

Backlog-Based Random Access in Wireless Networks: Fluid Limits and Delay Issues*

Niek Bouman[†], Sem Borst^{†§}, Johan van Leeuwen[†], Alexandre Proutiere[‡]

[†]Eindhoven University of Technology, [‡]KTH Royal Institute of Technology, [§]Alcatel-Lucent Bell Labs

Abstract—We explore the spatio-temporal congestion dynamics of wireless networks with backlog-based random-access mechanisms. While relatively simple and inherently distributed in nature, suitably designed backlog-based access schemes provide the striking capability to match the optimal throughput performance of centralized scheduling algorithms in a wide range of scenarios. In the present paper, we show that the specific activity functions for which maximum stability has been established, may however yield excessive queue lengths and delays. The results reveal that more aggressive/persistent access schemes can improve the delay performance, while retaining the maximum stability guarantees in a rich set of scenarios. In order to gain qualitative insights and examine stability properties we will investigate fluid limits where the system dynamics are scaled in space and time. As it turns out, several distinct types of fluid limits can arise, exhibiting various degrees of randomness, depending on the structure of the network, in conjunction with the form of the activity functions. We further demonstrate that, counter to intuition, additional interference may improve the delay performance in certain cases. Simulation experiments are conducted to illustrate and validate the analytical findings.

I. INTRODUCTION

Emerging wireless mesh networks typically lack any centralized control entity for regulating access and coordinating transmissions. Instead, these networks vitally rely on the individual nodes to operate autonomously and to efficiently share the medium in a distributed fashion. This requires the nodes to schedule their individual transmissions and decide on the use of a shared medium based on knowledge that is locally available or only involves limited exchange of information. A popular mechanism for distributed medium access control is provided by the so-called Carrier-Sense Multiple-Access (CSMA) protocol, various incarnations of which are implemented in IEEE 802.11 networks. In the CSMA protocol each node attempts to access the medium after a certain back-off time, but nodes that sense activity of interfering nodes freeze their back-off timer until the medium is sensed idle.

While the CSMA protocol is fairly easy to understand at a local level, the interaction among interfering nodes gives rise to quite intricate behavior and complex throughput characteristics on a macroscopic scale. In recent years relatively parsimonious models have emerged that provide a useful tool in evaluating the throughput characteristics of CSMA-like networks. These models essentially assume that the interference constraints can be represented by a general conflict graph,

*This work was supported by Microsoft Research through its PhD Scholarship Programme.

and that the various nodes activate asynchronously whenever none of their neighbors are presently active. Such models were originally considered by Boorstyn *et al.* [3], and pursued in the context of IEEE 802.11 systems by Wang & Kar [24] and Durvy *et al.* [5], with several extensions and refinements in [4], [8]. Although the representation of the IEEE 802.11 back-off mechanism in the above-mentioned models is far less detailed than in the landmark work of Bianchi [2], the general interference graph offers greater versatility and covers a broad range of topologies. Experimental results in [15] demonstrate that these models, while idealized, provide throughput estimates that match remarkably well with measurements in actual IEEE 802.11 systems.

Despite their asynchronous and distributed nature, CSMA-like algorithms have been shown to offer the capability of achieving the full capacity region and thus match the optimal throughput performance of centralized scheduling algorithms operating in slotted time [13], [14], [16]. More specifically, any throughput vector in the interior of the convex hull associated with the independent sets in the underlying interference graph can be achieved through suitable back-off rates and/or transmission lengths. Based on this observation, various clever algorithms have been developed for finding the back-off rates that yield a particular target throughput vector or that optimize a certain concave throughput utility function in scenarios with saturated buffers [13], [14], [18]. In the same spirit, several powerful algorithms have been devised for adapting the transmission lengths based on backlog information, and been shown to guarantee maximum stability [12], [20].

Roughly speaking, the maximum-stability guarantees were established under the condition that the activity factors of the various nodes behave as logarithmic functions of the backlogs. Unfortunately, such activity factors can induce excessive backlogs and delays, which has triggered a strong interest in developing approaches for improving the delay performance [11], [17], [19], [21]. Motivated by this issue, Ghaderi & Srikant [9] recently showed that it is in fact sufficient for the *logarithms* of the activity factors to behave as logarithmic functions of the backlogs, divided by an arbitrarily slowly increasing, unbounded function. These results suggest that activity functions can be used that are essentially linear for all practical values of the backlogs in order to reduce the delays while preserving maximum-stability guarantees.

In the present paper we will establish a lower bound for the mean delay in full interference graphs for concave activity

functions and a corresponding upper bound for convex activity functions. Simulation experiments indicate that the lower and upper bounds are in fact remarkably tight. These results provide detailed insight into how the mean delays depend on the activity functions, and explicitly reveal how aggressive/persistent access schemes improve the delay performance. In particular, the mean delay in a heavy-traffic regime is exponentially larger for logarithmic activity functions than for linear ones. Since the bounds are difficult to extend to general interference graphs, we provide further arguments which suggest that qualitatively similar observations apply for a broader set of topologies. As a side-result, we show that, counter to intuition, additional interference may improve the delay performance in certain cases!

As it turns out, maximum stability is ensured for arbitrarily aggressive/persistent access schemes in the full interference graph as well as several further scenarios. However, the existing throughput optimality results for general topologies require the activity functions to grow relatively slowly. Specifically, the results in [9] allow the activity functions to be essentially linear for all practical values of the backlogs, but still require them to grow slower than any positive power of the backlogs asymptotically. This raises the issue how fast the activity functions are allowed to grow, depending on the topology, while retaining throughput optimality. In order to examine that issue, we will investigate fluid limits where the system dynamics are scaled in space and time. As we will show, several distinct types of fluid limits can arise, exhibiting various degrees of randomness, depending on the structure of the network, in conjunction with the form of the activity functions.

The remainder of the paper is organized as follows. In Section II we present a detailed model description. We analyze delay issues in Section III and provide lower and upper bounds for the mean delay in a full interference graph for concave and convex activity functions, respectively. In Section IV we investigate fluid limits in order to gain qualitative insight and examine stability issues. The analysis identifies a strong trichotomy, as governed by the mixing properties of the system, which is corroborated through simulation experiments. In Section V we make some concluding remarks and identify various topics for future research.

