On Scheduling Algorithms Robust to Heavy-Tailed Traffic

Krishna Jagannathan
Department of Electrical Engineering
IIT Madras
Chennai, India 600036
Email: krishnaj@ee.iitm.ac.in

Libin Jiang
California Institute of Technology
Pasadena, CA 91106
Email: libinj@caltech.edu

Eytan Modiano
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: modiano@mit.edu

***Abstract*—Adaptive CSMA algorithms have attracted considerable attention, due to their throughput optimality, utility maximizing properties, and amenability to simple distributed implementation. In this paper, we investigate the impact of heavy-tailed traffic on the performance of a queue length based adaptive CSMA algorithm. Specifically, we consider two conflicting links that share a channel using adaptive CSMA. One of the links serves heavy-tailed traffic, while the other link serves light-tailed traffic. Our main contribution is in demonstrating a threshold phenomenon in the relationship between the arrival rates and the queue backlog distributions. In particular, we show that when the arrival rate of the light-tailed traffic is less than a threshold value, the light-tailed traffic experiences a light-tailed queue backlog at steady state, whereas for arrival rates above the same threshold, the light-tailed traffic experiences a heavy-tailed queue backlog.**

Comparing this result to a corresponding result for max-weight scheduling [8], we conclude that adaptive CSMA is potentially more robust to bursty traffic, compared to max-weight scheduling.

I. Introduction

In the context of communication networks, link scheduling for maximum throughput is a well studied problem. Maximum weight scheduling, first proposed in [17], [18], forms the basis for most of the literature on this topic. Unfortunately, finding the maximum weight schedule is in general NP-complete, and requires global information of the queue lengths, making such algorithms difficult to implement. Randomized versions of the max-weight algorithm such as the ones proposed in [5], [16] are computationally simpler, but still require centralized control. In [12], a distributed throughput optimal algorithm, based on randomized scheduling and gossip-based information exchange is proposed. This algorithm, though simple and distributed, requires network-wide information exchange for each new scheduling decision.

Greedy approximations to max-weight scheduling which are simpler to implement in a distributed fashion have been proposed in [1], [2], [19], but these approaches often lead to guaranteed stability only in a fraction of the stability region. Sufficient conditions on the network topology are derived in [3], [11] for throughput optimality of greedy algorithms. Random access techniques such as ALOHA and Carrier Sense Multiple Access (CSMA) have also been widely studied for many years, but their potential optimality properties were not well understood until recently.

In a series of recent papers [10], [13]–[15], adaptive CSMA based algorithms have been proposed to achieve maximum throughput. The key idea of adaptive CSMA scheduling is to adjust the transmission aggressiveness (TA) of each link according to its local queue length. Specifically, when the queue length of a link increases, the link transmits more aggressively by using smaller backoff time or larger transmission time; and the link does the opposite when its queue length decreases.[1] Adaptive CSMA based algorithms are expected to find wide-spread applications in wireless networks, owing to their optimality, extreme simplicity of operation, and inherent scalability.

In this paper, we study the performance of an adaptive CSMA algorithm in a queueing network that serves highly heterogeneous traffic. Such a study is motivated by the fact that packet switched networks serve a wide variety of bursty as well as benign traffic sources, and otherwise desirable network control policies (such as adaptive CSMA) must be evaluated for robustness under such varied traffic characteristics. To this end, we consider a simple queueing network consisting of two conflicting links that access a server using adaptive CSMA. One of the links serves heavy-tailed traffic, while the other link serves light-tailed traffic. We prove the existence of a threshold arrival rate λ^* such that when the arrival rate of the light-tailed traffic is less than λ^*, the light queue has light-tailed queue backlog in steady-state. When the arrival rate of the light-tailed traffic exceeds λ^*, the light-tailed traffic suffers a heavy-tailed queue backlog in steady-state.

In a closely related paper [8] regarding the performance of max-weight scheduling in the presence of heavy-tailed traffic, it was shown that maximum weight scheduling induces heavy-tailed delays on the light-tailed traffic, for any non-zero traffic arrival rates. Thus, max-weight scheduling may be undesirable in networks carrying bursty or heterogeneous traffic, its throughput optimality notwithstanding. In comparison, adaptive CSMA induces heavy-tailed delays on the light-tailed traffic only when the arrival rate of the light-tailed traffic

[1]Despite the intuitive idea, the proof to establish throughput-optimality is sophisticated. It utilizes the product-form stationary distribution of the transmission states under CSMA (roughly equivalent to the so-called Glauber Dynamics in statistical physics), and relates that to either a maximal-entropy problem [10] or max-weight scheduling [15].

