This paper was presented as part of the main technical program at IEEE INFOCOM 2011

Throughput-Optimal Random Access with Order-Optimal Delay

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

***Abstract*—In this paper, we consider CSMA policies for scheduling packet transmissions in multihop wireless networks with one-hop traffic. The main contribution of the paper is to propose a novel CSMA policy, called Unlocking CSMA (U-CSMA), that enables to obtain both high throughput and low packet delays in large wireless networks. More precisely, we show that for torus interference graph topologies with one-hop traffic, U-CSMA is throughput optimal and achieves order-optimal delay. For one-hop traffic, the delay performance is defined to be order-optimal if the delay stays bounded as the network-size increases. Simulations that we conducted suggest that (a) U-CSMA is throughput-optimal and achieves order-optimal delay for general geometric interference graphs and (b) that U-CSMA can be combined with congestion control algorithms to maximize the network-wide utility and obtain order-optimal delay. To the best of our knowledge, this is the first time that a simple distributed scheduling policy has been proposed that is both throughput/utility optimal and achieves order-optimal delay.**

I. Introduction

One of the most intriguing research challenges in the context of wireless networking is the design of a scheduling policy that

a) is throughput optimal,
b) achieves a low packet delay, and
c) has a simple and fully distributed implementation.

From a complexity theoretic viewpoint, unless $\mathbf{NP} \subseteq \mathbf{BPP}$ or $\mathbf{P} = \mathbf{NP}$, there does not exist a universal scheduling policy that has the above three properties for all possible network topologies [1]. However, it is still possible to design a policy that has the above properties for a subset class of network topologies. This seems to be true for geometric networks [2], [3], in which only links that are geometrically close interfere with each other. These networks closely approximate a wide range of practical wireless networks, and are known to admit Polynomial-Time Approximation Scheme (PTAS) for several NP-hard optimization problems (see e.g., [2], [4]). For this reason, we focus in this paper on the design of scheduling policies for large geometric wireless networks.

There are two main approaches for the design of scheduling policies in wireless networks: matching policies [2], [3], [5]–[16] and random access policies [17]–[36]. Despite the past efforts that have significantly advanced our understanding of these policies and their performance, to the best of our knowledge, there is no instance of these policies that realizes all of the three properties mentioned earlier, even for geometric networks.

Matching policies are able to achieve throughput optimality [5]–[7], [9] and can achieve order-optimal delay [15]. However, these properties are obtained assuming that an NP-hard problem can be solved in each scheduling round. Reducing the complexity of matching policies, in general, comes at the price of losing throughput optimality [10]–[13] or a large delay [1]. The design of a matching policy that achieves all three properties given above remains an open research challenge (see Section II for further discussion on matching policies).

Random access policies are simple and can be implemented in distributed manner. Among random access policies, the classical CSMA policy is known to be throughput-optimal [24], [29]–[31]. However, as we discuss below, the delay performance of CSMA policies can be poor.

As a motivating example consider a wireless network with $L = n^2$ links, and for which the *interference graph* [2], [29]–[31] is given by an $n \times n$ torus (see Fig. 1). Furthermore, assume that all traffic is one-hop traffic and that the link packet arrival rates are uniform, i.e., the packet arrival rate to each link is equal to λ. Let ρ be the corresponding load[1], and define

$$\epsilon = 1 - \rho.$$

For this simple topology, a mixing-time analysis [31] upper-bounds the one-hop packet delay under the classical CSMA policy as

$$O\Big(\Big[\frac{1}{\epsilon}\Big]^{c_u L}\Big),$$

where $c_u > 1$ is a constant. For small ϵ, a similar analysis [37] lowerbounds the one-hop packet delay under the classical CSMA policy as

$$\Omega(e^{c_l L/(log L)^2}),$$

for some constant $c_l > 0$.

The above delay-bounds show that the classical CSMA policy exhibits a *threshold behaviour* in the sense in order to achieve a high throughput, i.e., to make ϵ small, one has to tolerate a delay that *exponentially grows* with the network-size L. The threshold behaviour and the exponential growth are related to the *phase transition* phenomenon[2] in the hard-core lattice gas model [38], [39]. Due to such threshold behaviours, even in mid-sized simple topologies, the classical CSMA policy cannot support a high throughput with low delay (see Section IV-A).

In this paper, we propose Unlocking CSMA (U-CSMA) as a new CSMA policy that overcomes the threshold behaviour

[1]In the limit of large torus, the maximum uniform throughput is 0.5, and load ρ in the limit becomes $\frac{\lambda}{0.5}$. See Section III-C for the definition of ρ.

[2]Phase transition has also been reported as the cause of border effects that persist in 2D under the classical CSMA policy [27].

978-1-4244-9921-2/11/$26.00 ©2011 Crown

of the classical CSMA policy. While being simple and distributed, U-CSMA has the following properties for geometric networks with one-hop traffic [2], [3].

a) It enables to achieve a high throughout/utility arbitrarily close to the optimal with a low packet delay.
b) The packet delay under this policy is order-optimal, i.e., it stays bounded as the network-size increases to infinity.

We provide analytical results for the torus interference topology with uniform packet arrival rate as considered earlier, and show that for large network-size L, the delay under U-CSMA is order-optimal and is

$$O\Big(\Big[\frac{1}{\epsilon}\Big]^3\Big). \quad (1)$$

It is important to note that the above delay bound is independent of the network size L, in sheer contrast to the delay under the classical CSMA policy that exponentially increases with the network-size L. This means that U-CSMA does not suffer from the threshold behaviour and is indeed able to provide high throughput with low delay for arbitrarily large torus topologies.

In our simulation study, we use U-CSMA jointly with a congestion control algorithm to maximize a network-wide utility in large random geometric networks. We show that using U-CSMA, we can assign packet arrival rates closely to the optimal with a low packet delay that stays bounded as the network-size increases, and hence, a delay that exhibits order-optimality. As far as we are aware, it is for the first time that a simple distributed scheduling policy is proposed that can operate close to the optimal with order-optimal low packet delay.

We believe that the design principle of U-CSMA and the novel approach taken to study its performance open up a new direction into the design and study of scheduling policies for large-scale wireless networks. The main significance of our study in this paper is that it realizes the possibility of having large-scale wireless multihop networks that can be maintained in a simple distributed manner and that can provide high throughput/utility, arbitrarily close to the optimal, with order-optimal low packet delay.