II. MODEL DESCRIPTION

Network, interference, and traffic models. We consider a network of several nodes sharing a wireless medium according to random multi-access protocols. The network is described by an undirected graph (V, E) where the set of vertices $V = \{1, \dots, M\}$ represents the various nodes of the network and the set of edges $E \subset V \times V$ indicates which pairs of nodes interfere. Nodes that are neighbors in the interference graph are prevented from simultaneous activity, and the independent sets correspond to the feasible joint activity states. A node is said to be blocked whenever the node itself or any of its neighbors is active, and unblocked otherwise. Define $\Omega \subset \{0, 1\}^M$ as the set of all feasible joint activity states of the network, i.e., the

incidence vectors of the independent sets of the interference graph. Let $\sigma(t) \in \Omega$ represent the activity state of the network at time t , with $\sigma_i(t)$ indicating whether node i is active at time t or not. Packets arrive at node i as a Poisson process of rate λ_i . The packet transmission times of node i are independent and exponentially distributed with mean $1/\mu_i$. Denote by $\rho_i = \lambda_i/\mu_i$ the traffic intensity of node i .

Backlog-based CSMA protocols. We analyze the following general class of backlog-based random multi-access protocols. Denote by $L_i(t)$ the number of packets at node i at time t . When inactive, node i may start transmitting at the instants of a time-inhomogeneous Poisson process of intensity $f_i(L_i(t))$ at time t , where $f_i : [0, \infty) \mapsto [0, \infty)$ and $f_i(0) = 0$. It actually starts a transmission if it is unblocked. $f_i(\cdot)$ is referred to as the *activation function* of node i . When active, at the end of a packet transmission, say at time t^* , node i releases the medium with probability $p_i(L_i(t^*))$ or starts a new packet transmission with probability $1 - p_i(L_i(t^*))$, where $p_i : [0, \infty) \mapsto [0, 1]$ and $p_i(1) = 1$. In other words, node i releases the medium at the instants of a time-inhomogeneous Poisson process of intensity $g_i(t) = p_i(L_i(t))\mu_i$. $g_i(\cdot)$ is referred to as the *de-activation function*. We define the *activity function* $h_i(\cdot)$ of node i as $h_i(L_i) = f_i(L_i)/g_i(L_i)$ and $h_i(0) = 0$.

Network dynamics. Under any of the aforementioned queue-based CSMA protocols, $(L(t), \sigma(t), t \geq 0)$ with $L(t) = (L_1(t), \dots, L_M(t))$ is a continuous-time Markov process. We are interested in quantifying the mean delays, depending on the functions $f_i(\cdot)$ and $g_i(\cdot)$. We are also interested in deriving conditions on $\rho = (\rho_1, \dots, \rho_M)$ guaranteeing the ergodicity of this Markov process. It is well known [22] that ergodicity can be achieved only if the vector of traffic intensities ρ lies in Γ , defined as the set of ρ such that there exists η in $\text{conv}(\Omega)$ with $\rho < \eta$ component-wise. We therefore assume that ρ lies in Γ .

III. DELAY ISSUES

In this section we study the (mean) number of packets in the system for several access schemes and topologies. Note that the mean packet delay follows directly from the mean number of packets by Little's law. In fact, the delay distribution also follows from the distribution of the number of packets by virtue of the distributional form of Little's law. We start with an exact analysis for a full interference graph. Next we will provide a heuristic analysis for general topologies. Because of page constraints, we limit the attention to the case $g_i(\cdot) \equiv \mu$ for all $i = 1, \dots, M$, but similar results hold for other choices.

A. Full interference graph

Denote by $\rho = \sum_{i=1}^M \rho_i$ the total traffic intensity and by $\lambda = \sum_{i=1}^M \lambda_i$ the total arrival intensity. As at most one node can be active at a time in the full interference graph it follows that the necessary condition for ergodicity in Section II is $\rho < 1$. Also, it is easily verified that the system is stable for any $\rho < 1$ as long as each of the functions $f_i(\cdot)$ is unbounded. Before looking at general activation functions we first present

an exact analysis of the distribution of the number of packets for linear activation functions.

1) *Linear activation functions:* For linear activation functions $f_i(n) = \nu n$, we no longer need to distinguish between nodes in order to describe the evolution in time of the total number of packets in the system. In fact, the total number of packets in the system gives rise to a “one-and-a-half”-dimensional Markov process whose stationary distribution can be found analytically, as shown next.

Consider the continuous-time Markov process with state space $\{0, 1, 2, \dots\} \times \{0, 1\}$, where the first component represents the total number packets in the system, $L = \sum_{i=1}^M L_i$, and the second component indicates whether one of the nodes is active (state 1) or not (state 0). Transitions from $(n, 0)$ to $(n + 1, 0)$ occur at rate ρ , transitions from $(n, 0)$ to $(n, 1)$ occur at rate νn , transition from $(n, 1)$ to $(n + 1, 1)$ occur at rate ρ , and transitions from $(n + 1, 1)$ to $(n, 0)$ occur at rate 1. With $\pi(n, k)$ the stationary probability that the Markov process resides in states (n, k) , we obtain the balance equations

$$\begin{aligned} \lambda\pi(0, 0) &= \mu\pi(1, 1), \\ (\lambda + \mu)\pi(1, 1) &= \nu\pi(1, 0), \\ (\lambda + \nu n)\pi(n, 0) &= \lambda\pi(n - 1, 0) + \mu\pi(n + 1, 1), \quad n \geq 1, \\ (\lambda + \mu)\pi(n, 1) &= \lambda\pi(n - 1, 1) + \nu n\pi(n, 0), \quad n \geq 2. \end{aligned}$$

Introducing the generating functions $G_0(z) = \sum_{n=0}^{\infty} \pi(n, 0)z^n$ and $G_1(z) = \sum_{n=1}^{\infty} \pi(n, 1)z^n$, the balance equations lead to

$$\lambda(1 - z)G_0(z) + \nu z G_0'(z) = \frac{\mu}{z} G_1(z), \quad (1)$$

$$(\mu + \lambda - \lambda z)G_1(z) = \nu z G_0'(z), \quad (2)$$

yielding

$$G_1(z) = \frac{\rho z}{1 - \rho z} G_0(z) \quad (3)$$

(which corresponds to the Fuhrmann-Cooper decomposition property [7] for the distribution of the total number of packets in the system) and

$$G_0(z) = \frac{\nu z(\mu - \lambda z)}{\lambda z(\lambda + \mu - \lambda z)} G_0'(z). \quad (4)$$

