On the stability of flow-aware CSMA

T. Bonald[a], M. Feuillet[b]

[a]*Telecom ParisTech, Paris, France*
[b]*INRIA, Rocquencourt, France*

Abstract

We consider a wireless network where each *flow* (instead of each link) runs its own CSMA (Carrier Sense Multiple Access) algorithm. Specifically, each flow attempts to access the radio channel after some random time and transmits a packet if the channel is sensed idle. We prove that, unlike the standard CSMA algorithm, this simple distributed access scheme is optimal in the sense that the network is stable for all traffic intensities in the capacity region of the network.

Keywords: Wireless network, conflict graph, CSMA, flow-level dynamics, stability, throughput performance.

1. Introduction

The CSMA (Carrier Sense Multiple Access) algorithm is one of the most common medium access schemes in today's networks, both wired (e.g. IEEE 802.3) and wireless (e.g. IEEE 802.11). However, this algorithm is known to be inherently unfair, as illustrated by the two scenarios of Fig. 1. The first scenario relates to the downstream vs. upstream bandwidth sharing for a single access point. In the presence of n active mobiles on the upstream, the access point competes with n nodes for accessing the channel, resulting in a downstream to upstream bandwidth ratio of $1/n$, independently on the number of active flows on the downstream. The second scenario illustrates the impact of interference on bandwidth sharing. The center access point cannot transmit if one of the edge access points is active and thus gets much less transmission opportunities. Moreover, the resulting bandwidth sharing is inefficient since the edge access points can access the channel alternately, preventing the center access point from sending its traffic. Thus the CSMA algorithm is not able to fully utilize network capacity, a statement that will be made more precise later in the paper.

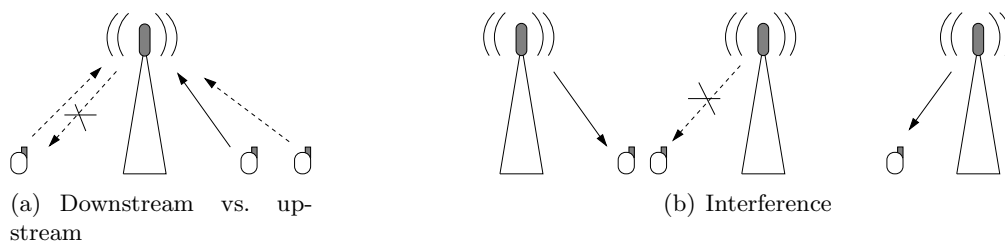(a) Downstream vs. upstream (b) Interference

Figure 1: Unfairness of standard CSMA.

We propose a slight modification of the standard CSMA algorithm that consists in running the algorithm for each *flow* instead of each transmitter. In this paper, we refer to a flow as

any file transfer from a source to a destination; it can typically be identified through the usual 5-uple: IP source and destination addresses, source and destination ports, protocol. For a single access point, each *flow* (either downstream or upstream) runs the CSMA algorithm and thus gets the same bandwidth share. The whole system can then be viewed as a unique, evenly shared wireless link. The focus of the present paper is rather on the second scenario where some links suffer from high interference. Specifically, we show that the flow-aware CSMA algorithm is optimal in the sense that it stabilizes the network whenever possible. In the example of Fig. 1, the center access point is likely to access the channel when it has a high number of active flows; at the end of the corresponding activity period, the edge access points can access the channel and will likely be *simultaneously* active, which is a necessary condition for fully utilizing network capacity.

The main result of the paper is to demonstrate that the flow-aware CSMA algorithm is optimal for *any* network topology. We consider a general model consisting of an arbitrary number of wireless links whose mutual interference is represented by some conflict graph. Flows of random size arrive at random at each link. In order to study the flow-level dynamics, we calculate the throughput of each flow granted by the CSMA algorithm under the usual time-scale separation assumption. We then prove that, provided there exists some schedule of the links that stabilizes the network, the flow-aware CSMA algorithm will do so, in a purely distributed and asynchronous way.

The rest of the paper is organized as follows. Related work is presented in the next section. We then present the model and analyse its stability under standard and flow-aware CSMA, respectively. The impact of network load on the mean throughput of each flow under flow-aware CSMA is considered in Section 6. Section 7 concludes the paper.

