From Glauber Dynamics To Metropolis Algorithm: Smaller Delay in Optimal CSMA

Chul-Ho Lee, Do Young Eun
Dept. of ECE, NC State University, Raleigh, NC, USA
Email: {clee4, dyeun}@ncsu.edu

Se-Young Yun, Yung Yi
Dept. of EE, KAIST, Daejeon, Korea
Email: syyun@comis.kaist.ac.kr, yiyung@kaist.edu

***Abstract*—Glauber dynamics, a method of sampling a given probability distribution via a Markov chain, has recently made considerable contribution to the MAC scheduling research, providing a tool to solve a long-standing open issue – achieving throughput-optimality with light message passing under CSMA. In this paper, we propose a way of reducing delay by studying *generalized* Glauber dynamics parameterized by $\beta \in [0, 1]$, ranging from Glauber dynamics ($\beta = 0$) to the Metropolis algorithm ($\beta = 1$). The same stationary distribution is sustained across this generalization, thus maintaining the long-term optimality. However, a different choice of β results in a significantly different second-order behavior (or variability) that has large impact on delay, which is hardly captured by the recent research focusing on delay in the large n (the number of nodes) asymptotic. We formally study such second-order behavior and its resulting delay performance, and show that larger β achieves smaller delay. Our results provide new insight into how to operate CSMA for large throughput and small delay in real, finite-sized systems.**

I. Introduction

Since the seminal work by Tassiulas and Ephremides on throughput-optimal scheduling [29], referred to as Max-Weight, a huge array of research has been made to develop distributed MAC scheduling with high performance guarantee and low complexity. The tradeoff between complexity and efficiency has been, however, observed in many cases, or even throughput-optimal algorithms with polynomial complexity have turned out to require heavy message passing (see, e.g., [32]). A breakthrough has been recently made, where just locally controlling the classical CSMA parameters, which is modeled by Glauber dynamics*, is enough to achieve throughput-optimality, see e.g., [12], [15], [19], [27]. We call this "optimal CSMA" for brevity.

In addition to throughput or utility, delay is another key performance metric in MAC scheduling. Delay research in MAC scheduling with performance guarantee has been studied with mathematical tools such as large deviation theory, heavy traffic approximation, and Lyapunov bound (see, e.g., [32] and references therein). However, delay in Glauber-dynamics based CSMA (or optimal CSMA) has been under-explored, where only a small set of work has been published with emphasis on the asymptotic results. Shah et al. [24] showed that it is unlikely to expect a simple MAC protocol such as CSMA to have high throughput and low delay. Motivated by such a "negative" result, Shah and Shin [25] proposed a modified CSMA requiring *coloring operation* that achieves $O(1)$ delay with throughput-optimality for networks with geometry (or polynomial growth). Lotfinezhad and Marbach [16] proved that a *reshuffling* approach, which periodically reshuffles all on-going schedules under time synchronized CSMA†, leads to both throughput-optimality and $O(1)$ delay for torus (inference) topologies. Jiang et al. [11] proved that a discrete-time parallelized Glauber dynamics achieves $O(\log n)$ delay for a limited set of arrival rates.

Despite these nice results on the delay asymptote for large-scale networks, it still remains questionable *how to improve the delay performance of (standard) Glauber-dynamics based CSMA for unscaled, fixed networks without loss of other important metrics such as throughput and complexity*. It is also unclear which tools to use for such purpose. While mixing time has been a popular toolkit for delay analysis [25], [11], it was shown very recently [28] that mixing time based approach may not be the right way to capture delay dynamics even in the asymptotic sense. On the other hand, the development of optimal CSMA, in principle, is equivalent to constructing a (reversible) Markov chain to achieve a given, desired stationary distribution under some constraints due to the interference. We note that Glauber dynamics is just *one* such instance. There can be many other Markov chains with the same stationary distribution (thus leading to throughput-optimality) and no additional complexity, but potentially higher efficiency for smaller delay under the same constraints.

In this paper, we propose, as extensions of the Glauber dynamics, a class of algorithms with a tunable parameter $\beta \in [0, 1]$, named *generalized Glauber dynamics*, ranging from the Glauber dynamics ($\beta = 0$) to the Metropolis algorithm ($\beta = 1$). We then show that the generalized Glauber dynamics or corresponding reversible Markov chain achieves the same stationary distribution *regardless of* the choice of β, while the Markov chain, when $\beta \in (0, 1]$, is *more efficient* than that under the Glauber dynamics ($\beta = 0$) in the sense of *Peskun ordering*, i.e., a partial order between off-diagonal elements of transition matrices of different Markov chains. Due to the invariant stationary distribution property, the generalized Glauber dynamics, when it comes into play for the problem of optimal CSMA, guarantees the same long-term throughput and also achieve throughput-optimality under mild conditions.

*Glauber dynamics is a Markov Chain Monte Carlo (MCMC) method for sampling a given probability distribution by constructing a Markov chain achieving the desired distribution as its unique stationary distribution [14].

†Thus, this is not a Glauber-dynamics based CSMA.

Algorithm 1 Glauber Dynamics (at Time Slot t)

Choose a node $v \in \mathcal{N}$ uniformly at random
For node v:
if $\sum_{w \in N_v} \sigma_w(t-1) = 0$ **then**
  $\sigma_v(t) = 1$ with probability $\frac{\lambda_v}{1+\lambda_v}$
  $\sigma_v(t) = 0$ with probability $\frac{1}{1+\lambda_v}$
else
  $\sigma_v(t) = 0$
end if
For any node $w \in \mathcal{N} \setminus \{v\}$: $\sigma_w(t) = \sigma_w(t-1)$

Despite the same long-term throughput, their 'second-order' behavior can be quite different. This in turn leads to different queueing delay performance, especially under the network of a reasonable size, which is *hardly* captured by any asymptotic order-wise analysis. However, thanks to the Peskun ordering and its relationship with efficiency ordering, we are able to demonstrate, in theory and simulation, that the original Glauber dynamics ($\beta = 0$) in fact gives the worst queueing delay performance among the generalized Glauber dynamics, and there are infinitely many different variants that have the same long-term throughput, but with *better* queueing delay performance as β increases, culminating in the 'Metropolised' version with $\beta = 1$ for *any* finite-sized networks.

II. Preliminaries

A. Glauber Dynamics for the Hard-core Model

Consider a connected, undirected graph $G = (\mathcal{N}, \mathcal{E})$ with a finite set of nodes (or vertexes) $\mathcal{N} = \{1, 2, \ldots, n\}$ and an edge set $\mathcal{E}$. Let $N_v = \{w \in \mathcal{N} : (v, w) \in \mathcal{E}\}$ be the set of neighbors of node v. We define by $\boldsymbol{\sigma}$ a *configuration* of the nodes in G, which is given by $\boldsymbol{\sigma} = \{\sigma_v, v \in \mathcal{N}\}$ with $\sigma_v \in \{0, 1\}$ for all v. A configuration is said to be *feasible* if the set $\{v \in \mathcal{N} : \sigma_v = 1\}$ is an independent set of G where no two nodes in the set are adjacent (or neighbor of each other), i.e., if $\sigma_v + \sigma_w \le 1$ for all $(v, w) \in \mathcal{E}$. Let $\Omega \subseteq \{0, 1\}^n$ also be the set of all *feasible* configurations on G. This model under the constraint of independent sets is called the hard-core model [14].

The (single-site update) Glauber dynamics for the hard-core model with heterogeneous fugacities $\{\lambda_v, v \in \mathcal{N}\}$, defined in Algorithm 1, leads to a (discrete-time) Markov chain achieving the following stationary distribution $\boldsymbol{\pi} = \{\pi(\boldsymbol{\sigma})\}$ over Ω:

$$\pi(\boldsymbol{\sigma}) = \frac{1}{Z} \prod_{v \in \mathcal{N}} \lambda_v^{\sigma_v}, \tag{1}$$

with a normalizing constant $Z = \sum_{\boldsymbol{\sigma} \in \Omega} \prod_{v \in \mathcal{N}} \lambda_v^{\sigma_v}$. Note that $\lambda_v > 0$ for all v, ensuring that $\pi(\boldsymbol{\sigma}) > 0$ for all $\boldsymbol{\sigma} \in \Omega$. Specifically, $\boldsymbol{\sigma}(t) = \{\sigma_v(t), v \in \mathcal{N}\}$ denotes the state of the Markov chain (or a feasible configuration by the Glauber dynamics) at time slot t. It is known that $\{\boldsymbol{\sigma}(t)\}_{t \ge 0}$ is an irreducible, aperiodic Markov chain achieving the stationary distribution $\boldsymbol{\pi}$ in (1) on the finite state space Ω [14], [27], [25]. The Markov chain $\{\boldsymbol{\sigma}(t)\}$ is also reversible with respect to $\boldsymbol{\pi}$, i.e., $\pi(\boldsymbol{\sigma}) Q(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \pi(\boldsymbol{\sigma'}) Q(\boldsymbol{\sigma'}, \boldsymbol{\sigma})$ for all $\boldsymbol{\sigma}, \boldsymbol{\sigma'} \in \Omega$, where $Q(\boldsymbol{\sigma}, \boldsymbol{\sigma'})$ is the transition probability from state $\boldsymbol{\sigma}$ to state $\boldsymbol{\sigma'}$.

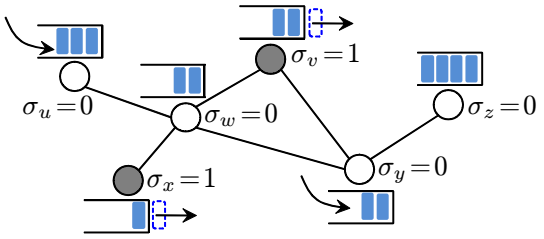

Fig. 1. Illustration of a wireless multihop scheduling driven by the Glauber dynamics over a conflict graph G with $|\mathcal{N}| = 6$. Links v, x are active while all others are silent, forming one instance of independent set on G.

B. CSMA and Glauber Dynamics

We present how CSMA in wireless multi-hop networks can be modeled by the Glauber dynamics. In the context of wireless multihop scheduling (or, simply scheduling), define a *link* as an (ordered) transmitter-receiver pair. It is said that two links *conflict* with each other if they cannot be "active" for communication at the same time due to the interference. Consequently, we can define a *conflict graph* $G = (\mathcal{N}, \mathcal{E})$ in which each node represents a link, while an edge between two nodes (or links) exists if they conflict with each other. Given a graph G, the scheduling governed by Glauber dynamics determines which nodes to be active or available for communication, forming one instance of independent sets (feasible configuration) over G at each time t in a distributed manner. For each node $v \in \mathcal{N}$, if $\sigma_v(t) = 1$, then node v is active, i.e., the transmitter of link (or node) v can transmit a packet to its receiver pair, and node v should be silent, if otherwise. See Fig. 1 for an illustrative example. Throughout this paper, the graph G refers to a conflict graph.

The Glauber dynamics in the context of scheduling is typically considered under continuous-time (or asynchronous) setting as used in [27], [25], which is also our target scenario. Specifically, each node is equipped with its own Poisson clock of rate 1, leading to the uniform node selection in Algorithm 1, and then decides its transmission schedule (or updates its status) accordingly. Here, the 'master' clock is Poisson with rate n and each (master) clock tick corresponds to a discrete-time slot in Algorithm 1. It is not difficult to see that the Glauber dynamics captures the following CSMA features: 1) *random back-off*: the transmitter of link v waits an exponentially distributed period of time with mean $(1+\lambda_v)/\lambda_v$ before transmitting (provided that the channel is sensed 'idle'); 2) *channel holding time*: once the transmitter of link v grabs the channel for transmission, it keeps the channel for an exponential distributed period of time with mean $1+\lambda_v$. For the sake of simplicity, our subsequent analysis is based on the discrete-time model, which is then supported by extensive numerical simulations under the continuous-time model.