A key step to obtain the delay bound in (1) is where we show that the schedule under the classical CSMA policy quickly converges to a maximum schedule in geometric networks. Using techniques from mean field theory [40], we show that for large torus and lattice topologies with large uniform attempt-rates, the *distance* (see Section V-A) to the maximum schedules as a function of time t drops as $\frac{1}{\sqrt{t}}$. To the best of our knowledge, our result is the first that analytically characterizes the fast convergence behaviour of the classical CSMA policy. As this convergence is independent of network-size L, it is fundamentally different than the convergence time to the steady-state (i.e., the mixing time) of the dynamics of the classical CSMA policy, which can be exponentially large in L [37].

The rest of the paper is organized as follows. In the next section, we briefly review the related work. In Section III, we present the network model and the classical CSMA policy model. In Section IV, we provide an overview of our main results, including the description of U-CSMA and simulation results. In section V, we provide a formal statement of our analytical results in this paper.

II. Related Work

In this section, we provide a brief, by no means exhaustive, overview of the work in the area of wireless scheduling that is closest to ours in this paper. We consider two main classes, i.e., the matching policies and random access policies.

Matching Policies: Maximum Weight Matching (MWM) policy was first proposed in the seminal work in [5]. This policy is perhaps the first policy that is throughput-optimal in a wide range of settings [5]–[7], [9]. MWM policy at any time slot maximizes a weighted summation of queue-sizes in the network, which can be an NP-hard optimization problem [2]. Despite its complexity, simulations [16] show that MWM policy is close to the optimal in terms of delay for one-hop traffic. For multihop traffic, the delay under MWM policy is $O(\frac{L}{\epsilon})$, and for one-hop traffic is order-optimal as $O(\frac{1}{\epsilon})$, under certain conditions [15] that hold for geometric networks. The delay bound in our paper for one-hop traffic is $O([\frac{1}{\epsilon}]^3)$, which includes a multiplicative factor of $[\frac{1}{\epsilon}]^2$ as well as $\frac{1}{\epsilon}$. This factor can be interpreted as the scheduling-time needed to find schedules that are ϵ close to the optimality. However, we note that the delay performance in [16] and the $O(\frac{1}{\epsilon})$ bound in [15] are obtained assuming that the NP-hard problem of MWM policy can be solved at every time slot.

Greedy Maximal Matching (GMM) policy is a simple and distributed alternative for MWM policy, see e.g., [2], [11]. While GMM policy is not throughput-optimal in general, a number of local pooling results [3], [10], [13] indicate that for a noticeable subset of topologies, GMM policy is indeed throughput-optimal. However, GMM requires message passing, and it is an open area to investigate the delay performance of GMM policy. Maximal Matching (MM) policy is simpler than GMM policy and has order-optimal delay of $O(\frac{1}{\epsilon})$ for one-hop traffic [12]. However, this policy is not throughput-optimal and is guaranteed to stabilize only half of the capacity region. See [14] for a comparison of different matching policies.

Random Access Policies: Random access policies started with the classical Aloha protocol [17], for which an optimality result was first established in [18]. The capacity of random access policies under collision detections, acknowledgements, or backoff schemes have been studied in [19], [20], [22]. The recent work in [26] chooses access probabilities in an Aloha-like policy based on queue backlogs to achieve the capacity region of slotted Aloha. In [25], [33], distributed protocols are proposed that assign access probabilities to maximize a network utility under an Aloha-like protocol. Due to their simplicity, Aloha-like protocols have been also used in mobile networks [23]. These protocols however are not throughput-optimal [26].

CSMA policies are a special class of random access policies that assume nodes can sense whether their neighbours are transmitting. Performance of these policies as defined in 802.11 standard for a specific network setup is studied in [21]. For a special class of networks with primary interference, it is known that 1) CSMA polices are throughput-optimal [24], and 2) for a subclass of these networks such as the $n \times n$ switch, the delay to access the channel becomes memoryless under CSMA policies, leading to an $O(\frac{1}{\epsilon})$ (normalized) packet delay [32].

Throughput-optimality of CSMA policies extends to networks with arbitrary interference graphs [29]–[31]. The throughput-optimal CSMA policies in [29]–[31] are based on a continuous time Markov chain that prevents collisions. This is addressed by considering contention resolution [30], [34].

Both in [29] and [30], it is assumed that there is a time-scale separation and, hence, CSMA dynamics quickly converges to its steady-state faster than the rate by which queues change over time. The authors of [31] and later those of [36] show that as long as attempt rates of nodes change sufficiently slowly, throughput optimality can be achieved. A related work [35] divides the time axis into frames, and updates parameters of CSMA policy only at the beginning of each frame. However, delay performance under the above throughput-optimal schemes is not investigated, and the upperbound on the delay inferred from these papers increases with the network-size.

Before concluding this section, we note that there are numerous results that study link starvation under CSMA policies, e.g., see [28] and references therein. In particular, the work in [27] shows that in 2D, the phase transition phenomenon makes the CSMA policy *lock into* a certain similar set of states for a long time, causing large packet delays. Using this insight, we propose U-CSMA which has a simple and distributed implementation, and provides both high throughput and a low packet delay.

III. Network and Classical CSMA Policy Model

In this section, we introduce the network and classical CSMA policy model that we use in this paper.

A. Network Model

We consider a fixed wireless network consisting of a set $\mathcal{N}$ of nodes, and a set $\mathcal{L}$ of links with cardinality L. We refer to L as the *network size*. A link $l = (n, m) \in \mathcal{L}$ indicates that transmitter node n and receiver node m are within transmission range of each other and can exchange data packets. Each link $l = (n, m)$ corresponds to a *queue* that is maintained by its transmitter node n.

We model the contention between links by an *interference graph* $G(\mathcal{L}, \mathcal{E})$ [2], [29]–[31], [35], 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. We define a *geometric interference graph* [2]–[4] to be a graph whose vertices can be considered as points on the plane, and where two vertices are connected by an edge if and only if the distance between them is less than the *interference range* r where $r > 0$. We define a *geometric network* as a network with geometric interference graph. We define a *random geometric network* as a geometric network for which the vertices of its interference graph are points that are distributed according to a uniform stochastic process over a convex region in the plane.