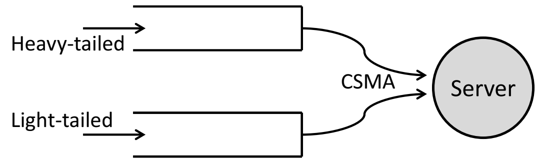

Fig. 1. A queueing system with two conflicting links. One of the links receives heavy-tailed traffic, while the other receives light-tailed traffic.

is 'high.' In other words, adaptive CSMA seems to exhibit superior robustness to heterogeneous and bursty traffic sources compared to max-weight scheduling, in addition to being much simpler to implement.

The remainder of this paper is organized as follows. In Section II, we describe the system model and the requisite mathematical preliminaries. In Section III, we state and prove our main results. In Section IV, we discuss the implications of our main result and draw conclusions.

II. System Description and Preliminaries

Consider a system consisting of two queues, sharing a common server. The queues access the server through conflicting wireless links, i.e., the two queues cannot be served at the same time. Time is slotted, with bursts of packets arriving at random to each queue at the beginning of each slot. The server is capable of serving one packet per time slot. Although we postpone the precise assumptions on the traffic statistics, let us loosely say that one of the queues receives light-tailed traffic, while the other receives heavy-tailed traffic (Fig. 1). We will refer to the two queues as the light queue and the heavy queue, respectively. Before we specify the precise assumptions on the arrival processes, we pause to present some relevant definitions.

Heavy-tailed and light-tailed random variables

Definition 1: A non-negative random variable X is said to be *light-tailed* if there exists $\theta > 0$ for which $\mathbb{E}\left[\exp(\theta X)\right] < \infty$. A random variable is *heavy-tailed* if it is not light-tailed. In other words, a light-tailed random variable is one that has a well defined moment generating function in a neighborhood of the origin. The complementary distribution function of a light-tailed random variable decays at least exponentially fast. Heavy-tailed random variables are those which have complementary distribution functions that decay slower than any exponential. We now define the tail-coefficient of a random variable.

Definition 2: The *tail coefficient* of a random variable X is defined by

$$C_X = \sup\{c \geq 0 \mid \mathbb{E}\left[X^c\right] < \infty\}.$$

In words, the tail coefficient is the threshold where the power moment of a random variable starts to blow up. Note that the tail coefficient of a light-tailed random variable is infinite. On the other hand, the tail coefficient of a heavy-tailed random variable may be infinite (e.g., log-normal distribution) or finite (e.g., Pareto distribution). In this paper, we restrict our attention to the class of heavy-tailed random variables which have a finite tail coefficient.

We now state the precise assumptions on the arrival processes.

Assumptions about the arrival processes

1) The arrival processes to the two queues are independent of each other.
2) The arrival process to each queue is independent and identically distributed (i.i.d.) from slot-to-slot.
3) The number of packet arrivals to the light queue during any slot is a light-tailed random variable, with mean λ_L.
4) The number of packet arrivals to the heavy queue during any slot is a heavy-tailed random variable, with tail coefficient C_H ($1 < C_H < \infty$), and mean λ_H.

CSMA scheduling algorithm

Assume that each link runs a CSMA-based algorithm to transmit packets. Without loss of generality, we assume that the discrete-time CSMA algorithm in [13] is utilized, where time is slotted and one unit of data can be transmitted in one slot. We note, however, that our main results are not affected if the continuous-time CSMA in [10] are adopted. As mentioned before, the key idea of adaptive CSMA scheduling is to adjust the transmission aggressiveness (TA) of each link according to its queue length. In this paper, we assume that the TA is adjusted at every frame boundary, where each frame has T time slots. Specifically, at the beginning of time slot jT ($j = 0, 1, 2, \dots$)[2], the TA of link k ($k = H, L$) is updated as according to,

$$r_k[j] = \min\left\{\frac{\alpha}{T}Q_k[j], r_{\max}\right\} \tag{1}$$

where $Q_k[j]$ is the queue length of link k at the beginning of time slot jT, $\alpha < 1$ is a constant, and the constant parameter $r_{\max}$ is the maximum possible TA.