2. Related work

The problem of optimal bandwidth sharing in wireless networks has first been tackled by Tassiulas and Ephremides, who showed in [19] that the so-called *maximal weight* scheduling policy, which activates a set of links that maximizes the total backlog of active links, stabilizes any network whenever possible. A number of distributed implementations of this policy have then been proposed, all relying on some message passing protocol between nodes, see e.g. [11, 16]. Simple heuristics based on greedy algorithms that require limited or no message passing have also been studied, most selecting schedules of *maximal size* (in terms of number of links) instead of maximal weight and, as such, being suboptimal [4, 6, 8, 9, 13, 21].

A new approach to optimal scheduling has recently been proposed by Jiang and Walrand, who introduced in [7] a distributed CSMA algorithm where at each link, the attempt rate is adapted to the arrival rate and service rate so as to meet the demand. The result is based on a time-scale separation assumption whereby the activity states of the links, which depend on the CSMA algorithm, evolve much faster than the attempt rates of the links. In practice, the algorithm used for adapting the attempt rates must be carefully designed in order to guarantee convergence and optimality [7, 14]. Similar problems arise for those adaptive CSMA algorithms where the attempt rates are functions of the *queue lengths* instead of some slowly varying estimates of the arrival rates and service rates [12, 15]: the algorithm converges only for some specific choices of these functions.

In all these papers, optimality is defined either in terms of stability, assuming exogenous random packets arrivals at each link, or in terms of utility maximization, cf. [7, 14]. The

flow-level dynamics are not considered, whereas they are key to understanding network performance [18]. In particular, it can be argued that the very notion of *congestion* should be defined at the flow level [1]. In a recent paper, van de Ven, Borst and Shneer have shown that the maximal weight scheduling policy, which is known to stabilize the network at the packet level, may be unable to stabilize the network at the flow level, which highlights the difference between the two notions of stability [20]. The main contribution of the present paper is to provide an algorithm that stabilizes the network *at flow level* whenever possible. With this objective in mind, it is very natural to think of flow-aware CSMA. The fact that it suffices for each flow to run its own CSMA algorithm is far from obvious, however. It is for instance well-known that maximizing the total throughput of the network at any time may make the network unstable at flow level [3]. It turns out that the fairness imposed by the proposed flow-aware CSMA is indeed sufficient to achieve stability.

Specifically, the flow-aware CSMA algorithm selects each feasible schedule in proportion to its *weight*, where the weight of a schedule is the *product* of the number of flows on the corresponding links. For a large number of flows, the selected schedules are close to the corresponding *maximal weight* schedule (with product weights instead of additive weights), a policy that turns out to be optimal. We note that a similar property is used by Ni, Bo and Srikant in [12] for proving the stability of queue-length based CSMA at packet level. The constraints imposed by the packet level, like the above mentioned problem of time-scale separation that restricts the set of eligible weight functions, make their algorithm very different from ours, however. Our model is purely asynchronous and stateless, the number of active flows at each link being determined by the packet headers in the corresponding buffer; moreover, the time-scale separation assumption is very natural in our case since the attempt rates are adapted at the flow time-scale, which is typically much slower than the packet time-scale.

3. Model

Wireless network. We consider the general model described in [7]. There are K links in the network, where each link is an ordered transmitter-receiver pair. The network is associated with a conflict graph $G = (V, E)$, where V is the set of vertices (each representing a link) and E is the set of edges (each representing a conflict). Two links k, l can be simultaneously active if and only if they do not conflict, that is if $(k, l) \notin E$. We refer to a *feasible schedule* as any set of links $S \subset V$ (possibly empty) that do not conflict with each other. We denote by N the number of distinct feasible schedules and by S_i the set of active links in schedule i, for all $i = 1, \ldots, N$. By convention, schedule 1 corresponds to the schedule where all nodes are idle, that is $S_1 = \emptyset$.

Consider the network of $K = 3$ links depicted by Fig.2 for instance. Two links conflict if and only if the distance between the transmitter or receiver of one link and the transmitter or receiver of the other link is less than some fixed threshold. The conflict graph is linear and there are $N = 5$ feasible schedules, corresponding to the sets of active links $\emptyset, \{1\}, \{2\}, \{3\}, \{1, 3\}$.

