no further links can be added to $\mathcal{C}_i$ without violating one of the above properties.

By the above definition, each cluster contains only odd active links or only even active links. We define an odd (resp. even) cluster to be a cluster consisting of odd (resp. even) links. We use $K^{(o)}(t)$ and $K^{(e)}(t)$ to be the number of odd and even clusters at time t, respectively. We define $I_o(t)$ to be the index set of all odd clusters at time t, i.e., $I_o(t) = \{i : \mathcal{C}_i \text{ is an odd cluster.}\}$. Similarly, we define $I_e(t)$ to be the index set of all even clusters at time t. The above properties define a cluster as a maximal set of even active links or odd active links, where each link in the cluster can reach any other in the cluster in a sequence of links, or path, in the cluster. Considering the mapping from $\mathcal{L}$ to $\mathbb{R}^2$ as explained in Section II-D, along the path, the euclidean distance of one link to the next is $\sqrt{2}$. In Fig. 1, colored links inside the inner polygon represent one cluster. We define $|\mathcal{C}_i|$ as the number of links in $\mathcal{C}_i$. We may refer to $|\mathcal{C}_i|$ as the size of cluster $\mathcal{C}_i$.

We also need to define the boundary of each cluster. To do so, first consider links that are inside the lattice, i.e., all links $l \notin \mathcal{B}(G_L)$. Any such link has four interfering links on the lattice. Considering the mapping from the links in $\mathcal{L}$ to the points in $\mathbb{R}^2$, for any link l inside the lattice, we define its *coverage area* A_l to be the *square* formed by its four closest links. In Fig. 1, we have shown the coverage area of one active link. For any link that is not inside the lattice, i.e, $l \in \mathcal{B}(G_L)$, we define A_l to be the intersection of the area $[0, n] \times [0, n]$ in $\mathbb{R}^2$ and the square that would exist if link l were also inside the lattice.

For each cluster $\mathcal{C}_i$, we can define its coverage area as

$$A_{\mathcal{C}_i} = \cup_{l \in \mathcal{C}_i} A_l,$$

i.e., the union of the coverage area of all links that belong to $\mathcal{C}_i$. The area $A_{\mathcal{C}_i}$ contains some points (i, j), where links of G_L may be located, and also some points in $\mathbb{R}^2$ where links are not located. For instance, the area inside the inner polygon in Fig. 1 is the area of one cluster. We also conveniently define the area of the lattice to be $\mathcal{A}(G_L) = L = (n+1)^2 > n^2$.

For the area $A_{\mathcal{C}_i}$, we define its boundary to be the set of all points in $\mathbb{R}^2$ that any neighbourhood of which contains points both in $A_{\mathcal{C}_i}$ and points not in $A_{\mathcal{C}_i}$. We use $\mathcal{B}(\mathcal{C}_i)$ to denote the boundary of $A_{\mathcal{C}_i}$ *restricted* to be inside the area $[0, n] \times [0, n]$ (see Fig. 1). We define $|\mathcal{B}(\mathcal{C}_i)|$ as the length of the boundary of cluster $\mathcal{C}_i$.

Finally, we define three quantities that define the average properties of odd clusters at a given time t. Recall that $K^{(o)}(t)$ is the number of odd clusters at time t, $I_o(t)$ is the index set of all odd clusters at time t, and $\mathcal{A}(G_L)$ is defined as the area of the lattice. We define

$$\overline{\rho_{L,\mathcal{C},t}} = \frac{K^{(o)}(t)}{\mathcal{A}(G_L)},$$

to be the *density* of odd clusters at time t. We also define

$$\overline{|\mathcal{C}_{L,t}|} = \frac{\sum_{i \in I_o(t)} |\mathcal{C}_i|}{K^{(o)}(t)},$$

to be the average number of links per odd cluster at time t. Similarly, we define

$$\overline{|\mathcal{B}(\mathcal{C}_{L,t})|} = \frac{\sum_{i \in I_o(t)} |\mathcal{B}(\mathcal{C}_i)|}{K^{(o)}(t)}$$

to be the average boundary-length per odd cluster at time t.

Based on the above definitions, we make the following two assumptions to study the transient behavior of CSMA policies. Our first assumption states that the average boundary-length of clusters does not grow faster than the square root of the average size of clusters. Without loss of generality, we have stated the assumption in terms of odd clusters.