If the TA's of the two links are fixed at r_H and r_L, then in the steady state of the CSMA Markov chain, the service rate of link k is given by [13]

$$s_k(\mathbf{r}) = \frac{\exp(r_k)}{1 + \exp(r_H) + \exp(r_L)},\ \ k = H, L. \tag{2}$$

Let us refer to the interval between the beginnings of time slots jT and $(j+1)T$ as "frame j". It is not difficult to see that $\{\mathbf{Q}[j], \sigma[j]\}_{j=0,1,\dots}$ forms a Markov chain [9], where $\sigma[j]$ is the transmission state (indicating whether link k ($k = H, L$) is transmitting) just before slot jT. In frame j, it has been shown in [9] that the "empirical service rate" $\tilde{s}_k[j]$ of link k (i.e., the units of data transmitted during frame j divided by T) satisfies

$$|\mathbb{E}_j\{\tilde{s}_k[j]\} - s_k(\mathbf{r}[j])| \leq b/T \tag{3}$$

for some constant $b > 0$, where $\mathbb{E}_j\{\cdot\}$ is a shorthand for the conditional expectation $\mathbb{E}\{\cdot|\mathbf{Q}[j], \sigma[j]\}$. Let us also define the quantity

$$\lambda^* = \frac{\exp(r_{\max})}{1 + 2\exp(r_{\max})}. \tag{4}$$

[2]Here, the indices of time slots start with 0

In light of (2), λ^* can be understood as the service rate afforded to each of the queues, when both queues attempt transmission with the maximum permissible TA of $r_{\max}$. The quantity λ^* will play an important role in our analysis.

The queue dynamics is given by

$$Q_k[j+1] = \{Q_k[j] - T \cdot \tilde{s}_k[j]\}_+ + a_k[j],\ k = H, L,$$

where $a_k[j]$ is the total number of packet arrivals to link k during frame j. Recall that $a_L[j]$ is assumed to be a light-tailed random variable, and $a_H[j]$ is assumed to be heavy-tailed.

III. Main results

Our first result shows that the adaptive CSMA algorithm can stabilize arrival rates that lie arbitrarily close to the largest possible stability region boundary, as long as the parameters $r_{\max}, T$ and α are chosen appropriately.

Theorem 1: Suppose that $r_{\max} > 1$, and that

$$\lambda_L + \lambda_H \leq \frac{\exp(r_{\max})}{1 + \exp(r_{\max}) + \exp(1)} - 2\epsilon$$

where $\epsilon > 0$. Choose $T \geq 2b/\epsilon$. Then both queues are stable.

Remark: Clearly, the stability region approaches the largest possible region ($\lambda_L + \lambda_H < 1$) as $r_{\max} \to \infty$ and $\epsilon \to 0$.

Proof: Consider the Lyapunov function $L[j] := Q_L[j] + Q_H[j] \geq 0$. We will show that if $L[j] \geq 2r_{\max}T/\alpha$, then the following holds:

$$\mathbb{E}_j\{L[j+1] - L[j]\} \leq -T\epsilon. \tag{5}$$

Suppose that $L[j] \geq 2r_{\max}T/\alpha$, then $Q_k[j] \geq r_{\max}T/\alpha$ for some k. Without loss in generality, assume that $Q_L[j] \geq r_{\max}T/\alpha$. So $r_L[j] = r_{\max}$. Since $r_{\max} > 1$ and $\alpha < 1$, one has $Q_L[j] \geq T$. Therefore, $Q_L[j+1] = Q_L[j] - T \cdot \tilde{s}_L[j] + a_L[j]$.

Now consider two cases.

Case 1: If $Q_H[j] \geq T$, then $Q_H[j+1] = Q_H[j] - T \cdot \tilde{s}_H[j] + a_H[j]$. Consequently,

$$\mathbb{E}_j\{L[j+1]\} = L[j] + T\{\lambda_L + \lambda_H - \mathbb{E}_j\{\tilde{s}_L[j] + \tilde{s}_H[j]\}\}.$$

With $T \geq 2b/\epsilon$, the RHS of (3) is less than or equal to $\epsilon/2$. Therefore,

$$\begin{aligned}\mathbb{E}_j\{\tilde{s}_L[j] + \tilde{s}_H[j]\} &\geq \frac{\exp(r_L[j]) + \exp(r_H[j])}{1 + \exp(r_L[j]) + \exp(r_H[j])} - \epsilon \\ &\geq \frac{\exp(r_{\max})}{1 + \exp(r_{\max})} - \epsilon.\end{aligned}$$

So (5) holds.