We define a *valid schedule* to be a subset of links in $\mathcal{L}$ no two of which interfere with each other. We define a *maximum schedule* to be a valid schedule with the largest number of links in $\mathcal{L}$. We also define a link to be *active* at time t, if the link is transmitting at time t. We define a *scheduling policy* to be an algorithm, randomized or deterministic, that determines which links are active at any given time.

Throughout the paper, we assume that traffic is one-hop. Let λ_l be the packet arrival rate for transmission over link l, which corresponds to a queue in the network, and let

$$\boldsymbol{\lambda} = (\lambda_l)_{l \in \mathcal{L}}$$

be the arrival rate vector for a given network. We assume arrivals are i.i.d so that every unit of time, one packet arrives to link l, $l \in \mathcal{L}$, with probability λ independent of any other arrival event in the network. Extension to non i.i.d arrivals is provided in the technical report [41]. Finally, we assume that the rate of transmission is the same for all links, and it takes one unit of time to transmit any one packet.

B. Classical CSMA Policy

For our analysis, we define the classical CSMA policy as follows, similar to the ones presented in [27], [29]–[31]. 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 $\mathcal{E}$. A link l senses the channel as *idle* at time t if all of its conflicting (interfering) links are not active and not transmitting at time t. If link l senses that any of its interfering links is transmitting, then it waits until all of its interfering links become silent. Once this happens, link l sets a backoff timer with a value that is exponentially distributed with mean $1/z_l$, $z_l > 0$, 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 transmission. Otherwise, link l simply waits until all of its interfering links become silent again, and repeats the above process. We define z_l to be the transmission *attempt-rate* of link l. We assume that all transmission times are independently and exponentially distributed with unit mean.

The above models an *idealized CSMA policy* in which 1) any link can always sense transmissions of all of its interfering links, and 2) there is no hidden-terminal problem that can create packet collisions as in [27], [29], [31]. These assumptions can be removed using the methods of [30], [34]. Hence, we continue assuming that the above two assumptions hold.

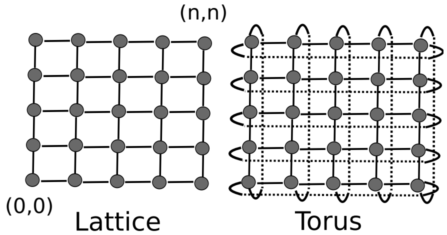

Fig. 1. Lattice and torus interference graphs. Each dark circle represents a link, and an edge between two dark circles shows that their corresponding links interfere with each other.

We characterize a classical CSMA policy by the vector $\mathbf{z} = (z_l)_{l \in \mathcal{L}}$ where z_l is the transmission attempt-rate of link l. Given vector $\mathbf{z}$, the network dynamics as which links are active over time can be represented by a Markov process [29]. Using this, we can define $\mu_l(\mathbf{z})$, $l \in \mathcal{L}$, as the service rate of link l under $\mathbf{z}$, i.e., $\mu_l(\mathbf{z})$ is the fraction of time that link l is active under the CSMA policy $\mathbf{z}$.

We say that the classical CSMA policy $\mathbf{z}$ *stabilizes* the network for a given packet arrival rate vector $\boldsymbol{\lambda}$ if [5]

$$\lambda_l < \mu_l(\mathbf{z}), \qquad l \in \mathcal{L}. \tag{2}$$

It is well-known that the classical CSMA policy is *throughput optimal* [24], [29]–[31], i.e., if there exists some policy, CSMA or not, that can stabilize the network for a given $\boldsymbol{\lambda}$, then there exists a classical CSMA policy $\mathbf{z}$ that stabilizes the network for $\boldsymbol{\lambda}$.

C. Lattice and Torus Interference Graphs with Uniform Attempt and Packet Arrival Rates

To obtain analytical results, we consider wireless networks with grid-like interference graphs. In particular, we consider the *lattice interference graph* $G_L = G_L(\mathcal{L}, \mathcal{E})$ and the *torus interference graph* $\mathcal{T}_L = \mathcal{T}_L(\mathcal{L}, \mathcal{E}_\mathcal{T})$. In both cases, the set $\mathcal{L}$ is the set of all links where each link $l \in \mathcal{L}$ can be represented by coordinates (i,j), $i,j \in \{0,...,n\}$, on the plane. See Fig. 1 for an illustration. Hence, the network-size, i.e., the total number of links, is given by $L = (n+1)^2$.

It remains to specify which links interfere with each other. For the lattice interference graph G_L, we assume that there exists an edge $e \in \mathcal{E}$ between any two links $l = (i,j)$ and $l' = (i',j')$, $l, l' \in \mathcal{L}$, iff link l and link l' differ in exactly one coordinate, i.e., we have that

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

For the torus interference graph $\mathcal{T}_L$, the edge set $\mathcal{E}_\mathcal{T}$ contains all edges defined for the lattice interference graph G_L. In addition, the set $\mathcal{E}_\mathcal{T}$ contains an edge between link $l = (i,0)$ and link $l' = (i,n)$, for $0 \le i \le n$, and also contains an edge between link $l = (0,j)$ and link $l' = (n,j)$, for $0 \le j \le n$. As a result, the torus interference graph $\mathcal{T}_L$ is the same as G_L with additional edges around the boundary of G_L so that every link has exactly four interfering links.

Given a lattice or torus interference graph, 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.

For the lattice and torus interference graphs G_L and $\mathcal{T}_L$, we focus on CSMA policies $\{\mathbf{z}\}$ with uniform transmission attempt-rates so that

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

for some $z > 0$. In addition, we focus on the case of uniform packet arrival rates, i.e., we let

$$\lambda_l = \lambda, \qquad 0 < \lambda < \mu_{max}(L), \ l \in \mathcal{L}. \tag{3}$$

where $\mu_{max}(L)$ is the *maximum uniform-throughput*, i.e., the maximum throughput that can be provided for *all* links by any policy in the network. For lattice interference graph G_L, we have that

$$\mu_{max}(L) = 0.5.$$

This throughput can be achieved, for instance, by alternating between two valid schedules $\mathcal{L}^{(o)}$ and $\mathcal{L}^{(e)}$ every unit of time, which allows every link to be active half of the time. For torus graph $\mathcal{T}_L$, due to boundaries being wrapped around, $\mathcal{L}^{(o)}$ and $\mathcal{L}^{(e)}$ are not valid schedules, but we can show that

$$\lim_{L \to \infty} \mu_{max}(L) = 0.5.$$

Having defined $\mu_{max}(L)$, for a given lattice or torus interference graph with L links, we define the network *load factor* or simply *load* ρ as

$$\rho = \rho(\lambda) = \frac{\lambda}{\mu_{max}(L)}. \tag{4}$$

We also define ϵ to be the distance to maximum load of $\rho = 1$:

$$\epsilon = \epsilon(\lambda) = 1 - \rho(\lambda). \tag{5}$$

We next provide an overview of our main results.