III. Comparing Reversible Markov Chains

There are potentially many other (discrete-time) reversible Markov chains with the same $\boldsymbol{\pi}$ in (1), all of which translate into distributed algorithms just like the one in Algorithm 1, as will be shown later. One important question would be how to compare these reversible Markov chains. As these algorithms have the same $\boldsymbol{\pi}$, they all guarantee the same long-term

throughput, while their 'second-order' behavior can be quite different, leading to different queueing delay performance.

A. Mixing Time of Reversible Markov Chains

Mixing time has been a popular criterion to compare competing reversible Markov chains with the same stationary distribution. The mixing time indicates the speed of convergence to the stationary distribution, and is generally defined based on the total variation distance. Consider an irreducible, reversible Markov chain on a finite state space $\mathcal{S} = \{1, 2, \ldots, n\}$ with a transition matrix $\mathbf{P} = \{P(i,j)\}_{i,j\in\mathcal{S}}$ and a stationary distribution $\boldsymbol{\pi} = \{\pi(1), \pi(2), \ldots, \pi(n)\}$. Let $P^t(i,B)$ be its t-step transition probability from state i to a subset $B \subseteq \mathcal{S}$, and let $1 = e_1 > e_2 \geq \cdots \geq e_n \geq -1$ be the eigenvalues of $\mathbf{P}$. As shown in [5], if we define the total variation distance $||P^t(i,\cdot) - \boldsymbol{\pi}||_{TV} \triangleq \max_{B\subseteq\mathcal{S}} |P^t(i,B) - \pi_B|$ where $\pi_B = \sum_{j\in B}\pi(i)$, then for all $i \in \mathcal{S}$ and all $t \geq 1$,

$$||P^t(i,\cdot) - \boldsymbol{\pi}||_{TV} \leq \frac{1}{2}\sqrt{\frac{1-\pi(i)}{\pi(i)}}e(\mathbf{P})^t \leq \frac{1}{2\sqrt{\pi_{min}}}e(\mathbf{P})^t,$$

where $e(\mathbf{P}) = \max\{e_2, |e_n|\}$ is the second largest eigenvalue modulus (SLEM) of $\mathbf{P}$, and $\pi_{min} = \min_{i\in\mathcal{S}}\pi(i)$. If we define the mixing time in a general form as

$$t_{mix}(\varepsilon) \triangleq \min\{t \geq 1 : \max_{i\in\mathcal{S}} ||P^t(i,\cdot) - \boldsymbol{\pi}||_{TV} \leq \varepsilon\},$$

then $t_{mix}(\varepsilon) \leq \log(1/(\varepsilon\pi_{\min}))/(1 - e(\mathbf{P}))$ [14]. All these imply that the mixing time is mainly determined by the SLEM. That is, smaller SLEM leads to smaller (faster) mixing time.

The mixing time has been also used to determine the upper bound of the average queue length (or delay) in [11]. They show that the mixing time of the *parallel* Glauber dynamics, which allows parallel updates of node status at each time slot (in contrast to the single-site update Glauber dynamics), is $O(\log n)$, and in turn, the average queue length is also $O(\log n)$ for a limited set of arrival rates. On the other hand, it is known that the mixing time of the (single-site update) Glauber dynamics is *at least* $\Omega(n \log n)$ [10], while its asymptotic upper bound is $O(n \log n)$ [31]. As mentioned before, the Glauber dynamics (not limited to the context of scheduling) is widely considered under an asynchronous or continuous-time setting [10], [27], [25]. In the continuous-time setting (with Poisson clock of rate 1), since the master clock rate is n, the time scale is scaled down by a factor of $1/n$, implying the similar parallel-update effect to the discrete-time parallel Glauber dynamics in [19], [11]. It is also shown in [10] that the mixing time of the *continuous-time* Glauber dynamics is at least $\Omega(\log n)$. Thus, in principle, there is no difference between the discrete-time parallel Glauber dynamics and its continuous-time counterpart, other than whether the system is synchronized or not.

It is worth noting that the asymptotic lower bounds obtained in [10] generally hold for *any* update rule from which the transition probabilities satisfy the reversibility condition with respect to $\boldsymbol{\pi}$ in (1). Thus, among all reversible chains with the same $\boldsymbol{\pi}$ (the same long-term throughput), it is *impossible* to obtain a chain with smaller mixing time than $O(n \log n)$ for discrete-time single site update (or $O(\log n)$ for continuous-time) by extending the Glauber dynamics. However, it is yet *unknown* whether such extensions lead to better "actual" queueing delay performance (not an asymptotic order-wise upper bound via mixing time) when they are used as distributed CSMA algorithms, which is the main thrust of this paper. This problem is quite important, since it was shown in [28] that the queueing delay of Glauber-dynamics based CSMA can be bounded independently of the network size n for a certain set of arrival rates, implying that delay bounds based on mixing time analysis may be loose.

B. Peskun Ordering and Efficiency Ordering

We look at the comparison of reversible Markov chains from a different, but important perspective. This is done based on a partial order called the Peskun ordering [20]. The Peskun ordering can order reversible Markov chains with the same $\boldsymbol{\pi}$ in terms of the asymptotic variance and eigenvalues of their transition matrices.

Definition 1 (Peskun ordering): [20] For two irreducible Markov chains on a finite state space $\mathcal{S}$ with transition matrices $\mathbf{P} = \{P(i,j)\}_{i,j\in\mathcal{S}}$ and $\tilde{\mathbf{P}} = \{\tilde{P}(i,j)\}_{i,j\in\mathcal{S}}$, it is said that $\tilde{\mathbf{P}}$ dominates $\mathbf{P}$ off the diagonal, written as $\mathbf{P} \preceq \tilde{\mathbf{P}}$, if $P(i,j) \leq \tilde{P}(i,j)$ for all $i, j \in \mathcal{S}$ $(i \neq j)$. □

Let $\{X(t)\}_{t\geq 0}$ and $\{\tilde{X}(t)\}_{t\geq 0}$ be irreducible Markov chains on a finite state space $\mathcal{S} = \{1, 2, \ldots, n\}$ with transition matrices $\mathbf{P}$ and $\tilde{\mathbf{P}}$, respectively. Suppose that the Markov chains $\{X(t)\}$ and $\{\tilde{X}(t)\}$ have the same stationary distribution $\boldsymbol{\pi} = \{\pi(1), \pi(2), \ldots, \pi(n)\}$. For a function $f : \mathcal{S} \to \mathbb{R}$, define an estimator $\hat{\mu}_m \triangleq \sum_{t=1}^m f(X(t))/m$ for $\mu = \mathbb{E}_{\boldsymbol{\pi}}(f) = \sum_{i\in\mathcal{S}} f(i)\pi(i)$. It is well known that $\lim_{m\to\infty}\hat{\mu}_m = \mu$ for any function f with $\mathbb{E}_{\boldsymbol{\pi}}(|f|) < \infty$ [13], [14]. The asymptotic variance of the estimate $\hat{\mu}_m$ is defined as

$$\nu(\mathbf{P}, f) \triangleq \lim_{m\to\infty} m \cdot \mathbb{VAR}\left(\hat{\mu}_m\right), \tag{2}$$

which is *independent* of the distribution of the initial state $X(0)$ [20]. We similarly define $\nu(\tilde{\mathbf{P}}, f)$ for the chain $\{\tilde{X}(t)\}$ with $\tilde{\mathbf{P}}$. As mentioned before, the estimate $\hat{\mu}_m$ based on any finite, irreducible Markov chain with the same $\boldsymbol{\pi}$ always becomes μ, as m goes to infinity. However, since the asymptotic variance decides approximately how many samples are required to achieve a certain accuracy of the estimate $\hat{\mu}_m$,[‡] it has been an important criterion to rank the *efficiency* among competing Markov chains with the same $\boldsymbol{\pi}$, especially for the MCMC samplers [20], [18]. It is also said that $\{\tilde{X}(t)\}$ is at least as *efficient* as $\{X(t)\}$ if $\nu(\mathbf{P}, f) \geq \nu(\tilde{\mathbf{P}}, f)$ for any f with $\mathbb{VAR}_{\boldsymbol{\pi}}(f) < \infty$ [18]. In particular, this efficiency ordering is still in effect even when both chains are already in their stationary regimes (already mixed). The efficiency ordering will be the key component in the delay analysis later in the paper. It is known that the Peskun ordering between two reversible $\mathbf{P}$ and $\tilde{\mathbf{P}}$ with the same $\boldsymbol{\pi}$ provides a sufficient

[‡]This is formally captured by the Central Limit Theorem for an ergodic, finite Markov chain under $\mathbb{VAR}_{\boldsymbol{\pi}}(f) < \infty$, saying that $\sqrt{m}(\hat{\mu}_m - \mu)$ converges to a Gaussian random variable with zero mean and variance $\nu(\mathbf{P}, f)$ as $m \to \infty$ [13].

condition for the efficiency ordering and enables to order the eigenvalues of $\mathbf{P}$ and $\tilde{\mathbf{P}}$ as follows.

Lemma 1: [20] If $\mathbf{P}$ and $\tilde{\mathbf{P}}$ are reversible with respect to $\boldsymbol{\pi}$, and $\mathbf{P} \preceq \tilde{\mathbf{P}}$, then $\nu(\mathbf{P}, f) \geq \nu(\tilde{\mathbf{P}}, f)$ for any f with $\mathbb{VAR}_{\boldsymbol{\pi}}(f) < \infty$. □

Lemma 2: [18] Suppose that $\mathbf{P}$ and $\tilde{\mathbf{P}}$ are reversible with respect to $\boldsymbol{\pi}$. Let $1 = e_1 > e_2 \geq \cdots \geq e_n \geq -1$ be the eigenvalues of $\mathbf{P}$, and $1 = \tilde{e}_1 > \tilde{e}_2 \geq \cdots \geq \tilde{e}_n \geq -1$ be the eigenvalues of $\tilde{\mathbf{P}}$. If $\mathbf{P} \preceq \tilde{\mathbf{P}}$, then $e_i \geq \tilde{e}_i$ for all i. □

It is shown in [18] that for $\mathbf{P}$ and $\tilde{\mathbf{P}}$ reversible with respect to $\boldsymbol{\pi}$, $\nu(\mathbf{P}, f) \geq \nu(\tilde{\mathbf{P}}, f)$ if and only if $e_i \geq \tilde{e}_i$ for all i. If $\mathbf{P}$ and $\tilde{\mathbf{P}}$ are also positive definite, i.e., all eigenvalues are positive, then $e_i \geq \tilde{e}_i$, or $\nu(\mathbf{P}, f) \geq \nu(\tilde{\mathbf{P}}, f)$, implying that $e(\mathbf{P}) \geq e(\tilde{\mathbf{P}})$. (Thus, efficiency ordering would give mixing time ordering.) However, this is not true in general. 'Nearly-periodic' Markov chains can be such examples, for which e_n values become close to -1 (thus very efficient with small asymptotic variance), but with larger SLEM (and thus larger mixing time) [22].

IV. Extending Glauber Dynamics for the Hard-core Model

We first explain a generalization of the procedures by Hastings [9], [21] for constructing a reversible Markov chain with a given, desired stationary distribution, which was originally introduced for the development of an MCMC sampler. We then show that such generalized procedure can be applied to extending the Glauber dynamics for the hard-core model, while the Glauber dynamics is nothing but one of many possible ways to achieve a desired stationary distribution under the independent set constraints. In particular, we demonstrate that there are *infinitely* many ways to extend the Glauber dynamics, leading to *more efficient* reversible Markov chains in the sense of Peskun ordering (and efficiency ordering).