Case 2: If $Q_H[j] < T$, we have $Q_H[j+1] \leq Q_H[j] + a_H[j]$. So

$$\mathbb{E}_j\{L[j+1]\} \leq L[j] + T\{\lambda_L + \lambda_H - \mathbb{E}_j\{\tilde{s}_L[j]\}\}.$$

Since $Q_H[j] < T$ we also have $r_H[j] < \alpha < 1$. Therefore

$$\begin{aligned}s_L(\mathbf{r}[j]) &= \frac{\exp(r_{\max})}{1 + \exp(r_{\max}) + \exp(r_H[j])} \\ &\geq \frac{\exp(r_{\max})}{1 + \exp(r_{\max}) + \exp(1)}.\end{aligned}$$

So (5) still holds.

Therefore, the Lyapunov function has a negative drift whenever $L[j] \geq 2r_{\max}T/\alpha$. Combined with the fact that $\mathbb{E}_j\{L[j+1]\} - L[j]$ is bounded, by the Foster-Lyapunov Criterion, we conclude that the queues are stable. ■

In the next theorem, we show that when the arrival rate to the light queue is less than λ^*, the steady-state queue occupancy at the light queue is light-tailed.

Theorem 2: Suppose that $\lambda_L \leq \lambda^* - \epsilon$ where $\epsilon > 0$. Choose $T \geq 2b/\epsilon$. Then the steady-state queue occupancy Q_L of the light queue is light-tailed.

Proof: We show that whenever $Q_L[j] > r_{\max}\frac{T}{\alpha}$, the light queue has negative drift. Suppose that $Q_L[j] > r_{\max}\frac{T}{\alpha}$. Then, according to (1), $r_L[j] = r_{\max}$. Since $r_H[j] \leq r_{\max}$, by (2), we have

$$s_L(\mathbf{r}[j]) \geq \frac{\exp(r_{\max})}{1 + 2\exp(r_{\max})} = \lambda^*.$$

Therefore, by (3),

$$\begin{aligned}\mathbb{E}_j\{\tilde{s}_L[j]\} &\geq s_L(\mathbf{r}[j]) - b/T \\ &\geq \frac{\exp(r_{\max})}{1 + 2\exp(r_{\max})} - \frac{\epsilon}{2}.\end{aligned}$$

So $\lambda_L - \mathbb{E}_j\{\tilde{s}_L[j]\} \leq -\epsilon/2 < 0$ (i.e., there is a negative drift). Since the arrival process is light-tailed, we can invoke Theorem 2.3 (Eq. (2.8)) in [6], to conclude that the light queue distribution is light-tailed in steady-state. ■

Our next result is a converse to Theorem 2. In other words, we show that when the arrival rate of the light-tailed traffic is greater than the threshold value λ^*, the steady-state queue occupancy at the light queue is heavy-tailed.

Theorem 3: Suppose $\lambda_L > \lambda^*$, and that the CSMA parameters are chosen such that the system is stable. Then, the steady-state occupancy of the light queue is heavy-tailed. with tail coefficient at most $C_H - 1$.

Proof: (Outline) We need to show that for any $\delta > 0$,

$$\mathbb{E}\left[Q_L^{C_H - 1 + \delta}\right] = \infty. \tag{6}$$

Since the formal proof of this result is rather lengthy and involved, we will present here an informal proof outline, and indicate how the arguments can be formalized. In our informal outline, we will 'show' that

$$\lim_{j \to \infty} \mathbb{E}\left[Q_L[j]^{C_H - 1 + \delta}\right] = \infty. \tag{7}$$

The above (Eq. 7) is the limit of the expectation of a sequence of random variables, whereas what we really want to show in (6) is regarding the expectation of the limiting random variable Q_L. Although it is by no means obvious that the limit and the expectation can be interchanged here, we will ignore this technical detail for now (see Chapter 5 of [7]).

To prove (6), we first note that the time intervals between two successive frame boundaries at which the system empties constitute renewal intervals. Let us denote by T_R the random variable representing the number of frames in a renewal interval. Since the system is stable, we have $\mathbb{E}[T_R] < \infty$.