Assumption 1. *For any time $t > 0$, odd clusters have a regular structure, in the sense that for some constant $c > 0$*

$$c\overline{|\mathcal{B}(\mathcal{C}_{L,t})|}^{\,2} < \overline{|\mathcal{C}_{L,t}|} < \infty. \tag{7}$$

Our second modeling assumption characterizes statistical structure of clusters' boundaries. Note that the boundary of each cluster is consisted of *line-segments* whose lengths are multiples of $\sqrt{2}$ (see Fig. 1). Moreover, moving along the boundaries is possible only in four directions. In the following, by the "normalized length" of a line-segment, we mean the length of the line-segment divided by $\sqrt{2}$. The following assumption states that the normalized lengths of line-segments are geometrically distributed, and the direction of line-segments are appropriately random:

Assumption 2. *Boundaries of clusters are consisted of line-segments with the following properties:*

a) *The normalized lengths of these line-segments are geometrically distributed with parameter $\kappa(t)$.*
b) *Independent of system history before time t, with probability approaching one as $L \to \infty$,*
   - *the average normalized length over all line-segments is $\kappa(t)^{-1}$, and*
   - *moving along the boundaries, at least a fraction $\upsilon\kappa(t)$ of line-segments have unit normalized length with the property that direction of the next line-segment is opposite to the direction of the segment before the current, where $\upsilon > 0$ is a constant.*

In Fig.1, the only line-segment that has the properties mentioned in Assumption 2(b) is the line-segment from link f to g, denoted by $f - g$.

B. Delay and Delay-throughput Trade-off: Analytical Results

In this section, we describe the main results of our analysis on the lattice graph G_L with uniform attempt-rates z as described in Section II-D. The proofs for these results will be provided in an extended version of this paper.

We first give a few definitions. For G_L with uniform attempt-rate z, we define $\rho_L(t) = \rho_L(t, z)$ to be the fraction of active links at time t, and $\bar{\rho}_L = \bar{\rho}_L(z)$ as the expected value of $\rho_L(t)$ as $t \to \infty$:

$$\bar{\rho}_L = \bar{\rho}_L(z) = \lim_{t\to\infty} \mathbb{E}[\rho_L(t, z)]. \tag{8}$$

We also define $\delta_L(t)$ as the gap between the density $\rho_L(t)$ at time t and $\bar{\rho}_L$:

$$\delta_L(t) = \bar{\rho}_L - \rho_L(t). \tag{9}$$

Our first result states how fast CSMA converges to the packed schedules. It states that, quite surprisingly, the convergence rate becomes independent of the network size, and that the rate does not decrease as the size of the network increases to infinity. More precisely, as a function of time t, the gap to the optimal density of active links drops as fast as $1/\sqrt{t}$, with high probability in the limit of large networks and attempt-rate z.

Theorem 3. ***Fast Convergence to Packed Schedules*** *Consider a sequence of $\{G_L\}$ with $L \to \infty$ and $z \to \infty$ such that $z = z(L) = o(L^{\frac{1}{3}})$. Under Assumptions 1-2, and assuming idealized CSMA on G_L, without the unlocking mechanism, there exists a constant $C_1 > 0$, such that for any $\tau < \infty$,*

$$\lim_{L\to\infty} P\left[\delta_L(t) < \frac{C_1}{\sqrt{t}},\ t \in (0, \tau]\right] = 1.$$

Through simulations, we have verified that $\delta_L(t)$ exhibits the above behavior and drops as $1/\sqrt{t}$ (neglecting constants). Interestingly, we have observed a similar behavior for more general settings of random interference graphs and non-uniform attempt-rates. This suggests that the above result is applicable to a much wider set of network scenarios.

Our second result provides a bound for the average queue sizes, and characterizes the throughput-delay trade-off of the modified CSMA. To state the result, let $Q_l(t)$ be the queue size of link l at time t. The theorem states that in order to get ϵ close to the throughput limit, the expected time-averaged $Q_l(t)$ becomes only $O(\frac{1}{\epsilon^3})$. This is achieved by choosing the unlocking period T to be on the order of $\frac{1}{\epsilon^2}$. By Little's theorem, we have that the packet delay is also $O(\frac{1}{\epsilon^3})$.

Quite surprisingly, the delay bound and the trade-off are valid for the limit of infinite lattice G_L, where $L = \infty$, which can be considered as the lattice of $\mathbb{Z}^2$. Hence, even in an infinite network, the modified CSMA can provide bounded delay. This implies that CSMA delay and the throughput-delay trade-off can be indeed made independent of the network size, as the network size grows to infinity.