A. Generalized Hastings' Approach

An objective here is to construct an irreducible, reversible (discrete-time) Markov chain $\{X(t)\}_{t\geq 0}$ on a finite state space $\mathcal{S}$ with its transition matrix $\mathbf{P} = \{P(i,j)\}_{i,j\in\mathcal{S}}$ in order to achieve a given, desired stationary distribution $\boldsymbol{\pi} = \{\pi(i), i \in \mathcal{S}\}$. Hastings [9], [21] suggested a procedure for constructing a transition matrix $\mathbf{P}$ satisfying the reversibility condition, i.e., $\pi(i)P(i,j) = \pi(j)P(j,i)$ for all $i, j \in \mathcal{S}$. We below explain several key steps of this approach, as they are necessary for extending the Glauber dynamics for the hard-core model. They will also be used to clarify the difference between the Glauber dynamics and the Meterpolis-Hastings algorithm [17], [9] for the hard-core model, since there seems to be some confusion in the networking literature in that these two are often considered to be identical, e.g., [27], [25].

Suppose that $P(i,j)$ has the form

$$P(i,j) = \gamma(i,j)\alpha(i,j) \quad (i \neq j), \tag{3}$$

with $P(i,i) = 1 - \sum_{j\neq i} P(i,j)$, where $\gamma(i,j)$ is a proposal probability and $\alpha(i,j)$ is an acceptance probability. Specifically, at the current state i of $X(t)$, the next state $X(t+1)$ is proposed with the *proposal* probability $\gamma(i,j)$ $(i \neq j)$, which is the transition probability of an arbitrary irreducible Markov chain on the same space $\mathcal{S}$. We call $\{\gamma(i,j)\}_{i,j\in\mathcal{S}}$ a proposal matrix. The proposed state transition to $X(t+1) = j$ is then accepted with probability $\alpha(i,j)$, and rejected with probability $1 - \alpha(i,j)$ in which case $X(t+1) = i$. In this framework, applying the form of $P(i,j)$ in (3) to the reversibility condition yields that

$$\pi(i)\gamma(i,j)\alpha(i,j) = \pi(j)\gamma(j,i)\alpha(j,i) \quad (i \neq j), \tag{4}$$

which means

$$\alpha(i,j) = \frac{\pi(j)\gamma(j,i)}{\pi(i)\gamma(i,j)}\alpha(j,i) = \theta(j,i)\alpha(j,i) \quad (i \neq j), \tag{5}$$

where $\theta(j,i) \triangleq \frac{\pi(j)\gamma(j,i)}{\pi(i)\gamma(i,j)}$. The proposal matrix $\{\gamma(i,j)\}$, in addition to the desired $\boldsymbol{\pi}$, is often given a priori, which is also our case as shall be explained shortly. Thus, one can see that *any* acceptance probability $\alpha(i,j) \in [0,1]$ satisfying (5) makes the resulting transition matrix $\mathbf{P}$ reversible with respect to $\boldsymbol{\pi}$. By noting that $\theta(i,j) = 1/\theta(j,i)$, the acceptance probability $\alpha(i,j)$ is generally in the form of $F(\theta(i,j))$, where $0 \leq F \leq 1$ is any function that satisfies $F(x) = F(1/x)/x$ for all x.

Among infinitely many possible choices for $F(x)$, Hastings [9], [21] considered a class of functions given by

$$F(x) = s/(1+x), \tag{6}$$

where s can be arbitrarily chosen so that $0 \leq F(x) \leq 1$. Specifically, the acceptance probability $\alpha(i,j)$ is given by

$$\alpha(i,j) = s(i,j)/[1 + \theta(i,j)], \tag{7}$$

where $s(i,j)$ is a symmetric function of i and j chosen to ensure $0 \leq \alpha(i,j) \leq 1$. Suppose that $\gamma(i,j)$ is given for all $i, j \in \mathcal{S}$ while satisfying $\gamma(i,j) = \gamma(j,i)$ $(i \neq j)$. Then, $\theta(i,j) = \pi(i)/\pi(j)$. First, setting $s(i,j) = 1$ gives

$$\alpha(i,j) = \pi(j)/[\pi(i) + \pi(j)], \tag{8}$$

which leads to Barker's sampling method [2] where the transition probability is given by

$$P(i,j) = \gamma(i,j)\alpha(i,j) = \gamma(i,j) \cdot \frac{\pi(j)}{\pi(i) + \pi(j)} \quad (i \neq j) \tag{9}$$

with $P(i,i) = 1 - \sum_{j\neq i} P(i,j)$. On the other hand, if $s(i,j) = 1 + \min\{\theta(i,j), \theta(j,i)\}$, then (7) becomes

$$\alpha(i,j) = \min\{1, \pi(j)/\pi(i)\}, \tag{10}$$

which corresponds to Metropolis algorithm [17]. Its transition probability, similarly to what was shown for the Barker's method, is written as

$$P(i,j) = \gamma(i,j) \cdot \min\{1, \pi(j)/\pi(i)\} \quad (i \neq j) \tag{11}$$

with $P(i,i) = 1 - \sum_{j\neq i} P(i,j)$.

Remark 1: After the Metropolis algorithm [17] was extended by Hastings [9] for *non-symmetric* proposal probabilities $\gamma(i,j) \neq \gamma(j,i)$, the extended version is called Metropolis-Hastings algorithm in the literature where the transition probability is given by $P(i,j) = \min\{\gamma(i,j), \gamma(j,i)\pi(j)/\pi(i)\}$ with $P(i,i) = 1 - \sum_{j\neq i} P(i,j)$. □

B. From Glauber Dynamics to Metropolis Algorithm

We now interpret the reversible Markov chain by the Glauber dynamics for the hard-core model under the above framework of MCMC sampling, and then show that the reversible chain is connected to the one by Barker's sampling method under the independent set constraints.

Consider a feasible configuration (or state) at time t by the Glauber dynamics for the hard-core model, i.e., $\boldsymbol{\sigma}(t) = \{\sigma_v(t), v \in \mathcal{N}\}$. Let $\mathbf{Q} = \{Q(\boldsymbol{\sigma}, \boldsymbol{\sigma'})\}_{\boldsymbol{\sigma},\boldsymbol{\sigma'} \in \Omega}$ be the transition matrix of the Markov chain $\{\boldsymbol{\sigma}(t)\}_{t \geq 0}$. We here evaluate the transition matrix $\mathbf{Q}$. Observe that only a single, uniformly chosen node can change its status ($\sigma_v(t-1) = 1 \to \sigma_v(t) = 0$, or vice versa) at time t under the independent set constraints. It means that for any configurations $\boldsymbol{\sigma}, \boldsymbol{\sigma'} \in \Omega$, if $\|\boldsymbol{\sigma} - \boldsymbol{\sigma'}\| \triangleq \sum_{v \in \mathcal{N}} |\sigma_v - \sigma'_v| > 1$, then $Q(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = 0$. It thus suffices to consider any $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma'}$ with $\|\boldsymbol{\sigma} - \boldsymbol{\sigma'}\| = 1$ in evaluating $\mathbf{Q}$. For each node v, we define by Φ_v a subset of all pairs of feasible configurations, which is given by

$$\Phi_v = \{(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) : \sigma_v = 0, \sigma'_v = 1, \sigma_u = \sigma'_u = 0 \text{ for all } u \in N_v, \text{ and } \sigma_w = \sigma'_w \text{ for all } w \in \mathcal{N} \setminus (N_v \cup \{v\})\}. \quad (12)$$

Each element of this set Φ_v is a pair of configurations in which only the status of node v is different, while the status of all the neighbors of v are all zero and all other nodes' status remain the same. Then, for any $(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) \in \Phi_v$, by definition, we have

$$Q(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \frac{1}{n} \cdot \frac{\lambda_v}{1 + \lambda_v} \quad \text{and} \quad Q(\boldsymbol{\sigma'}, \boldsymbol{\sigma}) = \frac{1}{n} \cdot \frac{1}{1 + \lambda_v}. \quad (13)$$

Here, the uniform node selection (line 1 in Algorithm 1) can be interpreted as follows: given that the current configuration is $\boldsymbol{\sigma}$, a transition from $\boldsymbol{\sigma}$ to $\boldsymbol{\sigma'}$ is *proposed* with probability

$$\gamma(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \gamma(\boldsymbol{\sigma'}, \boldsymbol{\sigma}) = 1/n, \quad (14)$$

which can be also considered as the proposal probability for the transition from $\boldsymbol{\sigma'}$ to $\boldsymbol{\sigma}$. The state transition from $\boldsymbol{\sigma}$ to $\boldsymbol{\sigma'}$ is then *accepted* with probability

$$\alpha(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \pi(\boldsymbol{\sigma'})/[\pi(\boldsymbol{\sigma}) + \pi(\boldsymbol{\sigma'})] = \lambda_v/(1 + \lambda_v), \quad (15)$$

where the second equality is from (1). Similarly obtain the acceptance probability $\alpha(\boldsymbol{\sigma'}, \boldsymbol{\sigma})$. Combining (14)–(15) with (13), one can see that (13) has exactly the same structure as the one by Barker's sampling method in (9) under the independent set constraints. Note that the Barker's sampling method in (9) does not assume any specific $\gamma(\boldsymbol{\sigma}, \boldsymbol{\sigma'})$. That is, in Algorithm 1, the uniform node selection leading to $\gamma(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \gamma(\boldsymbol{\sigma'}, \boldsymbol{\sigma}) = 1/n$ is not necessary, or rather, can be extended as follows: at time slot t, a node v is randomly chosen according to a node-selection probability distribution $\{q_v, v \in \mathcal{N}\}$, where $q_v > 0$ for all v, and $\sum_{v \in \mathcal{N}} q_v = 1$. In this extension, $\gamma(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \gamma(\boldsymbol{\sigma'}, \boldsymbol{\sigma}) = q_v$.

From the above argument along with the generalized Hastings' approach in Section IV-A, we observe that for a given $\{q_v\}$, there are infinitely many possible update rules for node status including the one by Glauber dynamics (lines 4–5 in Algorithm 1), each of which is simply one possible form of acceptance probabilities. Fix node v. Then, for any $(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) \in \Phi_v$, we choose a class of symmetric functions given by $s(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = [1 + \min\{\theta(\boldsymbol{\sigma}, \boldsymbol{\sigma'}), \theta(\boldsymbol{\sigma'}, \boldsymbol{\sigma})\}]^{\beta}$ in (7), parameterized by a *free* variable $\beta \in [0, 1]$, where $\theta(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \pi(\boldsymbol{\sigma})/\pi(\boldsymbol{\sigma'}) = 1/\lambda_v$ and $\theta(\boldsymbol{\sigma'}, \boldsymbol{\sigma}) = \lambda_v$. Note that if $\beta = 0$, then $s(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = 1$, which leads to (8), or (15). Also, if $\beta = 1$, then $s(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = 1 + \min\{\theta(\boldsymbol{\sigma}, \boldsymbol{\sigma'}), \theta(\boldsymbol{\sigma'}, \boldsymbol{\sigma})\}$, which corresponds to (10). That is, we have

$$\alpha(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \frac{[1 + \min\{\theta(\boldsymbol{\sigma}, \boldsymbol{\sigma'}), \theta(\boldsymbol{\sigma'}, \boldsymbol{\sigma})\}]^{\beta}}{1 + \theta(\boldsymbol{\sigma}, \boldsymbol{\sigma'})}. \quad (16)$$

After little algebraic manipulations, (16) becomes

$$\alpha(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \left(\frac{\pi(\boldsymbol{\sigma'})}{\pi(\boldsymbol{\sigma}) + \pi(\boldsymbol{\sigma'})}\right)^{1-\beta} \min\left\{1, \left(\frac{\pi(\boldsymbol{\sigma'})}{\pi(\boldsymbol{\sigma})}\right)^{\beta}\right\} = \left(\frac{\lambda_v}{1 + \lambda_v}\right)^{1-\beta} \min\{1, \lambda_v^{\beta}\}. \quad (17)$$

Similarly, we obtain

$$\alpha(\boldsymbol{\sigma'}, \boldsymbol{\sigma}) = \left(\frac{1}{1 + \lambda_v}\right)^{1-\beta} \min\left\{1, 1/\lambda_v^{\beta}\right\}. \quad (18)$$

Consequently, we achieve a class of algorithms with a controllable parameter $\beta \in [0, 1]$, named *generalized Glauber dynamics*, which is summarized in Algorithm 2. As a special case, if $\beta = 0$, then Algorithm 2 becomes identical to Algorithm 1 – the original Glauber dynamics for the hard-core model. Also, if $\beta = 1$, then it means that the Metropolis algorithm is applied properly for the hard-core model. The only difference between the generalized Glauber dynamics with $\beta \in (0, 1]$ and the original Glauber dynamics ($\beta = 0$) is that for a randomly chosen node v, if $\sigma_u(t-1) = 0$ for all $u \in N_v$, then $\sigma_v(t)$ is decided based on $\sigma_v(t-1)$ for any $\beta \in (0, 1]$, while $\sigma_v(t)$ is determined *independently of* $\sigma_v(t-1)$ for $\beta = 0$.