Let us now define the renewal reward function

$$R[j] = Q_L[j]^{C_H - 1 + \delta}.$$

By the key renewal theorem [4],

$$\lim_{j \to \infty} \mathbb{E}\left[R[j]\right] = \frac{\mathbb{E}\left[R\right]}{\mathbb{E}\left[T_R\right]},$$

where $\mathbb{E}[R]$ denotes the expected reward accumulated over a renewal interval, and $\mathbb{E}[T_R] < \infty$. It is therefore enough to show that[3]

$$\mathbb{E}\left[\sum_{i=0}^{T_R} Q_L[i]^{C_H - 1 + \delta}\right] = \infty. \tag{8}$$

To see why the expected reward over a renewal interval is infinite, we consider a busy period that commences with a burst of size B that arrives at the heavy queue. This event has non-zero probability. We next compute the expected reward over a renewal interval, conditioned on the busy period starting with a burst of size $B > B_0$, where B_0 is chosen suitably large.

For large enough B_0, the TA of the heavy queue will saturate at $r_{\max}$ starting from frame 1, and remain at $r_{\max}$ until the occuapncy of the heavy queue falls to $\frac{T r_{\max}}{\alpha}$. In other words, the TA of the heavy queue will remain at $r_{\max}$ for $O(B)$ time frames. During this $O(B)$ time interval, the light queue receives service at a rate of λ^* at best, according to (2), (3) and (4). Since the arrival rate λ_L is greater than λ^*, the light queue will build up with high probability during this $O(B)$ interval, at the rate of $\lambda_L - \lambda^*$. Indeed, the light queue will build up with high probability to an $O(B)$ level, over a duration lasting $O(B)$ time frames.

Conditioned on the above event, the reward is at least $O(B^{C_H - 1 + \delta})$ for $O(B)$ time slots, with high probability. We can then show that for some constant $\kappa > 0$,

$$\begin{aligned}\mathbb{E}\left[\sum_{i=0}^{T_R} Q_L[i]^{C_H - 1 + \delta}\right] &\geq \mathbb{E}\left[\kappa B \cdot B^{C_H - 1 + \delta}\right] \\ &= \mathbb{E}\left[\kappa B^{C_H + \delta}\right] = \infty,\end{aligned}$$

where the last expectation is infinite because the initial burst size has tail coefficient equal to C_H. We have thus shown (8), and from the key renewal theorem, (7) follows. In order to prove that the limit and expectation can legitimately be interchanged in (7), one needs to use a truncation argument, followed by repeated use of the monotone convergence theorem and the dominated convergence theorem. The above arguments can be made precise by adopting the methodology used in proving [7, Proposition 5.4]. ∎

IV. Discussion

In this section, we discuss the results proved in this paper, and compare them with the corresponding results derived in an earlier paper on max-weight scheduling. We considered a system consisting of two conflicting wireless links, where one of the links serves heavy-tailed traffic, while the other receives light-tailed traffic. We assumed that the links access a shared server using a recently proposed adaptive CSMA algorithm, which is throughput maximizing. We proved the existence of a threshold arrival rate λ^*, such that when the arrival rate of the light-tailed traffic is less than λ^*, the light queue has light-tailed queue backlog in steady-state. When the arrival rate of the light-tailed traffic exceeds λ^*, the light-tailed traffic suffers a heavy-tailed queue backlog in steady-state. Since λ^* is close to one half for large $r_{\max}$, our result is tantamount to saying that adaptive CSMA induces heavy-tailed backlog for the light queue only if the light-tailed traffic is responsible for more than half the total supportable traffic rate in the system.

In a recent paper [8], an analogous analysis was carried out for maximum weight scheduling in the presence of heavy-tailed traffic. Specifically, it was shown in [8] that maximum weight scheduing and its generalized version called max-weight-α scheduling induce a heavy-tailed queue backlog at the light queue. This is true for *all non-zero arrival rates of the heavy-tailed and light-tailed traffic*. Furthermore, it was shown that max-weight scheduling induces the worst possible asymptotics on the light queue among all non-idling policies. In comparison, the adaptive CSMA algorithm can perform better in terms of the delay faced by the light-tailed traffic. In addition, adaptive CSMA has the same throughput optimality property as max-weight scheduling.