Capacity region. Let φ_k be the physical rate of link k when scheduled, in bit/s. The throughput of link k when each schedule i is selected with probability p_i, with $\sum_{i=1}^{N} p_i = 1$, is given by:

$$\forall k = 1, \ldots, K, \quad \phi_k = \varphi_k \sum_{i:k \in S_i} p_i. \tag{1}$$
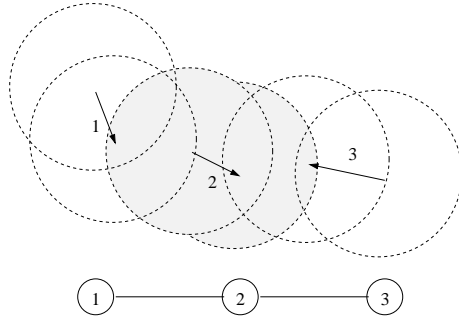
Figure 2: A 3-link network and its conflict graph.

Let ϕ be the corresponding vector. We refer to the *capacity region* as the set of vectors ϕ generated by all probability measures $p_1, \ldots, p_N$.

Flow-level dynamics. Assume that flows arrive according to a Poisson process of intensity $\lambda_k > 0$ at link k and have exponential flow sizes of mean $\sigma_k > 0$, in bits. We denote by $\rho_k = \lambda_k \sigma_k$ the traffic intensity at link k (in bit/s) and by ρ the corresponding vector. Let x_k be the number of active flows at link k. We refer to the vector x as the network state.

We shall consider random access algorithms that select each schedule i with some probability $p_i(x)$ that depends on the network state x, with $\sum_{i=1}^N p_i(x) = 1$. Under the time-scale separation assumption, the schedules change at a very high frequency compared to the flow-level time-scale, so that the throughput of link k in state x is given by:

$$\phi_k(x) = \varphi_k \sum_{i:k \in S_i} p_i(x). \tag{2}$$

The evolution of the network state then defines a Markov process $X(t)$ with transition rates λ_k from state x to state $x + e_k$ and $\mu_k(x) = \phi_k(x)/\sigma_k$ from state x to state $x - e_k$ (provided $x_k > 0$), where e_k denotes the K-dimensional unit vector on component k.

Stability condition. We are interested in the stability of the network in the sense of the positive recurrence of the Markov process $X(t)$. A necessary condition is that the vector traffic intensities ρ lies in the capacity region. We look for distributed access schemes that stabilize the network whenever possible, that is for all vectors of traffic intensities ρ in the interior of the capacity region. Such access schemes are referred to as *optimal*. For the sake of completeness, we first give an example showing the suboptimality of standard CSMA, that realizes some form of *maximal size* scheduling. We then prove the optimality of flow-aware CSMA.

4. Standard CSMA

Algorithm. We first consider a standard CSMA algorithm where each link waits for a period of random duration referred to as the *backoff time* before each transmission attempt. If the radio channel is sensed idle (in the sense that no conflicting link is active), a packet is transmitted; otherwise, the link waits for a new backoff time before the next attempt. Packets have random sizes of mean θ_k bits at link k and are transmitted at the physical rate φ_k; the backoff times are random with mean τ_k at link k. We denote by $\alpha_k = \theta_k/(\varphi_k \tau_k)$ the ratio of mean packet transmission time to mean backoff time at link k.

4

Equivalent scheduling. We look for the steady-state probability $p_i(x)$ that the set of active links corresponds to schedule i in state x. We assume that, in state x, each link k such that $x_k > 0$ takes all opportunities offered by the CSMA algorithm to transmit packets; any other link remains idle. If the packet sizes and the backoff times had exponential distributions and there were no conflict, the evolution of the set of active links S would form a reversible Markov process. A stationary measure of this Markov process is given by 1 if $S = \emptyset$ and:

$$\prod_{k \in S} \alpha_k \mathbb{1}(x_k > 0)$$

otherwise. By reversibility, the actual stationary measure induced by the conflict graph is the truncation of this measure to the set of feasible schedules. Specifically, the weight $w_i(x)$ of feasible schedule i in the stationary measure is given by:

$$w_1(x) = 1, \quad w_i(x) = \prod_{k \in S_i} \alpha_k \mathbb{1}(x_k > 0) \quad \text{for all } i = 2, \ldots, N.$$

We deduce that schedule i is selected in state x with probability:

$$p_i(x) = \frac{w_i(x)}{\sum_{j=1}^{N} w_j(x)}. \tag{3}$$

By the insensitivity property of the underlying loss network, this is also the probability that schedule i is selected in state x for arbitrary phase-type distributions of packet sizes and backoff times with the same means; such distributions are known to form a dense subset within the set of all distributions with real, non-negative support [2].