For any given $\{q_v\}$, let $\boldsymbol{\sigma}(t, \beta)$ be a configuration at time t by the generalized Glauber dynamics with $\beta \in [0, 1]$. One can see that $\{\boldsymbol{\sigma}(t, \beta)\}_{t \geq 0}$ is a finite Markov chain with a transition matrix $\mathbf{Q}_{\beta} = \{Q_{\beta}(\boldsymbol{\sigma}, \boldsymbol{\sigma'})\}_{\boldsymbol{\sigma},\boldsymbol{\sigma'} \in \Omega}$. We say that the Markov chain is ergodic if $\pi(\boldsymbol{\sigma'}) = \lim_{t \to \infty} Q_{\beta}^t(\boldsymbol{\sigma}, \boldsymbol{\sigma'})$, where $Q_{\beta}^t(\boldsymbol{\sigma}, \boldsymbol{\sigma'})$ is the t-step transition probability from state $\boldsymbol{\sigma}$ to state $\boldsymbol{\sigma'}$. We next show properties of the generalized Glauber dynamics.

Algorithm 2 Generalized Glauber Dynamics with $\beta \in [0, 1]$ (at Time Slot t)

1: Choose a node $v \in \mathcal{N}$ according to a given $\{q_v\}$
2: For node v:
3: **if** $\sum_{w \in N_v} \sigma_w(t-1) = 0$ **then**
4: **if** $\sigma_v(t-1) = 0$ **then**
5: $\sigma_v(t) = 1$ with probability $\left(\frac{\lambda_v}{1+\lambda_v}\right)^{1-\beta} \min\left\{1, \lambda_v^{\beta}\right\}$
6: $\sigma_v(t) = 0$, otherwise.
7: **else**
8: $\sigma_v(t) = 0$ with probability $\left(\frac{1}{1+\lambda_v}\right)^{1-\beta} \min\left\{1, 1/\lambda_v^{\beta}\right\}$
9: $\sigma_v(t) = 1$, otherwise.
10: **end if**
11: **else**
12: $\sigma_v(t) = 0$
13: **end if**
14: For any node $w \in \mathcal{N} \setminus \{v\}$: $\sigma_w(t) = \sigma_w(t-1)$

Theorem 1: For any given $\{q_v\}$ and $\beta \in [0,1]$, the Markov chain $\{\boldsymbol{\sigma}(t,\beta)\}$ with $\mathbf{Q}_\beta$ is ergodic and reversible with respect to $\boldsymbol{\pi}$ in (1). In addition, for any given $\{q_v\}$ and $0 \le \beta_1 \le \beta_2 \le 1$, $\mathbf{Q}_{\beta_1} \preceq \mathbf{Q}_{\beta_2}$. □

Proof: See Appendix A ■

V. Generalized Glauber Dynamics for Smaller Delay in Optimal CSMA

A. Throughput Optimality

While our focus in this paper is to analyze the performance of each queue per node (in a conflict graph) when the generalized Glauber dynamics comes into play for the problem of optimal CSMA, we here explain the throughput-optimality of the generalized Glauber dynamics for two settings: *static* and *dynamic*.

In the *static* setting, we assume that the arrival rate vector for each node in the conflict graph is known to the system. Theorem 1 says that the stationary distribution $\boldsymbol{\pi}$ of the Markov chain $\{\boldsymbol{\sigma}(t,\beta)\}$ is *invariant* with respect to $\beta \in [0,1]$ and $\{q_v\}$. Thus, if the fugacity of each node λ_v can be chosen so that the long-term service rate (or capacity) at each queue is larger than its packet arrival rate, which is the typical case for delay analysis [11], [28], then 'throughput-optimality' or 'per-node stability' is achieved *irrespective of* the choice of $\beta \in [0,1]$ and $\{q_v\}$.

However, in reality, it may not be possible for each node v to adjust its fugacity λ_v based on the measured arrival and service rates. Hence, in the literature, the throughput-optimality has been defined and shown under the following *dynamic* setting: the fugacity is now a function of time t, which is given by $\lambda_v(t) = g(W_v(t))$ where g is some weight function and $W_v(t)$ is the queue-length at node v at time t. Even in this *dynamic* fugacity set-up, one can establish the throughput-optimality of the generalized Glauber dynamics with any given $\beta \in [0,1]$ and $\{q_v\}$. There are two different ways to prove the throughput-optimality in the literature. The first way is based on the time-scale decomposition assumption under which the system quickly converges to its stationary regime before its dynamics changes (see, e.g., [12], [19]). Under this assumption, Theorem 1 immediately implies that the generalized Glauber dynamics is throughput-optimal.

The second approach in [26] is taken without the time-scale decomposition assumption when $q_v = 1/n$ for all v and $\beta = 0$, but by choosing a proper weight function g, so that $\lambda_v(t) = g(W_v(t))$ changes much slower than the system dynamics. Theorem 2 formally states the throughput optimality of the generalized Glauber dynamics.

Theorem 2: For any $\beta \in [0,1]$, the generalized Glauber dynamics is throughput optimal when the fugacity $\lambda_v(t)$ of each node v at time t is set by:

$$\lambda_v(t) = \exp(\max\{f(W_v(\lfloor t \rfloor)), \sqrt{f(W_{\max}(\lfloor t \rfloor))}\}), \qquad (19)$$

where $f(\cdot) = \log\log(\cdot + e)$, and $W_{\max}(\cdot) = \max_v W_v(\cdot)$.

Proof: See Appendix B. ■

In [26], the throughput optimality of the case when $\beta = 0$ has been proved. Our proof for $\beta = (0,1]$ is similar to that in [26]. Henceforth, in Appendix, we just show the sketch of the proof 2 for completeness, with focus on the major difference from [26].

B. Delay Analysis

While not much is known yet about the queueing delay performance of optimal CSMA algorithms, we emphasize that the time-varying behavior of $\lambda_v(t)$ (in the dynamic fugacity set-up) makes the analysis even more intractable. So, as used in [11], [28], we here focus on the following case for delay analysis: the fugacity of each node λ_v is given and fixed, but possibly heterogeneous over $v \in \mathcal{N}$, such that the long-term service rate at each queue is larger than its packet arrival rate. We then demonstrate that higher efficiency in the extensions of Glauber dynamics, the choices of $\beta \in (0,1]$, give rise to better queueing delay performance, while maintaining the same long-term throughput. Specifically, the original Glauber dynamics with $\beta = 0$ in fact gives the *worst* queueing delay performance, and there are infinitely many different variants of 'throughput-optimal' algorithms with better queueing delay performance as β increases, culminating in the 'Metropolised' version with $\beta = 1$. We also support our analytical findings for the dynamic fugacity case through extensive simulations under various network topologies and arrival rates.

Fix $\{q_v\}$ and $\beta \in [0,1]$. Since we are interested in the long-term behavior of the queueing delay performance, without loss of generality, we can assume that the system is in its stationary regime.[§] Thus, the Markov chain $\{\boldsymbol{\sigma}(t,\beta)\}_{t\ge 0}$ is in the steady-state, i.e., $\mathbb{P}\{\boldsymbol{\sigma}(t,\beta) = \boldsymbol{\sigma}\} = \pi(\boldsymbol{\sigma})$ for all $t \ge 0$. We consider that a packet arrives in each node v at the *beginning* of each time slot according to a stationary 0–1 process $\{A_v(t)\}$ with rate μ_v in which $A_v(t) = 1$ if there is a packet arrival to node v with probability $\mathbb{P}\{A_v(t) = 1\} = \mu_v$ at time t, and $A_v(t) = 0$, otherwise.[¶] On the other hand, whenever node v is available for communication, i.e., $\sigma_v(t) = 1$, it transmits one packet backlogged in its FIFO queue (if any) *during* time slot t. The communication (or service) availability at node v is modeled as a 0–1 valued process governed by the generalized Glauber dynamics. That is,

$$S_v(t) = \begin{cases} 1 & \text{if node } v \text{ is available for service, i.e., } \sigma_v(t) = 1 \\ 0 & \text{otherwise} \end{cases}$$
$$= \mathbf{1}_{\{\boldsymbol{\sigma}(t,\beta) \,\in\, B_v\}} \qquad (20)$$

for $t = 0, 1, \ldots$, where $B_v \triangleq \{\boldsymbol{\sigma} \in \Omega : \sigma_v = 1\} \subseteq \Omega$. We define $\pi_{B_v} \triangleq \sum_{\boldsymbol{\sigma} \in B_v} \pi(\boldsymbol{\sigma})$, the long-term proportion of communication availability at node v or its 'service rate'. From the stationarity of the Markov chain $\{\boldsymbol{\sigma}(t,\beta)\}$, we have $\mathbb{P}\{S_v(t) = 1\} = \mathbb{P}\{\boldsymbol{\sigma}(t,\beta) \in B_v\} = \pi_{B_v}$ for all t. Thus, $\{S_v(t)\}_{t\ge 0}$ is a stationary 0–1 process. Also, $\{S_v(t)\}$

[§]Any initial transient fluctuation will disappear when computing steady-state metrics for queueing dynamics. For instance, the initial queue-length doesn't matter for the analysis of M/G/1 queue in the steady-state.

[¶]We assume very general class of arrival processes $\{A_v(t)\}$, satisfying the usual conditions for the large deviation (large buffer) asymptotic to hold [8], [3]. Such processes include not only Bernoulli arrivals, but correlated arrivals such as auto-regressive processes whose autocorrelation functions are summable.

is independent of $\{A_v(t)\}$. As mentioned before, we assume that $\pi_{B_v} > \mu_v$ for all $v \in \mathcal{N}$, ensuring that utilization is strictly less than one at each queue.

Without loss of generality, we below examine the queueing delay at an arbitrarily chosen node v. From now on, our exposition will be all about the queue in node v. So, for the sake of notational simplicity, we drop the subscript v and use μ, $A(t)$, $S(t)$, B, and π_B instead of μ_v, $A_v(t)$, $S_v(t)$, B_v and π_{B_v}, respectively, unless stated otherwise. We first evaluate the time interval between the two successive communication availabilities at node v, which corresponds to the service time of a single-server queueing model. To this end, we define

$$T_1 \triangleq \min\{t \geq 0 : S(t) = 1\},\ T_{i+1} \triangleq \min\{t > T_i : S(t) = 1\},$$

and $\tau_i \triangleq T_{i+1} - T_i$ $(i \geq 1)$. Here, $\{\tau_i\}_{i \geq 1}$ are such time intervals, all identically distributed from the stationarity of $S(t)$, and also called the *recurrence* times to the state 1 for $\{S(t)\}$. Then, we have the following.