Since $G_0(1) = 1 - \rho$, we obtain from (3) that $G_1'(1) = \rho(1 + G_0'(1))/(1 - \rho)$. Also $G_0'(1) = \frac{\mu}{\nu} G_1(1) = \frac{\lambda}{\nu}$. So the mean number of packets in the system is

$$\mathbb{E}\{L\} = \sum_{n=1}^{\infty} n(\pi(n, 0) + \pi(n, 1)) = G_0'(1) + G_1'(1) \quad (5)$$

$$= \frac{\rho + G_0'(1)}{1 - \rho} = \frac{\rho + \frac{\lambda}{\nu}}{1 - \rho} = \frac{\lambda(\mu + \nu)}{\nu(\mu - \lambda)}. \quad (6)$$

In fact, we can obtain the entire stationary distribution by solving the differential equation for G_0 in (4). The solution reads (using $G_0(1) = 1 - \rho$)

$$G_0(z) = (1 - \rho) \left(\frac{1 - \rho}{1 - \rho z} \right)^{\lambda/\nu} e^{(z-1)\lambda/\nu}, \quad (7)$$

which can be interpreted as a convolution of a negative binomial distribution and a Poisson distribution. This observation leads to

$$\pi(n, 0) = \frac{e^{-\lambda/\nu} (1 - \rho)^{\lambda/\nu + 1}}{\Gamma(\frac{\lambda}{\nu})} \sum_{k=0}^n \frac{\Gamma(k + \frac{\lambda}{\nu}) \rho^k (\frac{\lambda}{\nu})^{n-k}}{\Gamma(k + 1) \Gamma(n - k + 1)}.$$

2) *General activation functions:* For general activation functions $f_i(\cdot)$, Equation (3) continues to hold, but Equation (4) no longer does, making an exact analysis less tractable. Hence we now proceed to derive bounds for $\mathbb{E}\{L\}$.

Denote by $L_{i,I}$ the number of packets at node i at an arbitrary epoch during a non-serving interval per unit of time. Observing that the mean number of non-serving intervals equals the mean number of packets served per unit of time, we obtain

$$\sum_{i=1}^M \mathbb{E}\{f_i(L_{i,I})\} = \frac{\lambda}{1 - \rho}, \quad (8)$$

which may also be deduced directly from the balance equations for the above-described Markov process. In fact, it may be shown that

$$\mathbb{E}\{f_i(L_{i,I})\} = \frac{\lambda_i}{1 - \rho}, \quad i = 1, \dots, M. \quad (9)$$

Theorem 1: Let $f_i(\cdot) \equiv f(\cdot)$ with $f : [0, \infty) \mapsto [0, \infty)$ be a strictly increasing, unbounded and concave function. Under the above assumptions,

$$\mathbb{E}\{L\} \geq \frac{\rho}{1 - \rho} + M f^{-1} \left(\frac{1}{M} \frac{\lambda}{1 - \rho} \right). \quad (10)$$

Similarly, if $f_i(\cdot) \equiv f(\cdot)$ is a strictly increasing, continuous and convex function,

$$\mathbb{E}\{L\} \leq \frac{\rho}{1 - \rho} + M f^{-1} \left(\frac{1}{M} \frac{\lambda}{1 - \rho} \right). \quad (11)$$

Proof: First note that the system is stable as $\rho < 1$ and $f(\cdot)$ is unbounded. If $f(\cdot)$ is concave it follows by Jensen's inequality that

$$\sum_{i=1}^M \mathbb{E}\{f(L_{i,I})\} \leq M f \left(\frac{1}{M} \sum_{i=1}^M \mathbb{E}\{L_{i,I}\} \right).$$

Since $f(\cdot)$ is increasing, we get using (8),

$$\sum_{i=1}^M \mathbb{E}\{L_{i,I}\} \geq M f^{-1} \left(\frac{1}{M} \frac{\lambda}{1 - \rho} \right).$$

The Fuhrmann-Cooper decomposition property [7] (applied to the total number of packets in the system) implies

$$\mathbb{E}\{L\} = \frac{\rho}{1 - \rho} + \sum_{i=1}^M \mathbb{E}\{L_{i,I}\},$$

yielding (10). Equation (11) follows by symmetry. ■

Note that when $f(\cdot)$ is linear, all inequalities in the proof of Theorem 1 are in fact equalities (or note that this function is both concave and convex), so that we recover (6). Further note that for concave functions with $f(1) \leq \nu$, the mean

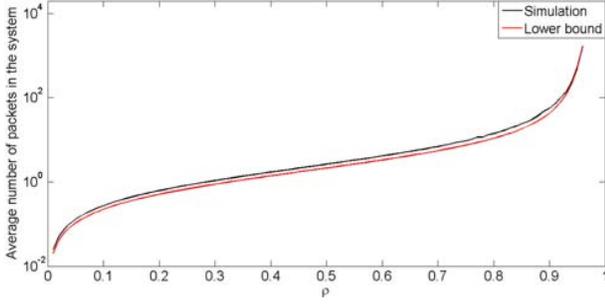


Fig. 1. Average number of packets in the system for $f(x) = \log(x + 1)$.

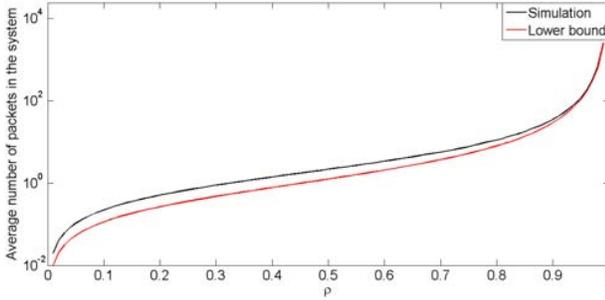


Fig. 2. Average number of packets in the system for $f(x) = \sqrt{x}$.

number of packets in the system is always larger than the mean number of packets in the system with the linear activation function $f(n) \equiv \nu n$. Conversely, for convex functions with $f(1) \geq \nu$, the mean number of packets in the systems is always less than the mean number of packets in the system with $f(n) \equiv \nu n$. More specifically, in heavy traffic as $\rho \uparrow 1$ the mean number of packets, and hence the mean delay, grows like $M \exp(\frac{\rho}{M(1-\rho)})$ for a logarithmic activation function and like $\frac{\rho}{1-\rho}$ for an exponential activation function. We thus see that more aggressive activation functions improve the delay performance.