Suboptimality. We provide simple examples showing the suboptimality of the standard CSMA algorithm. We consider unit physical rates, that is $\varphi_k = 1$ for all links k. For a single link, the optimal stability condition is $\rho_1 < 1$. In view of (2) and (3), the throughput is given by:

$$\phi_1(x) = \frac{\alpha_1}{1 + \alpha_1}.$$

We deduce the actual stability condition:

$$\rho_1 < \frac{\alpha_1}{1 + \alpha_1}.$$

This loss of efficiency is due to the backoff times, that must be chosen sufficiently small to limit the overhead of the CSMA algorithm.

Now consider the example of Fig. 2 with $K = 3$ links. The optimal stability condition is given by:

$$\rho_1 + \rho_2 < 1 \quad \text{and} \quad \rho_2 + \rho_3 < 1.$$

Assume for simplicity all links have the same mean packet sizes and mean backoff times, so that $\alpha_1 = \alpha_2 = \alpha_3 = \alpha$ for some $\alpha > 0$. In view of (2) and (3), the throughput of the links in state x are given by:

$$\phi_1(x) = \begin{cases} \frac{\alpha}{1+\alpha} & \text{if } x_2 = 0, \\ \frac{\alpha}{1+2\alpha} & \text{if } x_2 > 0, x_3 = 0, \\ \frac{\alpha+\alpha^2}{1+3\alpha+\alpha^2} & \text{if } x_2 > 0, x_3 > 0, \end{cases}$$

5

and

$$\phi_2(x) = \begin{cases} \frac{\alpha}{1+\alpha} & \text{if } x_1 = 0, x_3 = 0, \\ \frac{\alpha}{1+2\alpha} & \text{if } x_1 > 0, x_3 = 0, \text{ or } x_1 = 0, x_3 > 0, \\ \frac{\alpha}{1+3\alpha+\alpha^2} & \text{if } x_1 > 0, x_3 > 0. \end{cases}$$

The throughput of link 3 follows by symmetry. As for a single link, the backoff times must be chosen sufficiently small to limit the overhead of the algorithm. In the limit $\alpha \to \infty$, we get:

$$(\phi_1(x), \phi_2(x), \phi_3(x)) = \begin{cases} (1,0,1) & \text{if } x_1 > 0, x_3 > 0, \\ (1/2, 1/2, 0) & \text{if } x_1 > 0, x_2 > 0, x_3 = 0, \\ (1,0,0) & \text{if } x_1 > 0, x_2 = 0, x_3 = 0, \\ (0,1,0) & \text{if } x_1 = 0, x_2 > 0, x_3 = 0, \end{cases} \tag{4}$$

the other cases following by symmetry. Note that link 2 is not served when both links 1 and 3 are active. This is due to the fact that link 2 is in conflict with both links 1 and 3 and thus cannot access the channel for an infinitely small backoff time. This results in a suboptimal stability region:

Proposition 1. *The stability region is given by:*

$$\rho_1 < \frac{1+\rho_3}{2}, \ \rho_3 < \frac{1+\rho_1}{2}, \ \rho_2 < \pi_0 + \frac{\pi_{1,3}}{2},$$

or

$$\rho_1 < \frac{1+\rho_3}{2}, \ \frac{1+\rho_1}{2} \le \rho_3 < \frac{1+\rho_1}{2} + \frac{1-\rho_1}{2}\pi_{2,1}, \ \rho_2 < \frac{1-\rho_1}{2},$$

or

$$\rho_3 < \frac{1+\rho_1}{2}, \ \frac{1+\rho_3}{2} \le \rho_1 < \frac{1+\rho_3}{2} + \frac{1-\rho_3}{2}\pi_{2,3}, \ \rho_2 < \frac{1-\rho_3}{2},$$

where π_0, $\pi_{1,3}$, $\pi_{2,1}$ and $\pi_{2,3}$ are the respective probabilities that:

- *both links 1 and 3 are idle when link 2 is always active;*

- *one of the links 1 or 3 is idle when link 2 is always active;*

- *link 2 is idle given that link 1 is idle, when link 3 is always active;*

- *link 2 is idle given that link 3 is idle, when link 1 is always active.*

More precisely, the Markov process $X(t)$ is positive recurrent if the vector of traffic intensities ρ lies in this region and transient if it lies outside its closure.

