

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 [19], 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., [20]). 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., [4], [8], [12], [17]. 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., [20] 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. [15] 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 [16] 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 [9] 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. [3] 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 [16], [3], it was shown very recently [18] 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. Despite the same long-term throughput, their ‘second-order’ behavior can be quite different. This in turn leads to different

*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 [7].

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

Algorithm 1 Glauber Dynamics (at Time Slot t)

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

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) $\mathcal{N} = \{1, 2, \dots, n\}$ 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 σ a *configuration* of the nodes in G , which is given by $\sigma = \{\sigma_v, v \in \mathcal{N}\}$ 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 $\sigma_v + \sigma_w \leq 1$ 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 [7].

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 $\pi = \{\pi(\sigma)\}$ over Ω :

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

with a normalizing constant $Z = \sum_{\sigma \in \Omega} \prod_{v \in \mathcal{N}} \lambda_v^{\sigma_v}$. Note that $\lambda_v > 0$ for all v , ensuring that $\pi(\sigma) > 0$ for all $\sigma \in \Omega$. Specifically, $\sigma(t) = \{\sigma_v(t), v \in \mathcal{N}\}$ denotes the state of the Markov chain (or a feasible configuration by the Glauber dynamics) at time slot t . It is known that $\{\sigma(t)\}_{t \geq 0}$ is an irreducible, aperiodic Markov chain achieving the stationary distribution π in (1) on the finite state space Ω [7], [17], [16]. The Markov chain $\{\sigma(t)\}$ is also reversible with respect to π , i.e., $\pi(\sigma)Q(\sigma, \sigma') = \pi(\sigma')Q(\sigma', \sigma)$ for all $\sigma, \sigma' \in \Omega$, where $Q(\sigma, \sigma')$ is the transition probability from state σ to state σ' .

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

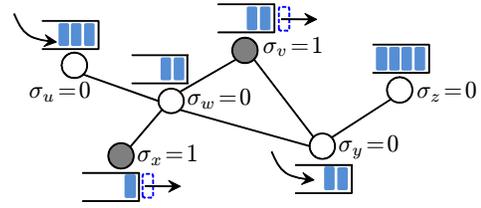


Fig. 1. Illustration of a wireless multihop scheduling driven by the Glauber dynamics over a conflict graph G with $|\mathcal{N}| = 6$. Links v, x are active while all others are silent, forming one instance of independent set on G .

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 [17], [16], 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 exponentially distributed period of time with mean $1 + \lambda_v$. It is worth noting that in the continuous-time setting, 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 [12], [3]. Although our target scenario is such continuous-time update for scheduling, in our subsequent analysis, time unit of interest is the unit of master clock ticks (so discrete time).

III. COMPARING REVERSIBLE MARKOV CHAINS

There are potentially many other (discrete-time) reversible Markov chains with the same π 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 π , they all guarantee the same throughput optimality, while their ‘second-order’ behavior can be quite different, leading to different queueing delay performance.

Mixing time has been a popular criterion to compare competing reversible Markov chains with the same stationary distribution. The mixing time of a (reversible) Markov chain indicates the speed of convergence to its stationary distribution and is mainly determined by the second largest eigenvalue modulus (SLEM) of its transition matrix [7]. Note that smaller SLEM leads to smaller (faster) mixing time. In this paper, however, we look at the comparison of reversible Markov chains from a different, but important perspective. This is done based on the following Peskun ordering and its relationship with efficiency ordering.

Definition 1 (Peskun ordering): [13] 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 $i, j \in \mathcal{S}$ ($i \neq j$). \square

Let $\{X(t)\}_{t \geq 0}$ and $\{\tilde{X}(t)\}_{t \geq 0}$ be irreducible Markov chains on a finite state space $\mathcal{S} = \{1, 2, \dots, n\}$ 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 $\pi = \{\pi(1), \pi(2), \dots, \pi(n)\}$. For a function $f : \mathcal{S} \rightarrow \mathbb{R}$, define an estimator $\hat{\mu}_m \triangleq \sum_{t=1}^m f(X(t))/m$ for $\mu = \mathbb{E}_\pi(f) = \sum_{i \in \mathcal{S}} f(i)\pi(i)$. It is well known that $\lim_{m \rightarrow \infty} \hat{\mu}_m = \mu$ for any function f with $\mathbb{E}_\pi(|f|) < \infty$ [5], [7]. The AVR of the estimate $\hat{\mu}_m$ is defined as

$$\nu(\mathbf{P}, f) \triangleq \lim_{m \rightarrow \infty} m \cdot \text{VAR}(\hat{\mu}_m), \quad (2)$$

which is *independent* of the distribution of the initial state $X(0)$ [13]. 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 π always becomes μ , as m goes to infinity. However, since the AVR 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 π , especially for the MCMC samplers [13], [11]. 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 $\text{VAR}_\pi(f) < \infty$ [11]. 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 π provides a sufficient condition for the efficiency ordering as follows.