IV. Overview of Main Results

In this section, we provide an overview of our main results. We first investigate the performance of the classical CSMA policy as defined in Section III-B, and explain why under this policy it is impractical to obtain both high throughput and low delay. We then propose and describe a novel CSMA policy called U-CSMA. We show that for geometric networks with one-hop traffic, U-CSMA policy overcomes the shortcomings of the classical CSMA policy and allows to obtain high throughput or utility, arbitrarily close to the optimal, with low packet delay that is order-optimal, i.e., stays bounded as the network-size increases to infinity.

A. Performance of Classical CSMA Policy

Consider a fixed wireless network with torus interference graph, as defined in Section III-C, having L links and a uniform packet arrival rate λ to each link, as defined in (3). It is well-known that [38] if all links use the same rate z, then the following holds for the achieved uniform throughput $\mu(z, L)$:

$$\mu_{max}(L) - \mu(z, L) = \boldsymbol{\Theta}(z^{-1}). \tag{6}$$

This means that to be $\Theta(\epsilon)$ away from the maximum uniform throughput $\mu_{max}(L)$, an attempt rate z of order $\frac{1}{\epsilon}$ is needed.

For the above network, two threshold behaviours exist, as explained in the following.

Threshold Behaviour as a Function of Attempt-rate z: It is well-known that for a fixed network size L, as the attempt rate z increases beyond a threshold, the delay of classical CSMA policy on the torus interference graph increases substantially. This increase is related to a *phase transition* phenomenon, in terms of the existence of more than one Gibbs measures for the infinite torus [39].

The currently best explicit characterization of the delay of the classical CSMA policy in terms of z shows that the delay is (see, e.g., the mixing time analysis in [31])

$$O(z^{c_u L}), \tag{7}$$

for some constant $c_u > 1$. While for $z < 1$, the above bound can be moderate for a moderate network size L, for $z > 1$, there will a rapid increase even for moderate values of L. Since by (6), a large attempt-rate is needed to support a high throughput, this explains why the classical CSMA policy cannot provide high throughput without incurring a large delay.

We note that by (6), the classical CSMA policy needs to use an attempt rate of order $1/\epsilon$ to support the load $\rho = 1-\epsilon$, which can be used to write the delay bound in (7) as

$$O\Big(\Big[\frac{1}{\epsilon}\Big]^{c_u L}\Big). \tag{8}$$

Threshold Behaviour as a Function of Network-size L: Depending on the value of a given attempt z, as we increase the network size L, the delay of the classical CSMA policy shows an undesirable threshold behaviour.

On one hand, there exists a constant $z_{c,1} > 0$ such that for all attempt-rates $z < z_{c,1}$, the delay is upperbounded as [42]

$$O(\log(L)). \tag{9}$$

This bound states that for low attempt rates resulting in low uniform-throughput, the delay increases only logarithmically in the network size L.

On the other hand, there exists a constant $z_{c,2} > 0$ such for any attempt-rate $z > z_{c,2}$, the delay is lowerbounded as [37]

$$\Omega(e^{c_l L/(\log L)^2}), \tag{10}$$

for some constant $c_l > 0$. Hence, for large attempt-rates required to support high throughput, the delay grows exponentially with the network-size L, which results in a threshold behaviour as L increases. It is this exponential increase in the delay that prevents the classical CSMA policy to provide high throughput with low packet delay as the network-size L increases.

Simulation: To illustrate the threshold behaviours, we have simulated a torus of size $L \in \{100, 400, 1600\}$ under the classical CSMA policy with uniform attempt rate z.

For a given network size L, to support the uniform arrival rate λ (see Section III-C) where

$$\lambda = (1-\epsilon)\mu_{max}(L), \qquad \epsilon > 0, \tag{11}$$

and consequently a load factor (as defined in (4)) of $\rho = (1-\epsilon)$, we have chosen the attempt rate z such that the resulting uniform throughput $\mu(z, L)$ is given by

$$\mu(z, L) = \mu_{max}(L)(1 - \frac{\epsilon}{2}) > \lambda. \tag{12}$$

Fig. 2(a) shows the resulting average queue size per link as a function of ρ in linear scale. This figure clearly illustrates the two threshold behaviours.

First, we see that for a given network-size L, for a small load ρ less than 0.3, the queue-sizes are small. However, as the load ρ increases towards 0.5, which requires a larger attempt-rate z, the queue-size increases from only few packets to thousands. While the classical CSMA policy is throughput-optimal and in principle can support a load ρ close to 1, we see that in practise, it cannot support loads as low as 0.5, i.e., it cannot reach the 50% utilization without incurring a large delay. For instance, for the 20×20 torus, the large delay becomes more than 1sec for a packet length of 2346 bytes and a channel rate of 54Mbs as in 802.11 standards.

Second, we see that for a given ρ, the queue-size shows two different behaviours. If $\rho < 0.4$, the queue-size is small and hardly changes with the network size. In contrast, for $\rho > 0.4$, the queue-size shows a threshold behaviour and drastically and exponentially increases with the network size. For instance, at $\rho = 0.44$, the queue-size almost doubles every time that the network size L increases by a factor of 4.

Intuition: By (6), in order to support a high uniform throughput, the classical CSMA policy needs to use a large attempt rate z. For a large attempt rate z, the network state will mainly alternate between two types of transmission patterns (valid schedules) 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 active (see Section III-C). However, as z and L increase, transitions between these two types of patterns occur very infrequently. This implies that the classical CSMA policy tends to *lock into* one type of transmission patterns for a very long time before it switches to the other type of patterns [39].