Theorem 3: For a given $\{q_v\}$, $\mathbb{E}\{\tau_i\} = 1/\pi_B$ for all β, and $\mathbb{E}\{(\tau_i)^2\}$ is decreasing in $\beta \in [0, 1]$. □

Proof: See Appendix C. ■

For a fixed $\{q_v\}$, the average recurrent time, $\mathbb{E}\{\tau_i\}$, remains the same for all β due to the invariance property of $\boldsymbol{\pi}$ with respect to β in Theorem 1, while the variance of the recurrence time is decreasing in β. In the standard queueing literature, the variance of the service time plays a major role in queueing performance. For example, it is well-known that, for M/G/1 queue, the variance of the service time solely determines the average queueing delay if the average service time is kept the same. Similarly, even for G/G/1 queue, more 'variable' service time leads to larger queueing delay [23]. However, our system is far more complicated than these standard queueing systems; the recurrence times $\{\tau_i\}$ can be possibly correlated over i for $|B| > 1$, as the time instant T_{i+1} depends on the configuration $\boldsymbol{\sigma}(T_i, \beta) \in B$ at time T_i. Such dependency, thus, makes the exact analysis of queueing delay performance intractable.‖ Nonetheless, for a given $\{q_v\}$, the 'marginal' distribution of the service time τ_i has smaller variance as β increases (with the same mean regardless of the choice of β), suggesting that larger β would lead to better delay performance.

In addition, we demonstrate that the efficiency ordering of $\mathbf{Q}_\beta$ for different β can still order the performance of queueing dynamics by directly taking into account the dependency structure in $\{\tau_i\}$ sequence. To proceed, let $W(t)$ be the queue length (or workload) at time t, satisfying Lindley recursion:

$$W(t+1) = \max\{0, W(t) + A(t+1) - S(t+1)\}. \quad (21)$$

From the large deviation theory, in considerable generality, the tail distribution of the steady-state queue length W is asymptotically exponential [8], [3], [4] with the asymptotic decay rate η given by

$$\eta = \lim_{x \to \infty} -\frac{1}{x} \log \mathbb{P}\{W > x\} > 0. \quad (22)$$

‖If $|B| = 1$, the exact analysis is possible, as $\{\tau_i\}$ forms an $i.i.d.$ sequence due to the strong Markov property. For example, if the arrival process is an independent Bernoulli process, the system is simply the discrete-time case of M/G/1 queue.

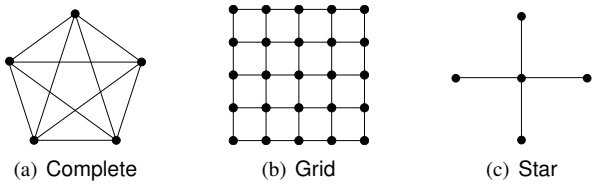

(a) Complete (b) Grid (c) Star

Fig. 2. Conflict graphs of Complete, Grid, and Star, in which each vertex represents a link and an edge between vertexes means that they interfere each other.

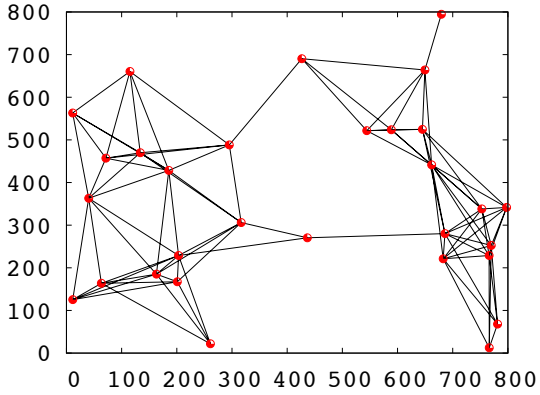

Fig. 3. Rand where a vertex is a node. In $800m \times 800m$, 30 nodes are unifomly located at random and links are generated when the inter-node distance is within 250m transmission rage.

Let $I(t) = A(t) - S(t)$ be the *net input* into the queue at time t. Let $\boldsymbol{I_t} \triangleq \sum_{i=1}^{t} I(i)$, $\boldsymbol{A_t} \triangleq \sum_{i=1}^{t} A(i)$, and $\boldsymbol{S_t} \triangleq \sum_{i=1}^{t} S(i)$. Clearly, $\mathbb{E}\{\boldsymbol{I_t}\} = \mathbb{E}\{\boldsymbol{A_t}\} - \mathbb{E}\{\boldsymbol{S_t}\} = t(\mu - \pi_B) < 0$. In this setup, we have the following:

Theorem 4: Suppose that the distribution of $\boldsymbol{I_t}$ is Gaussian for large t with $\lim_{t \to \infty} \mathbb{VAR}\{\boldsymbol{A_t}\}/t = \nu^* < \infty$. Then, η in (22) is increasing in $\beta \in [0, 1]$. □

Proof: See Appendix D. ■

Since $\boldsymbol{I_t} = \sum_{i=1}^{t} (A(i) - S(i))$ is the sum of t random variables, as long as their dependency over i is not so strong, it is reasonable to assume that $\boldsymbol{I_t}$ is roughly Gaussian for large t. Then, Theorem 4 tells us that larger β leads to faster decay of the tail distribution of the steady-state queue-length, again suggesting better queueing performance while preserving the throughput-optimality. The Gaussian approximation and Theorem 4 are also corroborated by simulation results.

As mentioned before, it may not be possible for each node v to choose its fugacity λ_v based on the measured arrival and service rates for per-node stability. Instead, the fugacity of each node needs to be an appropriate function of its (time-varying) queue-length. Nonetheless, if the corresponding temporal dynamics is relatively slow or is in 'almost-stationary' regime, from Theorems 3 and 4, we expect that the average queueing delay per node decreases in $\beta \in [0, 1]$ for a given $\{q_v\}$, which is also supported by simulation results.

VI. Simulation Results

A. Setup

Five topologies are considered in our simulation: Complete, Grid, Star, Rand(1), and Rand(2), shown in Fig. 2

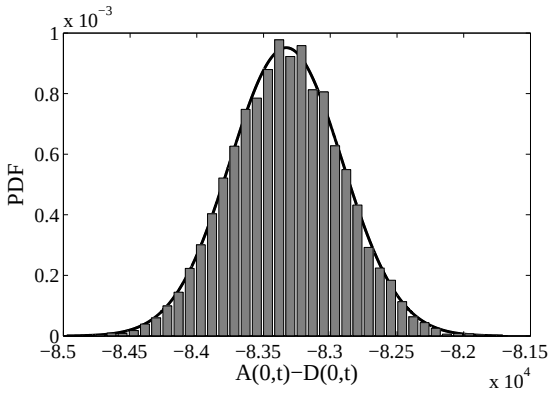

Fig. 4. Histogram of *net input* where topology is complete and fugacity is 1. (The line is $f(x) = 0.0009516 \times \exp(-(\frac{x+8.333\times10^4}{594})^2)$)

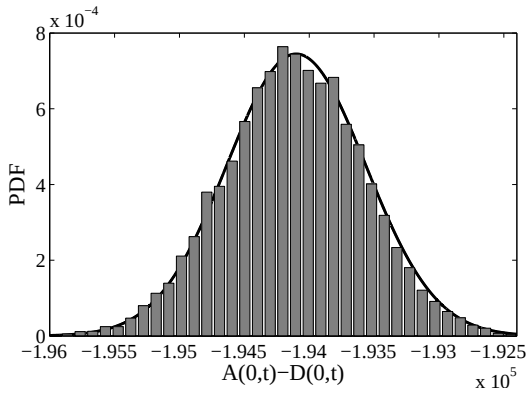

Fig. 5. Histogram of *net input* where topology is star and fugacity is 1. (The line is $f(x) = 0.0007448 \times \exp(-(\frac{x+1.941\times10^5}{756.4})^2)$)

and Fig. 3, where Rand(i) is a random graph with i-hop interference model**. Note that Complete, Grid, and Star are the *conflict graphs*, whereas Rand(1) and Rand(2) are the *original graphs*. In the random topologies, we uniformly place 30 nodes in the 800×800 m^2 area, and the each node's transmission range is set to be 250 m (a typical Wi-Fi range), where two nodes are connected if they are within the transmission range of each other. In all topologies, the link capacity is set to be one.

Generalized Glauber dynamics is simulated in random and asynchronous manner. In all simulations, every link has a Poisson clock with rate 1 and changes its schedule according to Algorithm 2 when its own clock expires. Thus, this scheduling scheme can be regarded as the generalized Glauber dynamics, where links are uniformly selected at random at each time slot and the time slot duration is exponentially distributed with the rate of n.

For packet arrivals at each time slot, we consider exogenous packet arrivals modeled by i.i.d. Bernoulli process and interrupted Bernoulli process (IBP) [7] with various average rates. All results are shown under the i.i.d. Bernoulli process, unless explicitly stated. In IBP, links are randomly on and off according to a two state Markov chain, where packets are generated according to a Bernoulli process only when it's in 'on' state. Note that on and off for IBP differs from active and inactive in scheduling. We set the stationary probability of on-state to be 0.1. IBP is used to test how the results vary for time-correlated arrival patterns with some bursts. Each packet size follows the i.i.d. exponential distribution with rate 1, i.e., inter-clock duration of each link is 1. Thus, one packet is served during one clock duarion. We simulate both static and dynamic cases, where in the dynamic cases the weight functions are $f(x) = x,\ \log(x+1),\ \log\log(x+e)$.

The main metric *average delay* indicates the average queueing delay per node. In the description, we sometimes use the expression of *normalized delay ratio* to refer to the delay ratio of $\beta = 1$ to $\beta = 0$, i.e., average delay($\beta = 1$)/average delay($\beta = 0$).

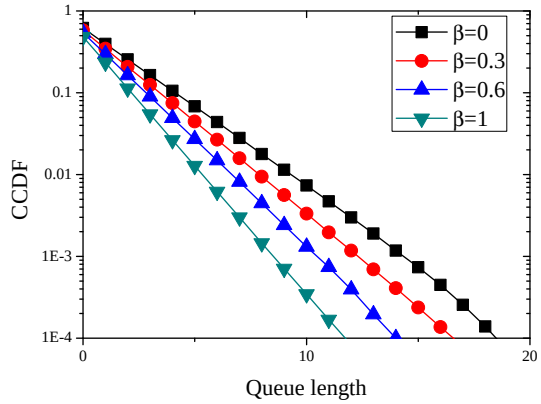

Fig. 6. CCDF of queue length for various β (0, 0.3, 0.6, and 1) where topology is complete and fugacity is 1

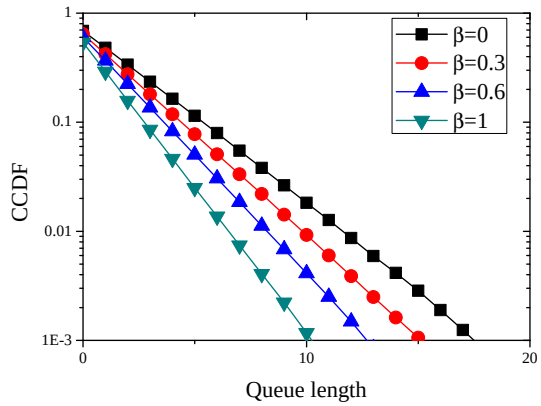

Fig. 7. CCDF of queue length for various β (0, 0.3, 0.6, and 1) where topology is star, fugacity is 1, and the queue is for the center node.