In order to investigate how tight the bounds derived in Theorem 1 are for non-linear activation functions, we performed several numerical experiments. Figures 1, 2 and 3 show the average number of packets in a simulation for a 4-node full interference graph for activation functions $f(x) = \log(x + 1)$,

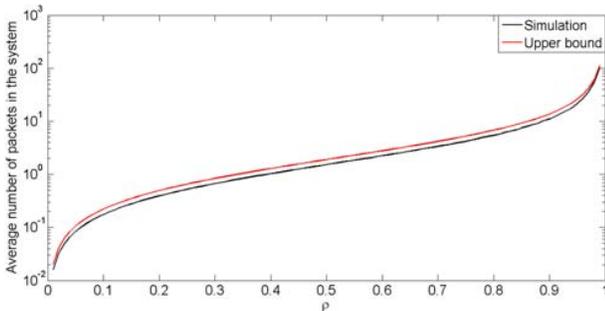


Fig. 3. Average number of packets in the system for $f(x) = e^x - 1$.

$f(x) = \sqrt{x}$ and $f(x) = e^x - 1$, respectively. In each figure we also plotted the bound derived in Theorem 1 and we see that the numbers from the simulation are close to the bound, i.e. the bound is rather tight. Further, for values of ρ close to 1 the bound is relatively tighter. Note that we used a log-linear scale.

B. General topologies

For general interference graphs, assuming the system to be stable, we may write $\rho_i = \phi_i \pi_i$, with $\phi_i = \mathbb{E}\{f_i(L_i) | E_i\} / \mu$, $\pi_i = \mathbb{P}\{E_i\}$, and E_i denoting the event that node i and all its neighbors are inactive. In the case of a full interference graph, we trivially have $\pi_i = 1 - \rho$, but unfortunately such a simple relationship does not hold in general. Hence we make the approximation that the activity process behaves as if each of the nodes operates according to a constant activation function that corresponds to its mean queue size. This provides approximations for both π_i and ϕ_i . First of all, $\pi_i \approx \hat{\pi}_i(\phi)$, with $\phi = (\phi_1, \dots, \phi_M)$, where $\hat{\pi}_i(\phi)$ pertains to a scenario with constant activity factors ϕ_1, \dots, ϕ_M . Now observe that replacing π_i by $\hat{\pi}_i(\phi)$ results in $\rho_i = \phi_i \hat{\pi}_i(\phi)$, which is tantamount to (ρ_1, \dots, ρ_M) being the long-term throughputs in case of constant activity factors ϕ_1, \dots, ϕ_M . As proved by Jiang & Walrand [14] and Van de Ven *et al.* [23], this uniquely determines each of the ϕ_i 's as some function $\phi_i(\rho)$ of $\rho = (\rho_1, \dots, \rho_M)$, rendering $\phi_i \approx \phi_i(\rho)$. In addition, $\phi_i \approx f_i(\mathbb{E}\{L_i\}) / \mu$. Combining the above two approximations, we obtain

$$\mathbb{E}\{L_i\} \approx f_i^{-1}(\mu \phi_i(\rho)).$$

Note that $\phi_i(\rho)$ does not depend on the activation functions $f_i(\cdot)$ at all. Thus the above approximation suggests that when $f_i(q)$ increases with q in a more aggressive manner, so that $f_i^{-1}(r)$ is smaller for a given value of r , the mean number of packets at node i will be smaller.

The above approximation seems reasonable when the activation functions are relatively flat, and the numbers of packets at the various nodes do not fluctuate too much and are concentrated around their mean values. In order to investigate the accuracy, we first revisit the case of a full interference graph and compare the approximations with the exact results of Section III-A. Next we will examine the approximations for two other networks.

1) *Full interference graph:* In this case, we have $\frac{\phi_i(\rho)}{1 + \sum_{j=1}^M \phi_j(\rho)} = \rho_i$, so that $\phi_i(\rho) = \frac{\rho_i}{1-\rho}$, yielding

$$\mathbb{E}\{L\} \approx \sum_{i=1}^M f_i^{-1}\left(\frac{\lambda_i}{1-\rho}\right). \quad (12)$$

With $f_i(n) = \nu n$ for all $i = 1, \dots, M$, the above approximation yields $\mathbb{E}\{L\} \approx \frac{\lambda/\nu}{1-\rho}$, which differs from the exact result (6). Upon closer inspection and comparison with Equation (9), we note that the relationship $\phi_i \approx \phi_i(\rho)$ is in fact exact in this case, and that the discrepancy is entirely due to the approximation $\phi_i \approx f_i(\mathbb{E}\{L_i\})$. Invoking the relationship $\sum_{i=1}^M \mathbb{E}\{L_i\} = \frac{M}{1-\rho} + \sum_{i=1}^M \mathbb{E}\{L_i | E_i\}$ based

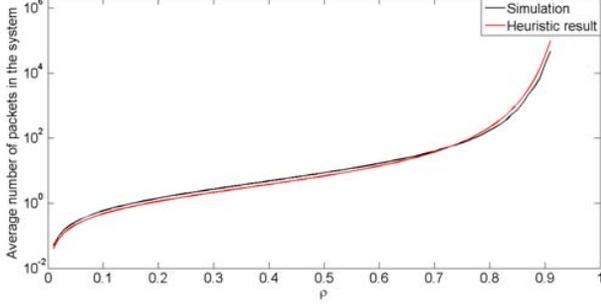


Fig. 4. Average number of packets in the 4-node ring network with $h(x) = \log(x+1)$.

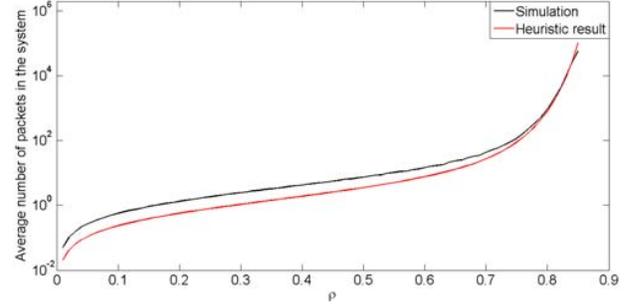


Fig. 5. Average number of packets in the 4-node linear network with $h(x) = \log(x+1)$.

on the Fuhrmann-Cooper decomposition instead, recovers the exact result.

We further observe that for less aggressive activation functions (12) performs better as expected. In particular, the ratio of (12) and the lower bound in Theorem 1 tends to 1 as $\rho \uparrow 1$ for strictly concave activation functions.