This *locking-in* behaviour of the CSMA policy immediately implies that while one type of links, e.g., even links, are active for a long time, the other type of links, e.g., odd links, cannot transmit for a long time. As a result, this locking-in behaviour leads to large queue-sizes and hence a large packet delay.

We next describe U-CSMA and provide theoretical and simulation results characterizing its performance.

B. U-CSMA and Its Performance

The main contribution of this paper is to propose U-CSMA that overcomes the threshold behaviours faced by the classical CSMA policy.

U-CSMA: The basic idea behind our proposed U-CSMA policy is very simple. U-CSMA uses a classical CSMA policy

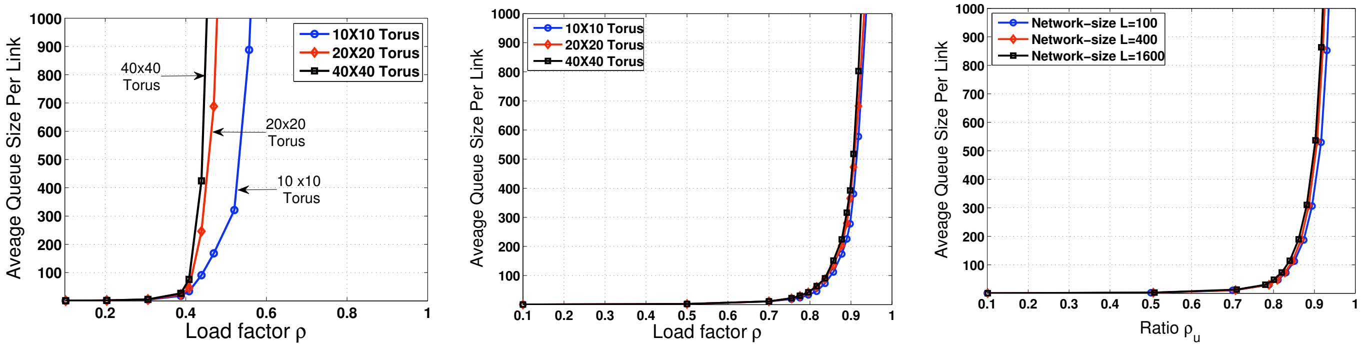

(a) Illustration of the threshold behaviours under classical CSMA policy for torus interference graph.

(b) Illustration of elimination of the threshold behaviours under U-CSMA for the torus interference graph.

(c) Performance under U-CSMA policy combined with congestion control in random geometric interference graphs, as a function of utility ratio ρ_u.

Fig. 2. Illustration of performance under the classical CSMA policy and U-CSMA.

z as described in Section III-B. However, periodically, i.e., at times

$$T_i = iT, \ i \in \{0, 1, 2, \cdots\},$$

U-CSMA resets, or *unlocks*, the transmission pattern of the classical CSMA policy by requiring all links to become silent, and then immediately restarts the classical CSMA protocol to operate as usual. In the rest, we refer to parameter T as the *unlocking period*. We note the in the limit of large T, U-CSMA reduces to the throughput-optimal classical CSMA policy.

The intuition behind the above unlocking mechanism is to prevent the threshold behaviour by preventing the policy from locking into a particular transmission pattern for too long. In technical report [41], we provide one approach to implement U-CSMA in a fully distributed and asynchronous manner.

Analytical Results: In order to characterize the performance of U-CSMA, we first need to know how to choose the unlocking period T. While a smaller T helps employ the unlocking mechanism more frequently leading to a smaller delay, it may also prevent the underlying classical CSMA policy used by U-CSMA from converging to a maximum schedule that is necessary to obtain a high throughput. Hence, as the first step, we need to study how fast the classical CSMA policy converges to a maximum schedule.

Our first analytical result (see Proposition 1 in Section V) shows that for the lattice and torus interference graphs with uniform attempt rate z, valid schedules under the classical CSMA policy quickly converge to a maximum schedule at a rate that becomes independent of network-size L for large networks and attempt-rates. Remarkably, this result shows that the *distance* to the maximum schedules roughly drops as

$$\frac{1}{\sqrt{t}}.$$

Our second analytical result (see Proposition 2 in SectionV-B) uses the above convergence result to stabilize networks with torus interference topology and uniform packet arrival rate λ. In particular this result shows that U-CSMA with unlocking period

$$T(\epsilon) = \mathbf{\Theta}\Big(\Big[\frac{1}{\epsilon}\Big]^2\Big), \tag{13}$$

and with large uniform attempt rate[3] z stabilizes the load $\rho = (1 - \epsilon)$ for large networks with torus interference graph. Hence, by the above choice for the unlocking period, U-CSMA stabilizes queues in the network, all of which have packet arrival rate of $\lambda = (1 - \epsilon)\mu_{max}(L)$.

Further, this result shows that by the above choice for the unlocking period $T(\epsilon)$, the average queue-size per link and, hence, average delay become order-optimal and independent of the network size L in the sense that for large L and attempt-rate z, they are upperbounded as

$$O\Big(\Big[\frac{1}{\epsilon}\Big]^3\Big). \tag{14}$$

Comparing the above delay bound with the ones in (8) and (10) for the classical CSMA policy, we see that U-CSMA does not suffer from the threshold behaviours. Specifically, we see that as a function of $1/\epsilon$, the queue-size under U-CSMA increases at most with exponent 3 as opposed to the exponent L under classical CSMA policy, as suggested by the bound in (8). Moreover, U-CSMA has changed a queue-size that exponentially grows with the network size L (see (10)) to a queue-size that does not depend on the network size L.