Theorem 4. ***Delay-Throughput Trade-off*** *Consider idealized CSMA with the (idealized) unlocking mechanism. Let $\epsilon = 0.5 - \lambda > 0$, where λ is links' packet arrival rate. Let the unlocking period be $T = \frac{C_2}{\epsilon^2}$, where $C_2 > 0$ is a sufficiently large constant. Suppose the interference graph is the infinite lattice of $\mathbb{Z}^2$. Then, for sufficiently large z and under Assumption 1 and Assumption 2, for any $l \in \mathcal{L}$ we have*

$$\lim_{\tau\to\infty} \frac{1}{\tau}\mathbb{E}\left[\int_{t=0}^{\tau} Q_l(t)dt\right] = O\Big(\frac{1}{\epsilon^3}\Big).$$

Through simulations for a 26×26 lattice and attempt rate $z = 100$, we have found a close match with the above throughput-delay behavior. We have also observed a similar utility-delay trade-off for the case where rate control is combined with the new CSMA, both in the lattice and random interference graphs. Therefore, we expect our results and the general trends observed here to be applicable to a wide range of network settings.

Theorem 3 and Theorem 4 state that even in topologies that original CSMA policies may fail to provide acceptable delay performance or delay-throughput trade-off, one can slightly change the operation of CSMA policies to achieve a delay as well as a delay-throughput trade-off that becomes independent of the network size. As stated earlier, we expect our results to hold in more general settings than the lattice interference graph.

VI. Conclusion

In this paper, we have studied the delay performance of CSMA policies in wireless networks. We first have defined an ideal network for which the CSMA access delay shows memoryless properties. Next, we have shown that in the limit, large bipartite graphs with small flows provide an example of ideal networks. Hence, we have verified that there are important topologies for which CSMA access delay exhibits desirable properties. We then have considered the lattice interference graph for which CSMA access delay can exponentially grow with the network size. For this topology, we have shown that using an unlocking mechanism, one can change the operation of CSMA policies so that in the limit of large networks and attempt rates, the rate of convergence to maximum schedules and the delay-throughput trade-off become independent of the network size. These results are encouraging as they suggest that CSMA policies not only are simple, distributed, and throughput optimal, but also can be used to ensure acceptable packet delay.

References

[1] L. Jiang and J. Walrand, "A distributed algorithm for maximal throughput and optimal fairness in wireless networks with a general interference model," EECS Department, University of California, Berkeley, Tech. Rep., Apr 2008. [Online]. Available: http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-38.html

[2] X. Lin and N. Shroff, "The impact of imperfect scheduling on cross-layer congestion control in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 14, no. 2, pp. 302–315, Apr. 2006.

[3] X. Wu and R. Srikant, "Regulated maximal matching: A distributed scheduling algorithm for multi-hop wireless networks with node-exclusive spectrum sharing," in *44th IEEE conference on decision and control, and European control conference CDC-ECC*, Dec. 2005.

[4] M. Durvy, O. Dousse, and P. Thiran, "Border effects, fairness, and phase transition in large wireless networks," in *IEEE INFOCOM'08*, April 2008, pp. 601–609.

[5] D. Aldous, G. Grimmett, C. Howard, H. Kesten, F. Martinelli, L. Saloff-coste, and J. Steele, *Probability on Discrete Structures*. Springer, 2004.

[6] M. Durvy, O. Dousse, and P. Thiran, "Self-Organization Properties of CSMA/CA Systems and Their Consequences on Fairness," *IEEE Transactions on Information Theory*, vol. 55, no. 3, 2009.

[7] S. Rajagopalan, D. Shah, and J. Shin, "Network adiabatic theorem: an efficient randomized protocol for contention resolution," in *SIGMETRICS '09: Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems*, 2009, pp. 133–144.

[8] D. S. Gaunt and M. E. Fisher, "Hard-Sphere Lattice Gases. I. Plane-Square Lattice," *Journal of Chemical Physics*, vol. 43, pp. 2840–2863, Oct. 1965.

[9] C. Borgs, J. Chayes, A. Frieze, J. H. Kim, P. Tetali, E. Vigoda, and V. H. Vu, "Torpid mixing of some monte carlo markov chain algorithms in statistical physics," in *40th Annual Symposium on Foundations of Computer Science*, 1999, pp. 218–229.

[10] M. Lotfinezhad and P. Marbach, "On channel access delay of CSMA policies in wireless networks with primary interference constraints," in *Allerton Conference*, Oct. 2009.

[11] P. Marbach, A. Eryilmaz, and A. Ozdaglar, "Achievable rate region of csma schedulers in wireless networks with primary interference constraints," *46th IEEE Conference on Decision and Control*, pp. 1156–1161, Dec. 2007.