2) *4-node ring network*: Consider a network of 4 nodes that are positioned in a ring, so that node 1 interferes with nodes 2 and 4, and so on. Assume that $\rho_1 = \rho_3 = \rho_{13}$ and $\rho_2 = \rho_4 = \rho_{24}$, then $\phi_i(\rho) \approx \frac{\sqrt{\rho_i(\sqrt{\rho_{13}} + \sqrt{\rho_{24}})}}{1-\rho}$, and we obtain

$$\mathbb{E}\{L\} \approx \sum_{i=1}^4 f_i^{-1}(\mu\phi_i(\rho)), \quad (13)$$

which for $\rho_1 = \rho_2 = \frac{\rho}{2}$ simplifies to

$$\mathbb{E}\{L\} \approx \sum_{i=1}^4 f_i^{-1}\left(\frac{\lambda}{1-\rho}\right).$$

In Figure 4 we check the accuracy of (13) with simulation results for $f_i(x) \equiv f(x) = \log(x+1)$. We observe that the approximation formula (13) is quite accurate for most values of ρ .

3) *4-node linear network*: Consider a network of 4 nodes that are positioned on a line, so that node 2 interferes with nodes 1 and 3, while node 3 further interferes with node 4. Assume $\rho_i \equiv \frac{\rho}{2}$, then $\phi_1(\rho) = \phi_4(\rho) = \phi_{14}(\rho) = \frac{\lambda}{2(1-\rho)}$, and $\phi_2(\rho) = \phi_3(\rho) = \phi_{23}(\rho) = \frac{(2-\rho)\lambda}{4(1-\rho)^2}$, and we obtain

$$\mathbb{E}\{L\} \approx f_1^{-1}(\phi_{14}) + f_2^{-1}(\phi_{23}) + f_3^{-1}(\phi_{23}) + f_4^{-1}(\phi_{14}). \quad (14)$$

In Figure 5 we compare (14) with simulation results for $f_i(x) \equiv f(x) = \log(x+1)$. We see that the approximation (14) is accurate for higher values of ρ .

Remarkably we observe that the 4-node linear network performs worse than the 4-node ring network in terms of mean delay in heavy traffic, even though the 4-node linear network is just the 4-node ring network in which nodes 1 and 4 do not interfere. Thus adding an interference constraint improves the delay performance in this case! This may be explained by noting that removing the interference between nodes 1 and 4 requires nodes 2 and 3 to have far higher activation rates and hence far larger numbers of packets in order to claim a fair share of the throughput.

IV. FLUID LIMITS

The results in the previous section demonstrated that more aggressive activation functions improve the delay performance in various specific scenarios. For arbitrary networks, however, the existing maximum stability results involve slowly varying activation functions, which raises the issue how aggressively the activation functions are allowed to grow, depending on the topology, while retaining throughput optimality. In order to address that issue, we will explore in this section the dynamics of the Markov process $Z(t) = (L(t), \sigma(t))$ using fluid limits.

Fluid limits may be interpreted as first-order approximations of the Markov process, and provide valuable qualitative insights and a powerful approach for establishing ergodicity properties. Fluid limits are obtained by scaling the relevant stochastic processes in both space and time. More specifically, we consider a sequence of processes $Z^N(t)$ indexed by a sequence of positive integers N , each governed by similar statistical laws as the process $Z(t)$, where the initial states satisfy $\sum_{i=1}^M L_i(0) = N$ and $L_i^N(0)/N \rightarrow q_i \geq 0$ as $N \rightarrow \infty$. The process $\bar{Z}^N(t) = (\frac{1}{N}L^N(Nt), \sigma^N(Nt))$ is called the fluid-scaled version of the process $Z^N(t)$. Note that the activity process is scaled in time as well, but not in space.

For conciseness, define $\xi^N(t) = \frac{1}{N}L^N(Nt)$. Then $\xi^N(t)$ can be viewed as a random element of $D(\mathbb{R}^+, \mathbb{R}_+^M)$, the set of cadlag functions with values in \mathbb{R}_+^M . There is a metric on $D(\mathbb{R}^+, \mathbb{R}_+^M)$ such that the latter set is a complete and separable space. Any random element of $D(\mathbb{R}^+, \mathbb{R}_+^M)$ which is a limit point of the sequence $\{\xi^N(t)\}_{N \geq 1}$, i.e., whose law is an accumulation point of the laws of $\{\xi^N(t)\}_{N \geq 1}$, is called a fluid limit.

A. Fluid limit trichotomy

Unlike in most queueing systems where fluid limits follow deterministic trajectories described by a set of differential equations, our system may exhibit fluid limits that are stochastic processes. To facilitate the discussion here, we assume that all queues are initially backlogged, i.e., $q_i > 0$ for all i .

The process $Z^N(t)$ has two interacting components, $L^N(t)$ and $\sigma^N(t)$, respectively. On the one hand, the evolution of $L^N(t)$ depends on the rate at which queues are served, and in turn depends on $\sigma^N(t)$. On the other hand, when the queues

$L^N(t)$ are fixed, the process $\sigma^N(t)$ is a reversible Markov process on the set of possible activation states whose transitions are functions of $L^N(t)$. As it turns out, we encounter different types of fluid limits depending on the *mixing* properties of the activity process $\sigma^N(t)$ as described next.

1) *Fast mixing - Deterministic fluid limits:* When all the transitions between the various activity states occur on a time scale faster than N , we obtain classical deterministic fluid limits. In such cases, $\sigma^N(t)$ evolves much faster than $L^N(t)$ as N grows large, and to obtain the rate at which queues are served in the fluid regime, the activity process $\sigma^N(t)$ is averaged. More precisely, we encounter fluid limits of the form

$$\xi_i'(t) = \lambda_i - \mu_i \phi_i(\xi(t)),$$

where $\phi_i(\xi)$ is the limiting fraction of time that node i is active in a saturated version of the system with node-dependent but time-invariant activation factors $h_i(N\xi_i)$ as $N \rightarrow \infty$. Now suppose that $h_i(\cdot) \equiv h(\cdot)$ for all $i = 1, \dots, M$. Also assume that $\lim_{N \rightarrow \infty} h(aN)/h(bN)$ exists for any $a, b > 0$, and that $h(aN) = o(h(bN)h(cN))$ for any $a, b, c > 0$ as $N \rightarrow \infty$. The latter assumption ensures that at the limit when $N \rightarrow \infty$, the schedules belong to the collection of dominant activity states, which corresponds to the collection of maximum-size independent sets defined by

$$\Omega^* := \left\{ \sigma \in \Omega : \sum_{i=1}^M \sigma_i = \max_{u \in \Omega} \sum_{i=1}^M u_i \right\}.$$

For any $\sigma \in \Omega^*$, define

$$\psi(\sigma; \xi) = \lim_{N \rightarrow \infty} \prod_{i=1}^M \left(\frac{h(N\xi_i)}{h(N)} \right)^{\sigma_i}.$$

For example, when $h(x) = x$, we obtain $\psi(\sigma; \xi) = \prod_{i=1}^M \xi_i^{\sigma_i}$, and when $h(x) = \log(x)$, we obtain $\psi(\sigma; \xi) = 1$.