B. Results

The distribution of* net input *is gaussin: Fig. 4 shows the histogram of *net input*, where topology is complete and fugacity is fixed to 1, is well fitted with gaussian distribution with mean -8.333×10^4 and variance 1.764×10^5. Because, in Theorem 4, we have shown that , when *net input* is Gaussian, asymptotic decay rate η for each queue increases with larger

**Any two links within i-hop distance interfere with each other.

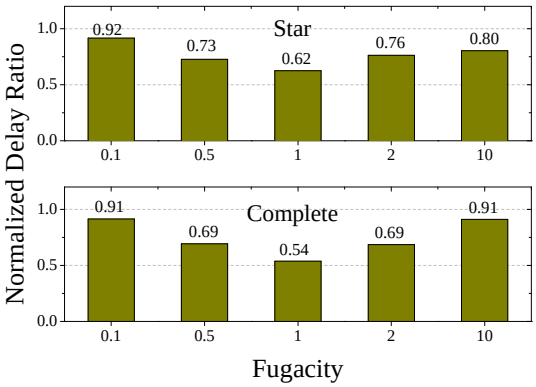

Fig. 8. Normalized delay ratio for static homogeneous fugacity as the fugacity(λ) varies, where the arrival rate of each link is the half of long-term service rate guaranteed by the fugacity.

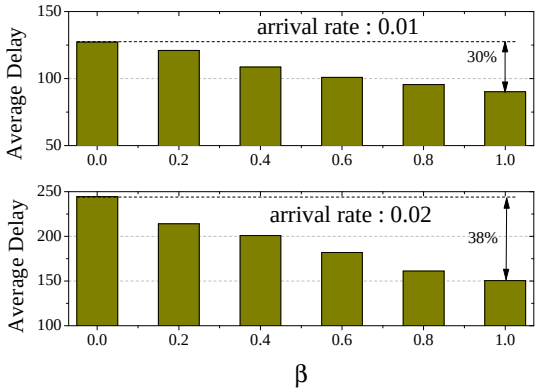

Fig. 9. Average delay for different β for dynamic fugacity with weight function, $f(x) = \log(x)$, in which topology is Rand(2).

β, the fitting results imply that the condition of Theorem 4 is satisfied. Therefore, we can expect that the semi-log graph for the complementary cumulative distribution function (CCDF) of queue length is linear and the slope becomes steeper with larger β. The asymptotic decay rate is shown in Fig. 6. As our expectation, semi-log plot of CCDF is linear and the slope increases as β increases, which indicates that with larger β delay is reduced. Not only for complete topology, but also for star star topology, we check simulation results suit with our expection. Fig. 5 and Fig. 7 show the histogram of *net input* and semi-log plot of CCDF, repectively, where topology is star. The both graphs well support our theorems as in complete topolgy.

Impact of fugacity for static cases: We first examine average delay for static fugacity. Fig. 8 shows the normalized delay ratio for Star and Complete. Note that smaller ratios imply larger reduction in average delay when $\beta = 1$. For each fugacity in the x-axis, the arrival rate on each link is set to be the half of available service rate provided by the fugacity. We observe that the average delay is improved when β changes from 0 to 1. In addition, in both topologies, the smallest delay ratio (i.e., the largest delay reduction which is about 50% delay reduction) is achieved when the fugacity is 1. This is because the gap of state transition probability in Algorithm 2 becomes largest when the fugacity is 1. The rest of the simulation results considers the dynamic cases when the fugacities are time-varying, set as a function of queue lengths.

Impact of β: Fig. 9 shows the average delay with changing β for Rand(2), $f(x) = \log(x)$, and the mean arrival rates 0.01 and 0.02. As suggested in Theorems 3 and 4, we observe that increasing β leads to decreasing average delay. The average delays for the arrival rates 0.02 and 0.01 decrease almost linearly from 244 to 150 and from 127 to 90, respectively (more than 30% decrease for both cases). This demonstrates that reduction in the average delay via Peskun (efficiency) ordering is significant, which is obtained just by modifying the local transition probabilities as in the generalized the glauber dynamics with no additional overhead.

Impact of $f(x)$ and topologies: We show the performance, measured by the normalized delay ratio, of dynamic cases for different $f(x)$ and topologies. Because feasible rate region

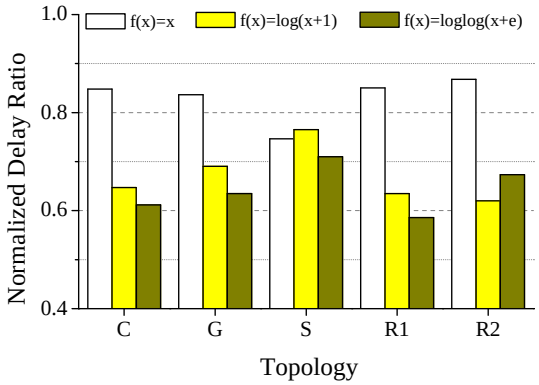

Fig. 10. Normalized delay ratios for different topologies (C:Complete, G:Grid, S:Star, R1:Rand(1), and R2:Rand(2)) when the arrival rate for each topology is 0.1, 0.1, 0.1, 0.05, and 0.02, respectively.

depends on topology, we set arrival rate differently for each topology. Fig. 10 shows the normalized average delay ratio of $\beta = 1$ by $\beta = 0$. We observe that the average delay reduction from $\beta = 0$ to $\beta = 1$ varies for different choices of weight functions. The average delay reduction for $f(x) = x$ is about 10%, whereas it ranges from about 30% to 40% for $f(x) = \log(x + 1), \log\log(x + e)$. This difference in average delay reduction is due to the fact that for an "aggressive" weight function, the fugacity is too high to be differentiated by β. This implies that more sluggish weight functions tend to induce larger delay reduction with increasing β.

Average delay and impact of arrival patterns: The results in the previous paragraph do *not* imply that the actual queue length for the aggressive weight function is large. Fig. 11 to Fig. 13 show such a fact, where the weight function $f(x)=x$ achieves smaller average delay over various arrival rates. In harmony with the earlier simulation results, larger β leads to smaller average delay. The result of smaller average delay for aggressive weight functions coincides with that in [27], where faster state changes in the glauber dynamics results in shorter average delay. Specifically, to have the fugacity of 10, we would need the queue length of 22026 for $f(x)=\log\log(x+e)$, whereas only queue length of 3 is enough for $f(x)=x$. We remark that the downside of aggressive weight functions may fail to achieve throughput optimality, as argued in [27],

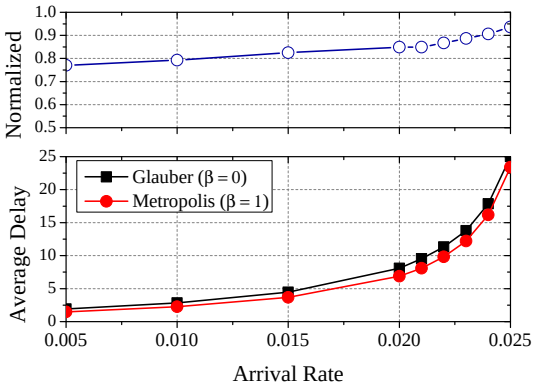

Fig. 11. Average delay with weight function $(f(x) = x)$ when β is 0 or 1 to show the effect of β. The above graph is normalized value (delay of $\beta = 0$ over delay of $\beta = 1$)

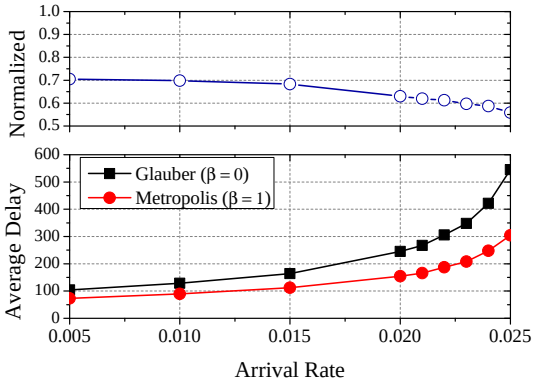

Fig. 12. Average delay with weight function $(f(x) = \log(x + 1))$ when β is 0 or 1 to show the effect of β. The above graph is normalized value (delay of $\beta = 0$ over delay of $\beta = 1$)

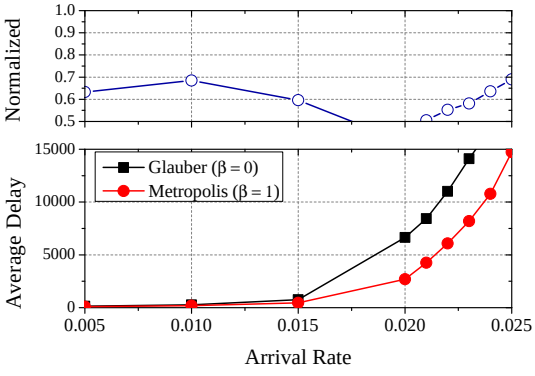

Fig. 13. Average delay with weight function $(f(x) = \log\log(x+e))$ when β is 0 or 1 to show the effect of β. The above graph is normalized value (delay of $\beta = 0$ over delay of $\beta = 1$)

mainly due to the lack of time to enjoy the "almost-stationary" regime for a given fugacity. We also observe that for IBP traffic arrival, the global trend does not change much, as shown in Fig. 14.

VII. Conclusion

We took a different direction, instead of relying on asymptotic delay analysis prevalent in recent studies, to achieve smaller delay in Glauber-dynamics based CSMA (or optimal CSMA) for finite-sized networks. By carefully exploring all possible variants of the traditional Glauber dynamics, we proposed generalized Glauber dynamics with no additional complexity, maintaining the same stationary distribution and thus rendering the long-term optimality unchanged in the context of scheduling. We then showed that our extensions lead to better queueing delay performance, by directly taking into account the second-order system behavior via a notion of Peskun (or efficiency) ordering and large deviation techniques.

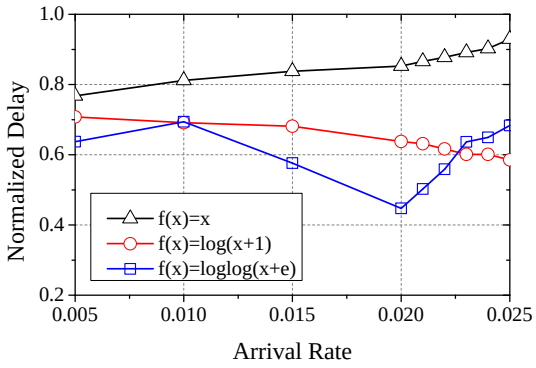

Fig. 14. Normalized value (delay of $\beta = 0$ over delay of $\beta = 1$) for various weight functions, when arrival patterns follow the interrupted Bernoulli process.

References

[1] D. Aldous and J. Fill, *Reversible Markov Chains and Random Walks on Graphs*. monograph in preparation.

[2] A. A. Barker, "Monte Carlo calculations of the radial distribution functions for a proton-electron plasma," *Australian Journal of Physics*, vol. 18, pp. 119–133, 1965.

[3] C. S. Chang, "Stability, queue length, and delay of deterministic and stochastic queueing networks," *IEEE Transactions on Automatic Control*, vol. 39, no. 5, pp. 913–931, 1994.

[4] G. L. Choudhury, D. M. Lucantoni, and W. Whitt, "Squeezing the most out of ATM," *IEEE Transactions on Communications*, vol. 44, no. 2, pp. 203–217, 1996.

[5] P. Diaconis and D. Stroock, "Geometric bounds for eigenvalues of Markov chains," *Annals of Applied Probability*, vol. 1, no. 1, pp. 36–61, 1991.

[6] R. Durrett, *Probability : Theory and Examples*, 2nd ed. Duxbury Press, 1996.