Simulation Results: To illustrate the performance of U-CSMA and compare it with the analytical results, we have simulated a torus of size $L \in \{100, 400, 1600\}$ under U-CSMA. We have set the uniform attempt rate at $z = 50$, and for a given uniform arrival rate

$$\lambda = (1 - \epsilon)\mu_{max}(L),$$

or load $\rho = 1 - \epsilon$, we have chosen the unlocking period T as

$$T = \frac{1.2}{\epsilon^2}. \tag{15}$$

Fig. 2(b) shows the resulting queue-sizes as a function of load ρ. We make the following two observations. First, comparing Fig. 2(b) with Fig. 2(a), we see that while the classical CSMA "hits the wall" and its queue-size becomes on the order of thousands of packets before reaching a low load of $\rho = 0.5$, U-CSMA can indeed get much closer to the

[3]Large attempt rates can be achieved by Glauber dynamics as in [30], [31].

maximum load of 1. In practical terms, for a packet length of 2346 bytes and a channel rate of 54Mbs as in 802.11 standards, the average packet delay under U-CSMA becomes 30ms and 90ms for 80% and 85% channel utilization, respectively, while the delay under the classical CSMA becomes more than 1sec before even reaching the 50% utilization. In addition, re-plotting the queue-size as a function of $\epsilon = 1 - \rho$ in log-log scale (see Fig. 3), we see that the average exponent by which queue-size increases as a function of $1/\epsilon$ is 3.02, which closely matches the exponent 3 as predicted by the analysis in (14).

Second and as remarkably predicted by the analysis, the queue-size does not change significantly with the network size. In fact, for 20×20 and 40×40 torus the queue-sizes are hardly distinguishable. This confirms that 1) U-CSMA eliminates the threshold behaviours that exist for the classical CSMA policy, and 2) the queue-size under U-CSMA is order-optimal in that it stays bounded as the network size increases.

To investigate whether the insight gained through the analysis for the torus interference graph carries over to general network setups, we have simulated a geometric interference graph [2]–[4], see Section III-A, in which $L \in \{100, 400, 1600\}$ links are randomly distributed over a square area of 10×10, 20×20, and 40×40, respectively. We have chosen the interference range r so that every link on the average interferes with six other links. As in [8], [11], [33], we have implemented a congestion control algorithm to tune the arrival rate to each link so that a *network-wide* logarithmic utility function U_{net} is maximized. This algorithm operates on top of U-CSMA (see technical report [41] for further details).

Fig. 2(c) plots the average queue-size as a function of ρ_u

$$\rho_u = \frac{U_{net}}{U_{opt}},$$

i.e., ρ_u is the ratio of the achieved network-wide utility to the optimum maximal utility U_{opt}. Remarkably, the delay behaviour is similar to the one illustrated by Fig. 2(b).

The main observation here is that the insight gained through the analysis for the torus interference graph also holds for the general case considered here. First, we observe that even in random topologies under a congestion control algorithm, we can use U-CSMA to assign arrival rates closely to the optimal without incurring a large delay. For instance, for a packet length of 2346 bytes and a channel rate of 54Mbs, the delay becomes 40ms to get to 80% of optimality. Interestingly, the exponent by which queue-size increases as a function of $1/(1-\rho_u)$ approaches 3, the same exponent in the delay bound of torus graph in (14) (see the technical report [41] for the corresponding log-log plot).

Second, we observe that the queue-size and hence the delay slightly change with the network-size. This means order-optimality of delay is preserved, and therefore, we can use U-CSMA jointly with congestion control to assign arrival rates close to the optimal with low packet delay in arbitrarily large networks.

Next, we provide formal statements of our analytical results.

V. Performance Analysis

In this section, we formally state the analytical results developed in this paper for lattice and torus interference graphs. These results characterize the rate by which the schedule under classical CSMA policy converges to maximum schedules, and characterize the delay-throughput tradeoff under U-CSMA. A more detailed discussion of these results with complete proofs and further simulation results is provided in technical report [41].

The analytical results presented in this section use two assumptions on the properties of schedules under the classical CSMA policy. Due to lack of space, we are unable to state these assumptions here. A formal statement of these assumptions and a detailed discussion as why these assumption are expected to hold are provided in technical report [41]. Simulation results are provided in order to investigate whether these assumptions indeed lead to correct qualitative results, not only for lattice and torus topologies, but also for random geometric networks

A. Convergence to Maximum Schedules Under Classical CSMA Policy

Our first result characterizes the rate by which the schedule under classical CSMA policy converges to maximum schedules. We consider the lattice or torus interference graph with L links, and a classical CSMA policy with uniform attempts rate z, as described in Section III.

To state our first result, we use the following notation. Let $\theta_L(t, z)$ be the *density*, i.e., fraction, of links that are active at time t, $t > 0$. Hence, if $N_a(t, z)$ is the total number of links that are active at the time t under a classical CSMA policy with uniform attempt rate z, then $\theta_L(t, z)$ is given by

$$\theta_L(t, z) = \frac{N_a(t, z)}{L}$$

We assume that the system is idle at time $t = 0$ such that

$$\theta_L(0, z) = 0, \qquad z > 0.$$

Let $\delta_L(t, z)$ be

$$\delta_L(t, z) = 0.5 - \theta_L(t, z). \tag{16}$$

Since 0.5 is the fraction of links that can be active under a maximum schedule in lattice or torus interference graphs in the limit of large L, we see that $\delta_L(t, z)$ can represent the *distance* between the schedule at time t and the limit maximum schedules.

Proposition 1 characterizes how fast the distance $\delta_L(t, z)$ approaches 0, or in other words, how fast the distance to maximum schedules drops to 0, in the limit of large L and z.

Proposition 1. *Suppose the interference graph is given by the lattice (or torus) interference graph G_L (or $\mathcal{T}_L$). Under Assumptions 1-2 in [41], there exists a positive constant C_1, independent of z and L, such that for any $\tau > 0$, we have that*

$$\liminf_{z\to\infty} \liminf_{L\to\infty} P\left[\sup_{t\in(0,\tau]}\left[\delta_L(t, z) - \frac{C_1}{\sqrt{t}}\right] \leq 0\right] = 1.$$

Proof: Proof is provided in technical report [41]. ■

Proposition 1 states that for every finite time-horizon $(0, \tau]$, with probability approaching one as first the network size L approaches infinity and then z approaches infinity, the distance $\delta_L(t, z)$ between $\theta_L(t, z)$ and the maximum fraction of active links 0.5 converges to 0 and drops as $O(\frac{1}{\sqrt{t}})$ for $t \in (0, \tau]$.

The above convergence has two important implications. First, under the classical CSMA policy, the distance to maximum schedules asymptotically drops as $O(\frac{1}{\sqrt{t}})$, only depending on time t. Second, as the $O(\frac{1}{\sqrt{t}})$ bound does not depend on the network-size L or attempt-rate z, the convergence is not negatively affected by a large L or large z. This is in a stark contrast to the results obtained for for the mixing time of CSMA policies, i.e., the rate at which CSMA policies reach their steady-state, which increases with attempt-rate z and can be exponential in the network size L [37].