Then

$$\phi_i(\xi) = \sum_{\sigma \in \Omega^*} \sigma_i \hat{\pi}(\sigma; \xi),$$

where $\hat{\pi}(\sigma; \xi) = \frac{1}{Z(\xi)} \psi(\sigma; \xi)$, with $\sigma \in \Omega^*$, and $Z(\xi) = \sum_{\sigma \in \Omega^*} \psi(\sigma; \xi)$.

2) *Slow mixing - Inhomogeneous Poisson fluid limits:*

When all the transitions between the various activation states occur on a time scale of order N , we will observe these in the fluid regime. More specifically, the fluid limit is differentiable almost everywhere in that case, except in some random set of measure zero which is produced by a time-inhomogeneous Poisson process as we will argue in Section IV-B.

3) *Torpid mixing - Pseudo-deterministic fluid limits:* Finally, when all the transitions between the various activation states occur on a time scale slower than N , the activation state seems to be frozen in the fluid regime. We obtain, until at least one queue empties, a fluid limit with deterministic (in fact linear) trajectories, but which trajectory is followed depends on the initial state and might be random. Transitions in the activation states can occur when a queue empties, and

may be random as well, yielding similar qualitative behavior as in [6].

It should be noted that one may construct examples of networks and activity functions such that the fluid limits correspond to a combination of fast, slow and torpid mixing. The above-described trichotomy will be discussed in greater depth below for the example of K -partite complete interference graphs. The strong qualitative difference in fluid limits raises the question how to determine whether the transitions between the dominant states occur on a time scale of the order N , faster, or slower. As it turns out, this is governed by the structure and size of the interference graph, in conjunction with the behavior of the activity factors as function of the backlog. Informally speaking, the more stable the maximum-size independent sets associated with the dominant states, and the higher the activity factors, the slower the transitions, see [10] for related results. A complete characterization of the transition rate for arbitrary graphs and arbitrary activity functions is outside the scope of the present paper.

B. K -partite complete interference graph

To illustrate the fluid limit trichotomy described above, we focus here on networks with a K -partite complete interference graph. Specifically, the set of nodes can be partitioned into K disjoint subsets V_1, \dots, V_K , of respective cardinalities M_1, \dots, M_K . Nodes within a given subset V_i do not interfere with each other, but do interfere with all other nodes. In other words, there is an edge $\{v, w\}$ between two nodes $v \in V_i$ and $w \in V_j$ if and only if $i \neq j$, i.e., $E = \bigcup_{i \neq j} V_i \times V_j$.

For K -partite networks, there are K maximal schedules. A maximal schedule consists in having all nodes from a given V_i simultaneously active. We assume here that $f_i(\cdot) \equiv f(\cdot)$ and $g_i(\cdot) \equiv g(\cdot)$. We further assume that $f(\cdot)$ and $g(\cdot)$ are chosen such that the system spends almost all of the time in maximal schedules when N tends to infinity.

We heuristically derive fluid limits, assuming that $q_i > 0$ for all $i \in V$. When the system has no active node, the next active maximal schedule is V_j with probability proportional to $\sum_{i \in V_j} f(Nq_i)$. Indeed, the latter schedule is determined by the first node grabbing the channel. To derive fluid limits, we need to quantify the time τ_j^N it takes when starting from maximal schedule V_j for all nodes in V_j to release the channel. In the following discussion, we consider time intervals whose durations are small enough so that the queue lengths do not evolve significantly, and hence can be considered as constant.

Define by T_j^N the time spent in schedule V_j between two successive instants where the system has no active node. We know from the theory of Markov processes [1] that

$$\mathbb{E}[T_j^N] = \frac{\prod_{i \in V_j} f(Nq_i)/g(Nq_i)}{\sum_i f(Nq_i)}.$$

We also have, for large N ,

$$\mathbb{E}[T_j^N] \approx \frac{\sum_{i \in V_j} f(Nq_i)}{\sum_i f(Nq_i)} \times \mathbb{E}[\tau_j^N],$$

and hence

$$\mathbb{E}[\tau_j^N] \approx \frac{\prod_{i \in V_j} f(Nq_i)/g(Nq_i)}{\sum_{i \in V_j} f(Nq_i)}.$$

In fact, one can even state that the distribution of τ_j^N is approximately exponentially distributed as N grows large. This is a classical property for exit times of Markov processes with asymptotically small transition rates. But it can be directly checked in our specific example as sketched below. We can prove it by induction on M_j . For $M_j = 1$, the result is immediate. Assume that the result holds for $M_j = H - 1 \geq 1$, and let us prove it for $M_j = H$. Label the H nodes in V_j by $1, \dots, H$. By induction, the time for the $H - 1$ first nodes to release the channel is approximately exponentially distributed, say with mean $1/r$. Once these nodes are inactive, the time it takes for one of these nodes to be active again is exponentially distributed with mean $1/s$, with $s = \sum_{i=1}^{H-1} f(Nq_i)$. This happens before the H -th node releases the channel with probability $s/(s + g(Nq_H))$. Similarly after releasing the channel, the H -th node becomes active again before the $H - 1$ other nodes release the channel with probability $f(Nq_H)/(r + f(Nq_H))$. It can be easily verified that when N grows large, $1/r \gg 1/s$ and $1/r \gg 1/f(Nq_H)$. Hence the transitions from a state where the H nodes are active to a state where no node is active occur at the instants of a Poisson process, which is the superposition of two Poisson processes of respective rates $r \times g(Nq_H)/(s + g(Nq_H))$ and $g(Nq_H) \times r/(r + f(Nq_H))$. Thus τ_j^N is approximately exponentially distributed.

Having characterized τ_j^N , we can now sketch how to derive the fluid limits

(i) If $\lim_{N \rightarrow \infty} \mathbb{E}[\tau_j^N]/N = 0$ for all j , then the time to switch maximal schedules is negligible compared to the time required for queue lengths to change. We have a separation of time scales, and the fluid limits are obtained by assuming that schedule V_j is used a fraction of time equal to the corresponding stationary distribution, as discussed in Section IV-A1.