The proof is given in the Appendix. Note that, when one of the links is always active, the two other links form a coupled system of two queues as considered by Fayolle and Iasnogorodski [5]. In particular, the stability region can be calculated exactly. In the symmetric case $\rho_1 = \rho_3$, the stability condition reduces to $\rho_1 < 1$, $\rho_2 < \pi_0 + \pi_{1,3}/2$. Fig. 3 shows that the corresponding stability region for equal mean flow sizes.
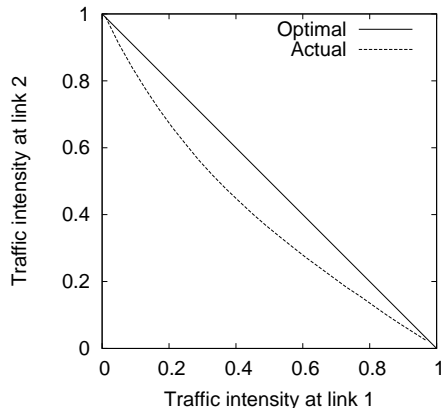
6

Figure 3: Stability condition for the network of Fig. 1 under standard CSMA ($\rho_1 = \rho_3$).

5. Flow-aware CSMA

Algorithm. We now consider the flow-aware CSMA algorithm where each *flow* (instead of each link) waits for a random backoff time before each transmission attempt. If the radio channel is sensed idle (in the sense that no conflicting link is active, nor any other flow on the same link), a packet of this flow is transmitted; otherwise, the flow remains idle for a new random backoff time before the next attempt. The backoff times have random durations of mean τ_k for each active flow at link k. We still denote by $\alpha_k = \theta_k/(\varphi_k \tau_k)$ the ratio of mean packet transmission time to mean backoff time at link k.

Equivalent scheduling. Again, we look for the steady-state probability $p_i(x)$ that the set of active links corresponds to schedule i in state x. We assume that all active flows take each opportunity offered by the CSMA algorithm to transmit packets. If the packet sizes and the backoff times had exponential distributions and there were no conflict, the evolution of the set of active links S would again form a reversible Markov process. Since there are x_k flows attempting to access the channel at link k, a stationary measure of this Markov process is given by 1 if $S = \emptyset$ and:

$$\prod_{k \in S} \alpha_k x_k$$

otherwise. By reversibility, the actual stationary measure induced by the conflict graph is the truncation of this measure to the set of feasible schedules. The weight $w_i(x)$ of feasible schedule i in the stationary measure is given by:

$$w_1(x) = 1, \quad w_i(x) = \prod_{k \in S_i} \alpha_k x_k \quad \text{for all } i = 2, \ldots, N.$$

Schedule i is then selected with probability $p_i(x)$ given by (3) in state x. By the insensitivity property of the underlying loss network, this probability remains the same for arbitrary phase-type distributions of packet sizes and backoff times with the same means, cf. [2].

Optimality. We now give the main result of the paper, that demonstrates the optimality of the above flow-aware CSMA algorithm.

7

Theorem 1. *The network is stable for all vectors of traffic intensities ρ in the interior of the capacity region.*

Proof. We apply Foster's criterion. Specifically, we look for some Lyapunov function $F(x)$ such that the corresponding drift, given by:

$$\Delta F(x) = \sum_{k=1}^{K} \lambda_k(F(x+e_k) - F(x)) + \sum_{k:x_k>0} \mu_k(x)(F(x-e_k) - F(x)),$$

satisfies:

$$\Delta F(x) \leq -\delta$$

for some $\delta > 0$, in all states x but some finite number.

If the vector of traffic intensities ρ lies in the interior of the capacity region, there exists some $\epsilon > 0$ and some probability measure $q_1, \ldots, q_N$ on the set of feasible schedules such that $q_i > 0$ for all $i = 1, \ldots, N$ and:

$$\forall k = 1, \ldots, K, \quad \rho_k = (1 - 2\epsilon)\varphi_k \sum_{i:k \in S_i} q_i. \tag{5}$$

Define:

$$F(x) = \sum_{k:x_k>0} \frac{\sigma_k}{\varphi_k} x_k (\log(\alpha_k x_k) - 1).$$

We get:

$$\Delta F(x) = G(x) + \sum_{k:x_k>0} \frac{\rho_k}{\varphi_k}(x_k + 1)(\log(1 + \frac{1}{x_k}) - 1)$$

$$+ \sum_{k:x_k>0} \frac{\phi_k(x)}{\varphi_k}(x_k - 1)(\log(1 - \frac{1}{x_k}) + 1), \tag{6}$$

with:

$$G(x) = \sum_{k:x_k>0} \frac{\rho_k - \phi_k(x)}{\varphi_k} \log(\alpha_k x_k).$$

Noting that, for any probability measure $p_1, \ldots, p_N$ on the set of feasible schedules:

$$\sum_{k:x_k>0} \sum_{i:k \in S_i} p_i \log(\alpha_k x_k) = \sum_{i=1}^{N} p_i \log(w_i(x)),$$

we get using (5):

$$G(x) = -\epsilon \sum_{i=1}^{N} q_i \log(w_i(x)) + \sum_{i=1}^{N} (q_i(1 - \epsilon) - p_i(x)) \log(w_i(x)).$$

We then need the following lemma.