Lemma 1: [13] If \mathbf{P} and $\tilde{\mathbf{P}}$ are reversible with respect to π , and $\mathbf{P} \preceq \tilde{\mathbf{P}}$, then $\nu(\mathbf{P}, f) \geq \nu(\tilde{\mathbf{P}}, f)$ for any f with $\text{VAR}_\pi(f) < \infty$. \square

It is worth nothing that efficiency ordering does not imply mixing time ordering in general, although there are related with each other [11]. (See more review on mixing time, Peskun ordering, and efficiency ordering in our technical report [6].)

[‡]This is formally captured by the Central Limit Theorem for an ergodic, finite Markov chain under $\text{VAR}_\pi(f) < \infty$, saying that $\sqrt{m}(\hat{\mu}_m - \mu)$ converges to a Gaussian random variable with zero mean and variance $\nu(\mathbf{P}, f)$ as $m \rightarrow \infty$ [5].

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  $\left(\frac{\lambda_v}{1+\lambda_v}\right)^{1-\beta} \min\{1, \lambda_v^\beta\}$ 
6:      $\sigma_v(t) = 0$ , otherwise.
7:   else
8:      $\sigma_v(t) = 0$  with probability  $\left(\frac{1}{1+\lambda_v}\right)^{1-\beta} \min\{1, 1/\lambda_v^\beta\}$ 
9:      $\sigma_v(t) = 1$ , otherwise.
10:  end if
11: else
12:    $\sigma_v(t) = 0$ 
13: end if
14: For any node  $w \in \mathcal{N} \setminus \{v\}$ :  $\sigma_w(t) = \sigma_w(t-1)$ 

```

IV. GENERALIZED GLAUBER DYNAMICS FOR SMALLER DELAY IN OPTIMAL CSMA

We now introduce a class of algorithms with a controllable parameter $\beta \in [0, 1]$, named *generalized Glauber dynamics*, as extensions of the (standard) Glauber dynamics in Algorithm 1. As shall be shown below, this generalization indicates that the Glauber dynamics is nothing but one of many possible ways to achieve the desired stationary distribution in (1) under the independent set constraints, while its extensions lead to *more efficient* reversible Markov chains in the sense of Peskun ordering (and efficiency ordering).

The generalized Glauber dynamics is summarized in Algorithm 2. This is achieved by judiciously employing a generalization of the procedures by Hastings [2] for constructing a reversible Markov chain with a given, desired stationary distribution, which was originally introduced for the development of an MCMC sampler. The details are omitted owing to space constraints but can be found in [6]. 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 [10] 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 $\beta = 0$. Note also that the node-selection probability distribution $\{q_v\}$ in Algorithm 2 can be arbitrary as long as $q_v > 0$ for all v and $\sum_{v \in \mathcal{N}} q_v = 1$.

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

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

Proof: See our technical report [6]. ■

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 briefly explain the throughput-optimality of the generalized Glauber dynamics. Theorem 1 says that the stationary distribution π of the Markov chain $\{\sigma(t, \beta)\}$ 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 [3], [18], 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 setting, especially for the original Glauber dynamics ($\beta=0$): the fugacity is now a function of time t , which is given by $\lambda_v(t) = \exp(f(W_v(t)))$ where f is some weighted 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 (and most popular) way is based on the time-scale decomposition under which the system quickly converges to its stationary regime before its dynamics changes (see, e.g., [4], [12]). Under this condition, Theorem 1 immediately implies that the generalized Glauber dynamics is throughput-optimal. On the other hand, the second approach in [17] is done without the time-scale decomposition when $q_v = 1/n$ for all v and $\beta=0$, but by choosing a proper weighted function f such as $f(\cdot) = \log \log(\cdot + e)$, so that $f(W_v(t))$ changes much slower than the system dynamics. The proof technique in [17] can be similarly used to establish the throughput-optimality of the generalized Glauber dynamics. A rigorous treatment on the throughput-optimality without the time-scale decomposition is another research topic beyond the scope of this paper.

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 [3], [18], 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 $\{\sigma(t, \beta)\}_{t \geq 0}$ is in its stationary regime, i.e., $\mathbb{P}\{\sigma(t, \beta) = \sigma\} = \pi(\sigma)$ for all $t \geq 0$. 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,