(ii) Assume without loss of generality that $M_1 \geq \dots \geq M_K$. Assume that for $j = 1, \dots, j_0$, $\lim_{N \rightarrow \infty} \mathbb{E}[\tau_j^N]/N = \alpha_j > 0$, and that for all $j > j_0$, $\lim_{N \rightarrow \infty} \mathbb{E}[\tau_j^N]/N = 0$. In such cases, in the fluid limits, the set of active schedules are V_1, \dots, V_{j_0} , and one switches from one schedule to another at random epochs. When schedule V_j , $j \leq j_0$, is used, we have $\xi'_i(t) = \lambda_i - \mu_i$, for all $i \in V_j$, and $\xi'_i(t) = \lambda_i$, for all $i \notin V_j$. The switching times between the various schedules are driven by time-inhomogeneous Poisson processes. That is, given backlogs q_i , the switching time in the fluid regime from schedule V_j to schedule $V_{j'}$ is exponentially distributed with mean

$$\lim_{N \rightarrow \infty} \frac{1}{N} \frac{\prod_{i \in V_j} f(Nq_i)/g(Nq_i)}{\sum_{i \in V_j} f(Nq_i)} \frac{\sum_{i \in V_{j'}} f(Nq_i)}{\sum_i f(Nq_i)}.$$

For example, in a 2-partite interference graph with $M_1 = M_2 = 2$, and $f(x) = x$, $g(x) = 1$, the average time to switch

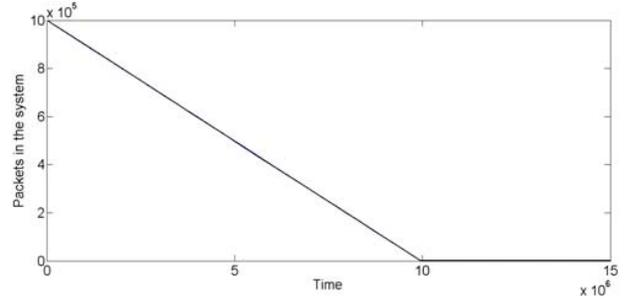


Fig. 6. Deterministic fluid limit for the bipartite interference graph with $M_1 = M_2 = 1$, $f(x) = x$ and $g(x) = 1$: fast mixing.

from schedule $V_1 = \{1, 2\}$ to schedule $V_2 = \{3, 4\}$ is

$$\frac{q_1 q_2}{q_1 + q_2} \frac{q_3 + q_4}{q_1 + q_2 + q_3 + q_4}.$$

This case corresponds to Section IV-A2.

(iii) Assume that $\lim_{N \rightarrow \infty} \mathbb{E}[\tau_j^N]/N = \infty$, and that initially schedule V_j is used. Then in the fluid regime we have, until some of the queues in V_j empty, $\xi'_i(t) = \lambda_i - \mu_i$ for all $i \in V_j$, and $\xi'_i(t) = \lambda_i$ for all $i \notin V_j$. Other behavior can occur for different initial schedules, e.g. if $\lim_{N \rightarrow \infty} \mathbb{E}[\tau_j^N]/N = \infty$ for all j and initially no node is active, then, with probability $\prod_{i \in V_{j'}} f(Nq_i) / \prod_i f(Nq_i)$ for all j' , $\xi'_i(t) = \lambda_i - \mu_i$ for all $i \in V_{j'}$, and $\xi'_i(t) = \lambda_i$ for all $i \notin V_{j'}$. This case corresponds to Section IV-A3.

We illustrate the three above possible scenarios through numerical experiments. That is, to investigate the fluid limit, we examine the evolution of the number of packets over time in a network that initially has a lot of packets at each node. More precisely, we consider bipartite graphs with initially 10^6 packets at every node. Further, we set $\lambda_i = 0.4$ and $\mu_i = 1$ for all i and initially no node is made active. Setting $f_i(x) \equiv f(x) = x$ and $g_i(x) \equiv g(x) = 1$ the empirical fluid limit for the bipartite graph whose subsets all have cardinality one is given in Figure 6. For the case where all subsets have cardinality two the empirical fluid limit is given in Figure 7 and for the case where all subsets have cardinality three the empirical fluid limit is given in Figure 8. These scenarios give fast mixing, slow mixing and torpid mixing respectively. Thus we observe that for the same activation function the three scenarios can be obtained by changing the topology.

Alternatively, the three different scenarios can be obtained by changing the activation function, $f(x)$, for the same topology. Consider the bipartite graph whose subsets all have cardinality three. In Figure 8 we found torpid mixing for $f(x) = x$. Setting $f(x) = \sqrt{x}$ we obtain slow mixing as seen in Figure 9 and setting $f(x) = \log(x + 1)$ we obtain fast mixing as seen in Figure 10.

V. CONCLUSION

We have explored the congestion dynamics of wireless networks with backlog-based CSMA mechanisms. Lower and upper bounds as well as heavy-traffic approximations were

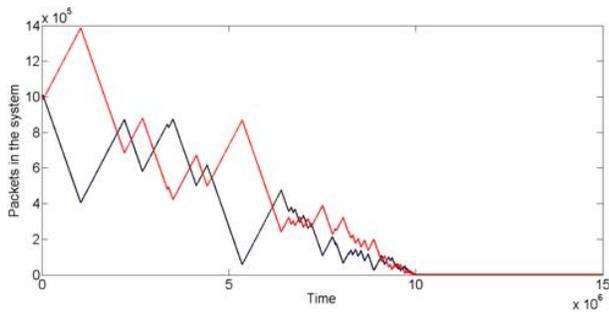


Fig. 7. Inhomogeneous Poisson fluid limit for the bipartite interference graph with $M_1 = M_2 = 2$, $f(x) = x$ and $g(x) = 1$: slow mixing.

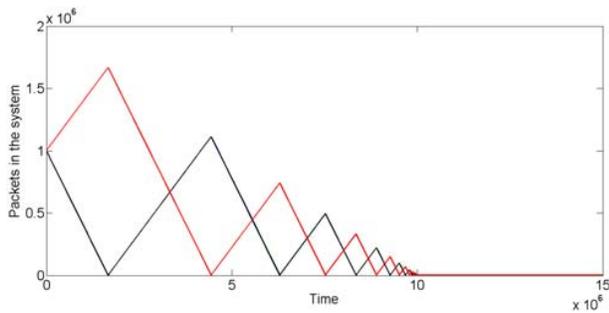


Fig. 8. Pseudo-deterministic fluid limit for the bipartite interference graph with $M_1 = M_2 = 3$, $f(x) = x$ and $g(x) = 1$: torpid mixing.