B. Delay-Throughput Trade-off of U-CSMA

Proposition 1 states that under classical CSMA policy, the distance to maximum schedules converges to zero at a rate independent of network size in the limit of large network sizes and attempt rates. Our second result stated in Proposition 2 characterizes the delay-throughput trade-off under U-CSMA for the torus interference graph with uniform attempt-rate z (see Section III-C). Intuitively, Proposition 2 states that in large networks, the delay-throughput trade-off under U-CSMA does not depend on the network-size L.

In order to formally state the throughput-delay trade-off for any given link in the network, irrespective of its position, we consider the torus interference graph $\mathcal{T}_L$ (see Section III-C) instead of the lattice interference graph G_L. For the lattice interference graph and similar topologies, it is well known that due to boundary effects, the throughput achieved by links in the network is not uniform over all links in the network when a uniform attempt rate z is used [27]. The torus interference graph is symmetric with respect to link positions, and as a result boundary effects do not exist. While we develop the analysis for the torus interference graph, the general insight gained through the analysis carries over to more general settings, as discussed in Section IV-B

To state Proposition 2, we introduce several definitions. We first note that by Proposition 1, for the torus interference graph $\mathcal{T}_L$ and a given $\tau > 0$, we can define a non-negative function $\epsilon_p(L, z, \tau)$ such that we have

$$P\left[\sup_{t\in(0,\tau]}\left[\delta_L(t,z) - \frac{C_1}{\sqrt{t}}\right] \le 0\right] \ge 1 - \epsilon_p(L, z, \tau), \quad (17)$$

and

$$\limsup_{z\to\infty}\limsup_{L\to\infty}\epsilon_p(L, z, \tau) = 0. \quad (18)$$

For a given $\epsilon' > 0$, the above limit allows us to define $z(\epsilon', \tau)$ and $L(z, \epsilon', \tau)$ such that for $z > z(\epsilon', \tau)$ and $L > L(z, \epsilon', \tau)$, we have

$$\epsilon_p(L, z, \tau) < \frac{1}{2}\epsilon'. \quad (19)$$

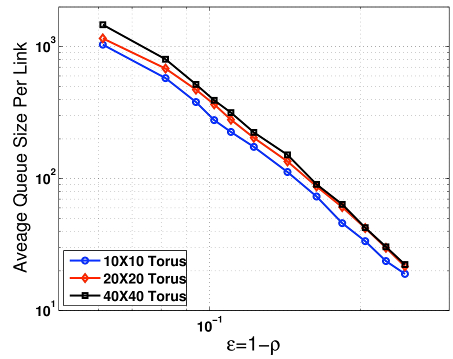

Fig. 3. Average queue-size as a function of distance ϵ to the maximum load $\rho = 1$, under U-CSMA for the torus interference graph.

Furthermore, for a given uniform packet arrival-rate λ, $0 < \lambda < 0.5$, and a given uniform attempt-rate z (see Section III-C), we define $Q_l(t, z, \lambda)$ as the queue size of link l at time t.

Proposition 2 is given as follows.

Proposition 2. *Consider the torus interference graph $\mathcal{T}_L$, and suppose Assumptions 1-2 in [41] hold. Let the uniform packet arrival rate to each link be λ where $0 < \lambda < 0.5$. Let the unlocking period $T(\lambda)$ used by U-CSMA be*

$$T(\lambda) = \frac{(16C_1)^2}{\epsilon^2} = \mathbf{\Theta}\left(\frac{1}{\epsilon^2}\right)$$

where $$\epsilon = \epsilon(\lambda) = 1 - \rho(\lambda),$$

and C_1 is a constant given in Proposition 1. Then, there exists a positive constant C_2 such that for $z > z(\epsilon, T(\lambda))$ and $L > L(z, \epsilon, T(\lambda))$, the time average of the queue size for any link l in $\mathcal{T}_L$ satisfies the following under U-CSMA with the unlocking period $T(\lambda)$:

$$\limsup_{t\to\infty}\mathbb{E}\left[\frac{1}{t}\int_0^t Q_l(t, z, \lambda)\ dt\right] < \frac{C_2}{\epsilon^3} = \mathbf{\Theta}\left(\frac{1}{\epsilon^3}\right)$$

Proof: Proof is provided in technical report [41]. ■

Proposition 2 states that in order to get ϵ close to the maximum load of $\rho = 1$, the expected time average of any queue-size in the network becomes only $O(\frac{1}{\epsilon^3})$, independent of network-size L for large L. 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 delay for any given link is also

$$O\left(\left[\frac{1}{\epsilon}\right]^3\right). \quad (20)$$

Quite surprisingly, the delay bound and the resulting throughput-delay trade-off are valid for arbitrarily large torus networks as long as $z > z(\epsilon, T(\lambda))$. Moreover, since C_2 in the proposition is a constant, the delay bound does not depend on the network-size L, and hence, we have an order-optimal delay. This makes the delay-throughput trade-off under U-CSMA independent of the network-size L for large L. As a result, U-CSMA can indeed provide high throughput with low delay for arbitrarily large wireless networks for which the interference graph is given by a torus.

To investigate the accuracy of the delay bound in (20), we have replotted the queue-size as a function of ϵ under the simulation setup of Section IV-B. The figure shows that the queue-size increases with (average) slop 3.02 in $\log$-$\log$ scale, which, as expected, is close to the exponent 3 given in (20)[4].

REFERENCES

[1] D. Shah, D. N. C. Tse, and J. N. Tsitsiklis, "Hardness of low delay network scheduling," 2009. [Online]. Available: http://www.mit.edu/∼devavrat/harddelay.pdf

[2] G. Sharma, R. R. Mazumdar, and N. B. Shroff, "On the complexity of scheduling in wireless networks," in *Proc. of the 12th annual international conference on Mobile computing and networking (MobiCom'06)*, 2006, pp. 227–238.

[3] C. Joo, X. Lin, and N. Shroff, "Understanding the capacity region of the greedy maximal scheduling algorithm in multi-hop wireless networks," in *Proc. IEEE INFOCOM'08*, Apr. 2008.