$$S_v(t) = \begin{cases} 1 & \text{if node } v \text{ is available for service, i.e., } \sigma_v(t) = 1 \\ 0 & \text{otherwise} \end{cases} \\ = \mathbf{1}_{\{\sigma(t, \beta) \in B_v\}} \quad (3)$$

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

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

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

[§]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 [1]. Such processes include not only Bernoulli arrivals, but correlated arrivals such as auto-regressive processes whose autocorrelation functions are summable.

Proof: See our technical report [6]. ■

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 π 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 [14]. 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 $\sigma(T_i, \beta) \in B$ 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}_β 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:

$$W(t+1) = \max\{0, W(t) + A(t+1) - S(t+1)\}. \quad (4)$$

From the large deviation theory, in considerable generality, the tail distribution of the steady-state queue length W is asymptotically exponential [1] with the asymptotic decay rate η given by

$$\eta = \lim_{x \rightarrow \infty} -\frac{1}{x} \log \mathbb{P}\{W > x\} > 0. \quad (5)$$

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

Theorem 3: Suppose that the distribution of \mathbf{I}_t is Gaussian for large t with $\lim_{t \rightarrow \infty} \mathbb{V}\mathbb{A}\mathbb{R}\{\mathbf{A}_t\}/t = \nu^* < \infty$. Then, η in (5) is increasing in $\beta \in [0, 1]$. □

Proof: See our technical report [6]. ■

Since $\mathbf{I}_t = \sum_{i=1}^t (A(i) - S(i))$ is the sum of t random variables, as long as their dependency over i is not so strong, it is reasonable to assume that \mathbf{I}_t is roughly Gaussian for large t . Then, Theorem 3 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 3 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

¶ 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.

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 2 and 3, 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. Due to space constraints, all simulation results are omitted and can be found in [6].

V. 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.

REFERENCES

- [1] 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.
- [2] W. K. Hastings, “Monte Carlo sampling methods using markov chains and their applications,” *Biometrika*, vol. 57, no. 1, pp. 97–109, 1970.
- [3] 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.
- [4] 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.
- [5] G. L. Jones, “On the Markov chain central limit theorem,” *Probability Surveys*, vol. 1, pp. 299–320, 2004.
- [6] C.-H. Lee, D. Y. Eun, S.-Y. Yun, and Y. Yi, “From Glauber dynamics to Metropolis algorithm: smaller delay in optimal CSMA,” North Carolina State University, Tech. Rep., Jan. 2012, available at <http://www4.ncsu.edu/~dyeun/pub/techrep-csma12.pdf>.
- [7] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*. American Mathematical Society, 2009.
- [8] 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.
- [9] M. Lotfinezhad and P. Marbach, “Throughput-optimal random access with order-optimal delay,” in *Proceedings of IEEE Infocom*, 2011.
- [10] 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.
- [11] A. Mira, “Ordering and improving the performance of Monte Carlo Markov chains,” *Statistical Science*, vol. 16, no. 4, pp. 340–350, 2001.
- [12] 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.
- [13] P. H. Peskun, “Optimum Monte-Carlo sampling using Markov chains,” *Biometrika*, vol. 60, pp. 607–612, 1973.
- [14] S. M. Ross, *Stochastic processes*, 2nd ed. John Wiley & Son, 1996.
- [15] D. Shah, D. N. C. Tse, and J. N. Tsitsiklis, “Hardness of low delay network scheduling,” *IEEE Trans. on Information Theory*, 2009, submitted.
- [16] D. Shah and J. Shin, “Delay optimal queue-based CSMA,” in *Proceedings of ACM Sigmetrics*, 2010.
- [17] J. Shin, D. Shah, and S. Rajagopalan, “Network adiabatic theorem: An efficient randomized protocol for contention resolution,” in *Proceedings of ACM Sigmetrics*, 2009.
- [18] V. Subramanian and M. Alanyali, “Delay performance of CSMA in networks with bounded degree conflict graphs,” in *IEEE ISIT*, 2011.
- [19] L. Tassioulas 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.
- [20] Y. Yi and M. Chiang, *Next-Generation Internet Architectures and Protocols*. Cambridge University Press, 2011, chapter 9: Stochastic network utility maximization and wireless scheduling.