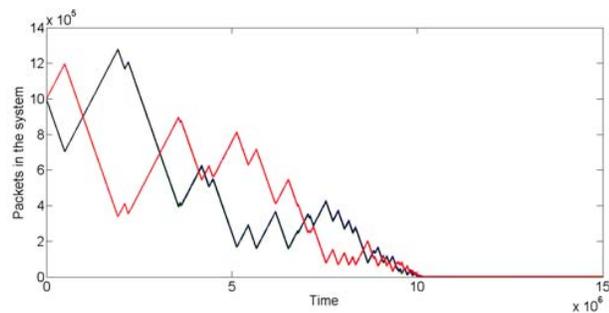


Fig. 9. Inhomogeneous Poisson fluid limit for the bipartite interference graph with $M_1 = M_2 = 3$, $f(x) = \sqrt{x}$ and $g(x) = 1$: slow mixing.

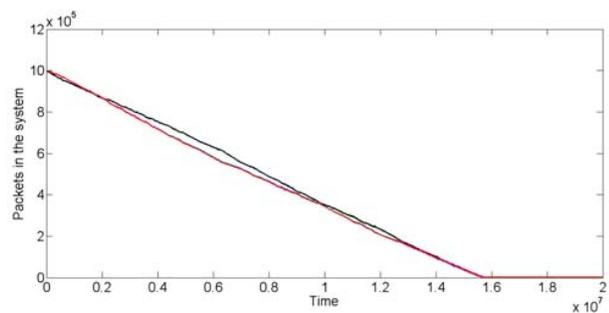


Fig. 10. Deterministic fluid limit for the bipartite interference graph with $M_1 = M_2 = 3$, $f(x) = \log(x + 1)$ and $g(x) = 1$: fast mixing.

obtained which provide detailed insight into how aggressive/persistent access schemes improve the delay performance. A key challenge for further research is to establish how fast the activity functions are allowed to grow, depending on the topology, while retaining maximum stability. As a first step in that direction, we have investigated fluid limits, and identified several distinct qualitative regimes that can arise, as governed by the mixing properties of the system.

REFERENCES

- [1] D. Aldous, J. Fill (2002). Reversible Markov Chains and Random Walks on Graphs. Available online at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- [2] G. Bianchi (2000). Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE J. Sel. Areas Commun.* **18** (3), 535–547.
- [3] R.R. Boorstyn, A. Kershbaum, B. Maglaris, V. Sahin (1987). Throughput analysis in multihop CSMA packet radio networks. *IEEE Trans. Commun.* **35**, 267–274.
- [4] M. Durvy, O. Dousse, P. Thiran (2007). Modeling the 802.11 protocol under different capture and sensing capabilities. In: *Proc. Infocom 2007*.
- [5] M. Durvy, P. Thiran (2006). A packing approach to compare slotted and non-slotted medium access control. In: *Proc. Infocom 2006*.
- [6] M. Feuillet, A. Proutiere, Ph. Robert (2010). Random capture algorithms: fluid limits and stability. In: *Proc. ITA Workshop*.
- [7] S.W. Fuhrmann, R.B. Cooper (1985). Stochastic decompositions in the $M/G/1$ queue with generalized vacations. *Opns. Res.* **33**, 1117–1129.
- [8] M. Garetto, T. Salonidis, E.W. Knightly (2008). Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks. *IEEE/ACM Trans. Netw.* **16** (4), 864–877.
- [9] J. Ghaderi, R. Srikant (2010). On the design of efficient CSMA algorithms for wireless networks. Preprint.
- [10] L. Jiang, M. Leconte, J. Ni, R. Srikant, J. Walrand (2010). Fast mixing of parallel Glauber dynamics and low-delay CSMA scheduling. Preprint.
- [11] L. Jiang, J. Ni, R. Srikant, J. Walrand (2010). Performance bounds of distributed CSMA scheduling. Preprint.
- [12] L. Jiang, D. Shah, J. Shin, J. Walrand (2010). Distributed random access algorithm: scheduling and congestion control. *IEEE Trans. Inform. Theory*, to appear.
- [13] L. Jiang, J. Walrand (2008). A distributed CSMA algorithm for throughput and utility maximization in wireless networks. In: *Proc. Allerton 2008*.
- [14] L. Jiang, J. Walrand (2010). A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Trans. Netw.* **18** (3), 960–972.
- [15] S.C. Liew, C.H. Kai, J. Leung, B. Wong (2010). Back-of-the-envelope computation of throughput distributions in CSMA wireless networks. *IEEE Trans. Mob. Comp.* **9** (9), 1319–1331.
- [16] J. Liu, Y. Yi, A. Proutière, M. Chiang, H.V. Poor (2008). Maximizing utility via random access without message passing. Tech. Rep. MSR-TR-2008-128, Microsoft Research.
- [17] M. Lotfinezhad, P. Marbach (2010). Throughput-optimal random access with order-optimal delay. Preprint.
- [18] P. Marbach, A. Eryilmaz (2008). A backlog-based CSMA mechanism to achieve fairness and throughput-optimality in multihop wireless networks. In: *Proc. Allerton 2008*.
- [19] J. Ni, B. Tan, R. Srikant (2010). Q-CSMA: queue-length based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks. In: *Proc. Infocom 2010 Mini-Conf*.
- [20] S. Rajagopalan, D. Shah, J. Shin (2009). Network adiabatic theorem: an efficient randomized protocol for content resolution. In: *Proc. ACM SIGMETRICS/Performance 2009*.
- [21] D. Shah, J. Shin (2010). Delay-optimal queue-based CSMA. In: *Proc. ACM SIGMETRICS 2010 Conf*.
- [22] L. Tassiulas, A. Ephremides (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Aut. Contr.* **37**, 1936–1948.
- [23] P.M. van de Ven, S.C. Borst, D. Denteneer, A.J.E.M. Janssen, J.S.H. van Leeuwen (2010). Equalizing throughputs in random-access networks. *Perf. Eval. Rev.* **38** (2), 39–41.
- [24] X. Wang, K. Kar (2005). Throughput modeling and fairness issues in CSMA/CA based ad-hoc networks. In: *Proc. Infocom 2005*.