As explained in [8], max-weight scheduling induces very poor queue backlog on the light-tailed link because large burst arrivals to the heavy-tailed link can starve the light queue for extended durations. On the other hand, with adaptive CSMA, all links have bounded TA values. As a result, even when large bursts arrive at a link carrying heavy-tailed traffic, the link cannot take over the server by attempting to transmit with arbitrary aggressiveness. This has the effect of 'shielding' the light-tailed traffic from the large bursts, at least when the arrival rate is smaller than the threshold value. Furthermore, adaptive CSMA does not need any a priori information about traffic statistics. In contrast, the policies proposed in [8] to mitigate the effect of heavy-tailed traffic need to have a priori information about which flow is heavy-tailed.

We wish to point out here that the cap on the TA values of each link in adaptive CSMA was originally intended as a mechanism to bound the mixing time of the CSMA Markov chain. In other words, capped TA values imply bounded 'fugacities' in the underlying Glauber dynamics, which leads to bounded mixing time. However in our context, the bounded TA values help in another way as well, by preventing the heavy-tailed link from attempting too aggressively.

In conclusion, our study suggests that adaptive CSMA has the potential to be more robust than max-weight scheduling, in queueing networks that serve highly bursty, heterogeneous traffic.

Acknowledgment

This work was supported in part by NSF grants CNS-0626781 and CNS-0915988, and by ARO Muri grant number

[3]Without loss of generality, we have considered a busy period that commences at time 0.

W911NF-08-1-0238.

REFERENCES

[1] P. Chaporkar, K. Kar, X. Luo, and S. Sarkar, "Throughput and fairness guarantees through maximal scheduling in wireless networks," *IEEE Transactions on Information Theory*, vol. 54, no. 2, p. 572, 2008.

[2] J. G. Dai and B. Prabhakar, "The throughput of data switches with and without speedup," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2. IEEE, 2000, pp. 556–564.

[3] A. Dimakis and J. Walrand, "Sufficient conditions for stability of longest-queue-first scheduling: Second-order properties using fluid limits," *Advances in Applied probability*, vol. 38, no. 2, pp. 505–521, 2006.

[4] R. Gallager, *Discrete stochastic processes*. Kluwer, 1996.

[5] P. Giaccone, B. Prabhakar, and D. Shah, "Randomized scheduling algorithms for high-aggregate bandwidth switches," *Selected Areas in Communications, IEEE Journal on*, vol. 21, no. 4, pp. 546–559, 2003.

[6] B. Hajek, "Hitting-time and occupation-time bounds implied by drift analysis with applications," *Advances in Applied probability*, pp. 502–525, 1982.

[7] K. Jagannathan, "Asymptotic performance of queue-length-based network control policies," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, USA, 2010.

[8] K. Jagannathan, M. Markakis, E. Modiano, and J. Tsitsiklis, "Queue length asymptotics for generalized max-weight scheduling in the presence of heavy-tailed traffic," in *IEEE INFOCOM*, Shanghai, China, 2011.

[9] L. Jiang, M. Leconte, J. Ni, R. Srikant, and J. Walrand, "Fast mixing of parallel glauber dynamics and low-delay csma scheduling," in *INFOCOM, 2011 Proceedings IEEE*. IEEE, 2011, pp. 371–375.

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

[11] C. Joo, X. Lin, and N. Shroff, "Understanding the capacity region of the greedy maximal scheduling algorithm in multihop wireless networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 17, no. 4, pp. 1132–1145, 2009.

[12] E. Modiano, D. Shah, and G. Zussman, "Maximizing throughput in wireless networks via gossiping," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 34, no. 1. ACM, 2006, pp. 27–38.

[13] 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 *INFOCOM, 2010 Proceedings IEEE*. IEEE, 2010, pp. 1–5.

[14] S. Rajagopalan and D. Shah, "Distributed algorithm and reversible network," in *Information Sciences and Systems, 2008. CISS 2008. 42nd Annual Conference on*. IEEE, 2008, pp. 498–502.

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

[16] L. Tassiulas, "Linear complexity algorithms for maximum throughput in radio networks and input queued switches," in *INFOCOM'98. Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2. Ieee, pp. 533–539.

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

[18] ——, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Transactions on Information Theory*, vol. 39, no. 2, pp. 466–478, 1993.

[19] X. Wu, R. Srikant, and J. Perkins, "Queue-length stability of maximal greedy schedules in wireless networks," in *Proceedings of Information Theory and Applications Inaugural Workshop*, 2006, pp. 6–10.