[7] F. Froehlic and A. Kent, *The Froehlich/Kent encyclopedia of telecommunications*. Marcel Dekker, 1995.

[8] P. W. Glynn and W. Whitt, "Logarithmic asymptotics for steady-state tail probabilities in a single-server queue," *Journal of Applied Probability*, vol. 31, pp. 131–156, 1994.

[9] W. K. Hastings, "Monte Carlo sampling methods using markov chains and their applications," *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.

[10] T. P. Hayes and A. Sinclair, "A general lower bound for mixing of single-site dynamics on graphs," *Annals of Applied Probability*, vol. 17, no. 3, pp. 931–952, 2007.

[11] L. Jiang, M. Leconte, J. Ni, R. Srikant, and J. Walrand, "Fast mixing of parallel Glauber dynamics and low-delay CSMA scheduling," in *Proceedings of IEEE Infocom*, 2011.

[12] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," *IEEE/ACM Transactions on Networking*, vol. 18, no. 3, pp. 960 –972, June 2010.

[13] G. L. Jones, "On the Markov chain central limit theorem," *Probability Surveys*, vol. 1, pp. 299–320, 2004.

[14] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*. American Mathematical Society, 2009.

[15] J. Liu, Y. Yi, A. Proutiere, M. Chiang, and H. V. Poor, "Towards utility-optimal random access without message passing," *Wireless Communications and Mobile Computing*, vol. 10, no. 1, pp. 115–128, Jan. 2010.

[16] M. Lotfinezhad and P. Marbach, "Throughput-optimal random access with order-optimal delay," in *Proceedings of IEEE Infocom*, 2011.

[17] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of state calculations by fast computing machines," *The Journal of Chemical Physics*, vol. 21, no. 6, pp. 1087–1092, 1953.
[18] A. Mira, "Ordering and improving the performance of Monte Carlo Markov chains," *Statistical Science*, vol. 16, no. 4, pp. 340–350, 2001.
[19] J. Ni, B. Tan, and R. Srikant, "Q-csma: Queue-length based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks," in *Proceedings of IEEE Infocom*, 2010.
[20] P. H. Peskun, "Optimum Monte-Carlo sampling using Markov chains," *Biometrika*, vol. 60, pp. 607–612, 1973.
[21] M. Richey, "The evolution of Markov chain Monte Carlo methods," *American Mathematical Monthly*, vol. 117, no. 5, pp. 383–413, 2010.
[22] J. S. Rosenthal, "Asymptotic variance and convergence rates of nearly-periodic Markov chain Monte Carlo algorithms," *Journal of the American Statistical Association*, vol. 98, no. 461, pp. 169–177, 2003.
[23] S. M. Ross, *Stochastic processes*, 2nd ed. John Wiley & Son, 1996.
[24] D. Shah, D. N. C. Tse, and J. N. Tsitsiklis, "Hardness of low delay network scheduling," *IEEE Trans. on Information Theory*, 2009, submitted.
[25] D. Shah and J. Shin, "Delay optimal queue-based CSMA," in *Proceedings of ACM Sigmetrics*, 2010.
[26] ——, "Randomized scheduling algorithm for queueing networks," *The Annals of Applied Probability*, vol. 22, no. 1, pp. 128–171, 2012.
[27] J. Shin, D. Shah, and S. Rajagopalan, "Network adiabatic theorem: An efficient randomized protocol for contention resolution," in *Proceedings of ACM Sigmetrics*, 2009.
[28] V. Subramanian and M. Alanyali, "Delay performance of CSMA in networks with bounded degree conflict graphs," in *Proceedings of IEEE ISIT*, 2011.
[29] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling for maximum throughput in multihop radio networks," *IEEE Transactions on Automatic Control*, vol. 37, no. 12, pp. 1936–1949, December 1992.
[30] ——, "Stability properties of constrained queueing systems and scheduling for maximum throughput in multihop radio networks," *IEEE Trans. on Automatic Control*, vol. 37, no. 12, pp. 1936–1949, 1992.
[31] E. Vigoda, "A note on the Glauber dynamics for sampling independent sets," *Electronic Journal of Combinatorics*, vol. 8, no. 1, pp. 1–8, 2001.
[32] Y. Yi and M. Chiang, *Next-Generation Internet Architectures and Protocols*. Cambridge University Press, 2011, chapter 9: Stochastic network utility maximization and wireless scheduling.

APPENDIX A
PROOF OF THEOREM 1

Recall that when $q_v = 1/n$ for all v, the Markov chain $\{\boldsymbol{\sigma}(t,0)\}$ (by the original Glauber dynamics) is irreducible. This implies that there exists a finite-length sample path of the Markov chain starting from a configuration (or state) to any other configuration *without* self-transitions. We also know that among all possible $\boldsymbol{\sigma}, \boldsymbol{\sigma'} \in \Omega$ ($\boldsymbol{\sigma} \neq \boldsymbol{\sigma'}$), only $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma'}$ with $\|\boldsymbol{\sigma} - \boldsymbol{\sigma'}\| = 1$ lead to $Q_\beta(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) > 0$ *regardless of* any choice of $\{q_v\}$ and any value of $\beta \in [0,1]$. Specifically, for a given node v, consider any $(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) \in \Phi_v$. Then, $\gamma(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = q_v > 0$, and also $\alpha(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) > 0$ as can be seen from (17). We thus have

$$Q_\beta(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \gamma(\boldsymbol{\sigma}, \boldsymbol{\sigma'})\alpha(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) > 0.$$

We similarly obtain $Q_\beta(\boldsymbol{\sigma'}, \boldsymbol{\sigma}) > 0$. Therefore, there must exist the *same* finite-length sample path of the Markov chain $\{\boldsymbol{\sigma}(t,\beta)\}$ as above, and thus $\{\boldsymbol{\sigma}(t,\beta)\}$ is irreducible for any $\{q_v\}$ and $\beta \in [0,1]$. In addition, for a configuration $\boldsymbol{\sigma''}$ such that $\sigma''_v = 0$ and $\sum_{w \in N_v} \sigma''_w \geq 1$, by the generalized Glauber dynamics, we have $Q_\beta(\boldsymbol{\sigma''}, \boldsymbol{\sigma''}) \geq q_v > 0$. This implies that $\{\boldsymbol{\sigma}(t,\beta)\}$ is aperiodic [23], [5]. By noting that $|\Omega| < \infty$, $\{\boldsymbol{\sigma}(t,\beta)\}$ is a finite, irreducible, aperiodic Makro chain, or the Markov chain $\{\boldsymbol{\sigma}(t,\beta)\}$ is ergodic for any $\{q_v\}$ and $\beta \in [0,1]$.

As to the reversibility, for every $(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) \in \Phi_v$ ($v \in \mathcal{N}$), one can easily check

$$\frac{Q_\beta(\boldsymbol{\sigma}, \boldsymbol{\sigma'})}{Q_\beta(\boldsymbol{\sigma'}, \boldsymbol{\sigma})} = \frac{\pi(\boldsymbol{\sigma'})}{\pi(\boldsymbol{\sigma})} = \lambda_v,$$

which implies that $\pi(\boldsymbol{\sigma})Q_\beta(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = \pi(\boldsymbol{\sigma'})Q_\beta(\boldsymbol{\sigma'}, \boldsymbol{\sigma})$ *regardless of* any choice of $\{q_v\}$ and any value of $\beta \in [0,1]$. Thus, the Markov chain $\{\boldsymbol{\sigma}(t,\beta)\}$ is reversible with respect to $\boldsymbol{\pi}$ for any $\{q_v\}$ and $\beta \in [0,1]$.

For a given $\{q_v\}$, fix $\beta_1, \beta_2 \in [0,1]$ such that $\beta_1 \leq \beta_2$. It again suffices to prove that $\mathbf{Q}_{\beta_1} \preceq \mathbf{Q}_{\beta_2}$ for any configurations $\boldsymbol{\sigma}, \boldsymbol{\sigma'} \in \Omega$ with $\|\boldsymbol{\sigma} - \boldsymbol{\sigma'}\| = 1$. Fix node v, and then consider any $(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) \in \Phi_v$. One can see that $\alpha(\boldsymbol{\sigma}, \boldsymbol{\sigma'})$ in (16) is *monotone increasing* in $\beta \in [0,1]$. Also, by noting that $\gamma(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) = q_v$ is the same for all $\beta \in [0,1]$, we have $Q_{\beta_1}(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) \leq Q_{\beta_2}(\boldsymbol{\sigma}, \boldsymbol{\sigma'})$. Similarly, we can show that $Q_{\beta_1}(\boldsymbol{\sigma'}, \boldsymbol{\sigma}) \leq Q_{\beta_2}(\boldsymbol{\sigma'}, \boldsymbol{\sigma})$. Thus, we have $\mathbf{Q}_{\beta_1} \preceq \mathbf{Q}_{\beta_2}$.

APPENDIX B
PROOF OF THEOREM 2

Unless confusion arises, we omit the time t in all notations in this proof. We first present Lemma 3 which states that the SLEM of $\boldsymbol{Q}_\beta$ is upper-bounded by a value irrespective of β and only depends on the number of nodes n and the maximum queue length at the corresponding time. The bound on the SLEM of $\boldsymbol{Q}_\beta$ implies that the mixing time can be bounded irrespective of β. Lemma 3 is the only difference from the proof of [26], so the rest of the proof becomes equivalent to [26].

Lemma 3: If $W_{\max} \geq \exp(2)$, for all $\beta \in [0,1]$, SLEM of $\mathbf{Q}_\beta$ is bounded by:

$$e(\mathbf{Q}_\beta)) \leq 1 - \frac{1}{n^4 2^{2n+6} \exp((2n+4)f(W_{\max}))}. \qquad (23)$$

Proof: First, we note that the set of eigen values of $\mathbf{Q}_\beta^2$ is $\{(e_i)^2 : 1 \leq i \leq n\}$, where $\{e_i : 1 \leq i \leq n\}$ is the set of eigen values of $\mathbf{Q}_\beta$. By Cheeger's inequality for $\mathbf{Q}_\beta^2$, the SLEM of $\mathbf{Q}_\beta$ $e(\mathbf{Q}_\beta)$ satisfies:

$$e(\mathbf{Q}_\beta)^2 = \max\{(e_2)^2, (e_n)^2\} \leq 1 - \frac{\Phi^2}{2},$$

where Φ is the conductance of $\mathbf{Q}_\beta^2$. We remark that the conductance of $\mathbf{Q}_\beta^2$ is defined as:

$$\Phi \triangleq \min_{A \subset \Omega : \pi(A) \leq 1/2} \frac{\sum_{\boldsymbol{\sigma} \in A, \boldsymbol{\sigma'} \notin A} \pi(\boldsymbol{\sigma}) \mathbf{Q}_\beta^2(\boldsymbol{\sigma}, \boldsymbol{\sigma'})}{\pi(A)}.$$

Then, the conductance is bounded as:

$$\begin{aligned}
\Phi &\geq \min_{\mathbf{Q}_\beta^2(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) \neq 0} \pi(\boldsymbol{\sigma}) \mathbf{Q}_\beta^2(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) \\
&\geq \left(\min_{\boldsymbol{\sigma} \in \Omega} \pi(\boldsymbol{\sigma}) \right) \cdot \min_{v \in N} \frac{1}{n^2} \frac{(\lambda_v)^\beta (1+\lambda_v)^{1-\beta} - 1}{(\lambda_v)^{2\beta}(1+\lambda_v)^{2-2\beta}} \\
&\geq \frac{1}{2^n (\lambda_{\max})^n} \frac{1}{n^2} \frac{(\lambda_{\max})^\beta (1+\lambda_{\max})^{1-\beta} - 1}{(\lambda_{\max})^{2\beta}(1+\lambda_{\max})^{2-2\beta}} \\
&\geq \frac{1}{n^2 2^{n+2} \exp((n+2)f(W_{\max}))}. \qquad (24)
\end{aligned}$$

From (23) and (24), we get:

$$e(\mathbf{Q}_\beta) \leq \sqrt{1-\frac{\Phi^2}{2}} \leq 1-\frac{\Phi^2}{4} \quad (25)$$

$$\leq 1-\frac{1}{n^4 2^{2n+6}\exp((2n+4)f(W_{\max}))}. \quad (26)$$

This completes the proof. ∎

We now provide the sketch of the proof by dividing it into three steps, as presented in [26].