[4] H. B. Hunt, M. V. Marathe, V. Radhakrishnan, S. S. Ravi, D. J. Rosenkrantz, and R. E. Stearns, "NC-Approximation schemes for NP- and PSPACE-hard problems for geometric graphs,," *Journal of Algorithms*, vol. 26, no. 2, pp. 238 – 274, 1998.

[5] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.

[6] R. Buche and H. J. Kushner, "Control of mobile communication systems with time-varying channels via stability methods," *IEEE Trans. Autom. Control*, vol. 49, no. 11, pp. 1954–1962, Nov. 2004.

[7] A. Eryilmaz, R. Srikant, and J. Perkins, "Stable scheduling policies for fading wireless channels," *IEEE/ACM Trans. Netw.*, vol. 13, no. 2, pp. 411–424, Apr. 2005.

[8] M. J. Neely, E. Modiano, and C.-P. Li, "Fairness and optimal stochastic control for heterogeneous networks," in *Proc. IEEE INFOCOM'05*, vol. 3, Mar. 2005, pp. 1723–1734.

[9] M. Neely, E. Modiano, and C. Rohrs, "Dynamic power allocation and routing for time-varying wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, pp. 89–103, Jan. 2005.

[10] A. Dimakis and J. Walrand, "Sufficient conditions for stability of longest-queue-first scheduling: second-order properties using fluid limits," *Adv. in Appl. Probab.*, vol. 38, no. 2, pp. 505–21, 2006.

[11] 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.

[12] M. J. Neely, "Delay analysis for maximal scheduling in wireless networks with bursty traffic," in *Proc. IEEE INFOCOM'08*, Apr. 2008.

[13] G. Zussman, A. Brzezinski, and E. Modiano, "Multihop local pooling for distributed throughput maximization in wireless networks," in *Proc. IEEE INFOCOM 2008.*, 2008, pp. 1139 –1147.

[14] Y. Yi, A. Proutière, and M. Chiang, "Complexity in wireless scheduling: impact and tradeoffs," in *ACM MobiHoc'08*, 2008.

[15] L. B. Le, K. Jagannathan, and E. Modiano, "Delay analysis of maximum weight scheduling in wireless ad hoc networks," in *43rd Annual Conference on Information Sciences and Systems, CISS'09*, March 2009, pp. 389–394.

[16] G. R. Gupta and N. B. Shroff, "Delay analysis for wireless networks with single hop traffic and general interference constraints," *IEEE/ACM Trans. Netw.*, vol. 18, no. 2, pp. 393–405, 2010.

[17] N. Abramson, "THE ALOHA SYSTEM: another alternative for computer communications," in *AFIPS '70 (Fall): Proceedings of the November 17-19, 1970, fall joint computer conference*. ACM, 1970, pp. 281–285.

[18] B. Hajek and T. van Loon, "Decentralized dynamic control of a multiaccess broadcast channel," *IEEE Trans. Autom. Control*, vol. 27, no. 3, pp. 559 – 569, jun 1982.

[19] F. P. Kelly and I. M. MacPhee, "The number of packets transmitted by collision detect random access schemes," *Ann. Probab*, vol. 15, no. 4, pp. 1557–1568, 1987.

[4]As mentioned in Section IV, we have also observed an exponent close to 3 when queue-size is plotted against $1 - \rho_u$ for the case where U-CSMA is combined with congestion control in random geometric topologies, implying that a variant of Proposition 2 should likely be true for this case.

[20] J. Hastad, T. Leighton, and B. Rogoff, "Analysis of backoff protocols for mulitiple access channels," *SIAM J. Comput.*, vol. 25, no. 4, 1996.

[21] G. Bianchi, "Performance analysis of the ieee 802.11 distributed coordination function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535–547, mar 2000.

[22] L. A. Goldberg, M. Jerrum, S. Kannan, and M. S. Paterson, "A bound on the capacity of backoff and acknowledgement-based protocols," Research Report 365, Coventry, UK, UK, Tech. Rep., 2000.

[23] F. Baccelli, B. Baszczyszyn, and P. Mhlethaler, "An Aloha protocol for multihop mobile wireless networks," *IEEE Trans. Inf. Theory*, vol. 52, pp. 421–436, 2006.

[24] 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.

[25] J.-W. Lee, M. Chiang, and A. Calderbank, "Utility-optimal random-access control," *IEEE Trans. Wireless Commun.*, vol. 6, no. 7, pp. 2741–2751, July 2007.

[26] A. L. Stolyar, "Dynamic distributed scheduling in random access networks," *Journal of Applied Probability*, vol. 45, no. 2, pp. 297–313, 2008.

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

[28] 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.*, vol. 16, no. 4, pp. 864–877, 2008.

[29] 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

[30] J. Ni and R. Srikant, "Distributed CSMA/CA algorithms for achieving maximum throughput in wireless networks," 2009. [Online]. Available: http://www.citebase.org/abstract?id=oai:arXiv.org:0901.2333

[31] 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.

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

[33] A. H. M. Rad, J. Huang, M. Chiang, and V. W. S. Wong, "Utility-optimal random access without message passing," *Trans. Wireless. Comm.*, vol. 8, no. 3, pp. 1073–1079, 2009.

[34] L. Jiang and J. Walrand, "Approaching throughput-optimality in a distributed CSMA algorithm: collisions and stability," in *ACM MobiHoc S3 '09*, 2009.

[35] A. Proutiere, Y. Yi, T. Lan, and M. Chiang, "Resource allocation over network dynamics without timescale separation," in *Proc. IEEE INFOCOM'10* , 2010.

[36] J. Ghaderi and R. Srikant, "On the design of efficient CSMA algorithms for wireless networks," 2010. [Online]. Available: http://arxiv.org/abs/1003.1364v1

[37] 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.

[38] 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.

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

[40] M. Benaim and J.-Y. Le Boudec, "A class of mean field interaction models for computer and communication systems," EPFL, Tech. Rep. LCA-REPORT-2008-010, 2008.

[41] M. Lotfinezhad and P. Marbach, "Throughput-optimal random access with order-optimal delay," arXiv:1009.5944[cs.IT].

[42] E. Vigoda, "A note on the glauber dynamics for sampling independent sets," *Electronic Journal of Combinatorics*, vol. 8, 2001.