Step 1: Let π denote the stationary distribution of generalized Glauber dynamics with fugacities $\{\lambda_v, v\in\mathcal{N}\}$, where λ_v is given by (19). Then, we can prove the following

$$\mathbb{E}_\pi\{\sum_{v\in N} f(W_v)\sigma_v\} \geq \left(1-\frac{\epsilon}{4}\right)(\max_{\boldsymbol{x}\in\Omega}\sum_{v\in N} f(W_v)x_v) - O(1),$$

implying that scheduling of nodes (in the conflict graph) according the probability π is close to so-called *max-weight* [30].

Step 2: Let $\mu(t)$ be the distribution of the configurations (i.e., schedules) at time t, given the initial configuration $\boldsymbol{\sigma}(0)$ and queue length $\{W_v(0)\}$. The proof of ***Step 2*** shows that the configuration chosen according to $\mu(t)$ is very close to the configuration from the max-weight. To that end, from ***Step 1***, we have to prove that $\pi(t)$, which is the stationary distribution of the configurations defined by $\lambda_v(t)$ as set by (19), is very close to $\mu(t)$. This is because for a sufficiently slow queue changes by the condition $f(\cdot)=\log\log(\cdot)$, (i) the SLEM is bounded by the order of $1-1/\log(W_{\max})$ from (25), and (ii) from the relation between the mixing time and the SLEM, we can have that the mixing time is polylog with respect to $W_{\max}$. Note that this holds for any $\beta\in[0,1]$. Then, roughly speaking, from the facts that the changing speed of fugacities is the order of $1/W_{\max}$ and the mixing time is polylog in $W_{\max}$, we have that: the amount of fugacity changes is the order of $\text{polylog}(W_{\max})/W_{\max}$, which goes to 0 for a sufficient large $W_{\max}$. This means that by using $f(\cdot)=\log\log(\cdot)$, the fugacity change over the mixing time becomes negligible for a large queue regime. This implies that $\mu(t)\approx\pi(t)$.

Step 3: By using the two properties, the stability for all feasible arrival rate can be shown by defining the following Lyapunov function,

$$L(\boldsymbol{X}(t)) = \sum_{v\in N} F(W_v(t)),$$

where $\boldsymbol{X}(t)=(\{W_v\}(t),\boldsymbol{\sigma}(t))$ and $F(x)=\int_0^x f(y)dy$, and the fact that the configurations from oCSMA (by $\mu(t)$) is very close to those by the max-weight from ***Step 1*** and ***Step 2***.

Appendix C
Proof of Theorem 3

For notational convenience, we define two different hitting times to the set B for the chain $\{\boldsymbol{\sigma}(t,\beta)\}$:

$$H_\beta \triangleq \min\{t\geq 0 : \boldsymbol{\sigma}(t,\beta)\in B\} = \min\{t\geq 0 : S(t)=1\},$$
$$H_\beta^+ \triangleq \min\{t>0 : \boldsymbol{\sigma}(t,\beta)\in B\} = \min\{t>0 : S(t)=1\}.$$

Here, the only difference between these two is that the former counts the case of $\boldsymbol{\sigma}(0,\beta)\in B$, i.e., $H_\beta=0$ if $\boldsymbol{\sigma}(0,\beta)\in B$, while the latter does not. We then define the mean *first hitting* time to B from a *stationary start* as

$$\mathbb{E}_{\boldsymbol{\pi}}\{H_\beta\} \triangleq \sum_{\boldsymbol{\sigma}\in\Omega}\mathbb{E}\{H_\beta \mid \boldsymbol{\sigma}(0,\beta)=\boldsymbol{\sigma}\}\cdot\pi(\boldsymbol{\sigma}). \quad (27)$$

This is the average time steps required for node v until to be available for service, when starting from a randomly chosen time slot. We also define the mean *return* time to B as

$$\mathbb{E}_{\pi_B}\{H_\beta^+\} \triangleq \mathbb{E}\{H_\beta^+ \mid \boldsymbol{\sigma}(0,\beta)\in B\} = \sum_{\boldsymbol{\sigma}\in B}\mathbb{E}\{H_\beta^+ \mid \boldsymbol{\sigma}(0,\beta)=\boldsymbol{\sigma}\}\cdot\pi_B(\boldsymbol{\sigma}), \quad (28)$$

where $\pi_B(\boldsymbol{\sigma}) \triangleq \frac{\pi(\boldsymbol{\sigma})}{\pi_B}$ for $\boldsymbol{\sigma}\in B$, and the last equality is from the stationarity of $\{\boldsymbol{\sigma}(t,\beta)\}$ in t. This is same as the average recurrence time to the state 1 for $\{S(t)\}$, i.e., $\mathbb{E}\{\tau_i\} = \mathbb{E}_{\pi_B}\{H_\beta^+\}$ for all $i\geq 1$. We similarly define $\mathbb{E}_{\pi_B}\{(H_\beta^+)^2\} \triangleq \mathbb{E}\{(H_\beta^+)^2 \mid \boldsymbol{\sigma}(0,\beta)\in B\} = \mathbb{E}\{(\tau_i)^2\}$ for all i.

For any given $\{q_v\}$, first fix $\beta\in[0,1]$. From Kac's formula [1], [6], we have

$$\mathbb{E}\{\tau_i\} = \mathbb{E}_{\pi_B}\{H_\beta^+\} = \frac{1}{\pi_B}, \quad (29)$$

$$\mathbb{E}\{(\tau_i)^2\} = \mathbb{E}_{\pi_B}\{(H_\beta^+)^2\} = \frac{2\mathbb{E}_{\boldsymbol{\pi}}\{H_\beta\}+1}{\pi_B}. \quad (30)$$

From Theorem 1, $\mathbb{E}\{\tau_i\}=1/\pi_B$ remains the same for all β. We then consult the extremal characterization of the mean first hitting time to a subset for a reversible Markov chain in [1, Ch.3, Proposition 41],[††] which says that for a reversible chain $\{\boldsymbol{\sigma}(t,\beta)\}$ and a subset $B\subseteq\Omega$,

$$\mathbb{E}_{\boldsymbol{\pi}}\{H_\beta\} = \sup_g\left\{\frac{1}{\mathcal{E}_{\mathbf{Q}_\beta,\boldsymbol{\pi}}(g,g)} : -\infty<g<\infty, g(\cdot)=1 \text{ on } B, \text{ and } \sum_{\boldsymbol{\sigma}\in\Omega}\pi(\boldsymbol{\sigma})g(\boldsymbol{\sigma})=0\right\}. \quad (31)$$

Here, the Dirichlet form $\mathcal{E}_{\mathbf{Q}_\beta,\boldsymbol{\pi}}(g,g)$ for functions $g:\Omega\to\mathbb{R}$ excluding $g\equiv 0$, is defined by

$$\mathcal{E}_{\mathbf{Q}_\beta,\boldsymbol{\pi}}(g,g) \triangleq \frac{1}{2}\sum_{\boldsymbol{\sigma},\boldsymbol{\sigma'}\in\Omega}\pi(\boldsymbol{\sigma})Q_\beta(\boldsymbol{\sigma},\boldsymbol{\sigma'})\left(g(\boldsymbol{\sigma})-g(\boldsymbol{\sigma'})\right)^2.$$

Now, choose $0\leq\beta_1\leq\beta_2\leq 1$. From Theorem 1, we have

$$\mathcal{E}_{\mathbf{Q}_{\beta_1},\boldsymbol{\pi}}(g,g) = \frac{1}{2}\sum_{\boldsymbol{\sigma},\boldsymbol{\sigma'}\in\Omega}\pi(\boldsymbol{\sigma})Q_{\beta_1}(\boldsymbol{\sigma},\boldsymbol{\sigma'})\left(g(\boldsymbol{\sigma})-g(\boldsymbol{\sigma'})\right)^2 \leq \frac{1}{2}\sum_{\boldsymbol{\sigma},\boldsymbol{\sigma'}\in\Omega}\pi(\boldsymbol{\sigma})Q_{\beta_2}(\boldsymbol{\sigma},\boldsymbol{\sigma'})\left(g(\boldsymbol{\sigma})-g(\boldsymbol{\sigma'})\right)^2 = \mathcal{E}_{\mathbf{Q}_{\beta_2},\boldsymbol{\pi}}(g,g)$$

for any given function g. Together with (31), this implies that $\mathbb{E}_{\boldsymbol{\pi}}\{H_{\beta_1}\}\geq\mathbb{E}_{\boldsymbol{\pi}}\{H_{\beta_2}\}$. Therefore, from (30), $\mathbb{E}\{(\tau_i)^2\}$ is decreasing in $\beta\in[0,1]$.

[††]The extremal characterization is originally given for a random walk on a weighted graph (allowing to have self-loops) in [1, Ch.3, Proposition 41]. Note that any reversible chain can be represented as a random walk on a weighted graph. See [1, Ch.3] for more details.

APPENDIX D
PROOF OF THEOREM 4

Observe that

$$\lim_{t\to\infty}\frac{1}{t}\mathbb{VAR}\{\boldsymbol{S_t}\}=\lim_{t\to\infty}\frac{1}{t}\mathbb{VAR}\left\{\sum_{i=1}^{t}\mathbf{1}_{\{\boldsymbol{\sigma}(i,\beta)\in B\}}\right\}=\nu(\mathbf{Q}_\beta,f)$$

where the first equality is from (20), and the last equality is from (2) with $f(\boldsymbol{\sigma}(i,\beta))=\mathbf{1}_{\{\boldsymbol{\sigma}(i,\beta)\in B\}}$. Since $\mathbb{VAR}\{\boldsymbol{I_t}\}=\mathbb{VAR}\{\boldsymbol{A_t}\}+\mathbb{VAR}\{\boldsymbol{S_t}\}$, we have $\lim_{t\to\infty}\mathbb{VAR}\{\boldsymbol{I_t}\}/t=\nu^*+\nu(\mathbf{Q}_\beta,f)\in(0,\infty)$.

Since $\boldsymbol{I_t}$ is Gaussian for all large t, its limiting log-moment generating function $\lim_{t\to\infty}\frac{1}{t}\log\mathbb{E}\{e^{\theta\boldsymbol{I_t}}\}$ is the same as that of Gaussian $\boldsymbol{I_t}$ with mean $t(\mu-\pi_B)<0$ and variance $\mathbb{VAR}\{\boldsymbol{I_t}\}$. Thus, from the large deviation results in [8, Theorem 1], the asymptotic decay rate η becomes

$$\eta=-2\lim_{t\to\infty}\frac{\mathbb{E}\{\boldsymbol{I_t}\}/t}{\mathbb{VAR}\{\boldsymbol{I_t}\}/t}=\frac{2(\pi_B-\mu)}{\nu^*+\nu(\mathbf{Q}_\beta,f)}>0. \qquad (32)$$

For given μ and ν^* (or a given arrival process), from Lemma 1 and Theorem 1, $\nu(\mathbf{Q}_\beta,f)$ is decreasing in $\beta\in[0,1]$, and therefore, η is increasing in $\beta\in[0,1]$.