Lemma 1. *Let:*

$$w(x) = \max_{i=1,\ldots,N} w_i(x).$$

Then, for all states x but some finite number,

$$\sum_{i=1}^{N} p_i(x) \log(w_i(x)) \geq (1 - \epsilon) \log(w(x)).$$

Proof. The proof is similar to that of [12, Proposition 2]. Let:

$$I(x) = \left\{ i = 1, \ldots, N : \log(w_i(x)) \geq (1 - \frac{\epsilon}{2}) \log(w(x)) \right\}.$$

We have:

$$\sum_{i=1}^{N} p_i(x) \log(w_i(x)) \geq (1 - \frac{\epsilon}{2}) \log(w(x)) \sum_{i \in I(x)} p_i(x).$$

Moreover,

$$\sum_{i \notin I(x)} p_i(x) = \frac{\sum_{i \notin I(x)} w_i(x)}{\sum_{i=1}^{N} w_i(x)},$$

$$\leq \frac{(N - |I(x)|) w(x)^{1 - \frac{\epsilon}{2}}}{w(x)},$$

$$= \frac{N - |I(x)|}{w(x)^{\frac{\epsilon}{2}}}.$$

Since $w(x)$ tends to $+\infty$ when $|x| = \sum_{k=1}^{K} x_k$ tends to $+\infty$, this quantity is less than $\epsilon/2$ for all states x but some finite number. We deduce that in all states x but some finite number:

$$\sum_{i=1}^{N} p_i(x) \log(w_i(x)) \geq (1 - \frac{\epsilon}{2})^2 \log(w(x)) \geq (1 - \epsilon) \log(w(x)).$$

$\square$

In view of Lemma 1, we have for all states x but some finite number:

$$G(x) \leq -\epsilon \sum_{i=1}^{N} q_i \log(w_i(x)) + (1 - \epsilon) \sum_{i=1}^{N} (q_i \log(w_i(x)) - \log(w(x))).$$

Since $w_i(x) \leq w(x)$ for all states x, we deduce that for all states x but some finite number:

$$G(x) \leq -\epsilon \sum_{i=1}^{N} q_i \log(w_i(x)).$$

Since $q_i > 0$ for all $i = 1, \ldots, N$, this expression tends to $-\infty$ when $|x| = \sum_{k=1}^{K} x_k$ tends to $+\infty$. The other terms of $\Delta F(x)$ in (6) being bounded, we deduce that there exists $\delta > 0$ such that $\Delta F(x) \leq -\delta$ for all states x but some finite number. $\square$

6. Throughput performance

This section is devoted to the throughput performance of flow-aware CSMA, under the stability condition. We are interested in the *mean throughput*, defined as the ratio of the mean flow size to the mean flow duration. By Little's law, the mean throughput at link k is given by:

$$\gamma_k = \frac{\rho_k}{\mathrm{E}[x_k]}. \tag{7}$$

We consider unit physical rates, that is $\varphi_k = 1$ for all links k.

Single link. We first analyse the impact of the mean backoff time on the mean throughput in the case of a single link. In the presence of x_1 flows, the total throughput is given by:

$$\phi_1(x_1) = \frac{\alpha_1 x_1}{1 + \alpha_1 x_1}.$$

The number of flows then behaves as the number of customers in a processor-sharing queue with state-dependent service rate. The corresponding stationary distribution is given by:

$$\pi(x_1) = \pi(0) \prod_{n=1}^{x_1} \frac{\rho_1}{\phi_1(n)},$$

under the stability condition $\rho_1 < 1$. The mean throughput then follows from (7). For $\alpha_1 \to \infty$, the throughput is constant and equal to 1 and the mean throughput is given by $\gamma_1 = 1 - \rho_1$; for $\alpha_1 = 1$, the system corresponds to a processor-sharing queue with an additional permanent customer representing the backoff times and we have $\gamma_1 = (1 - \rho_1)/2$; in general, we have $\gamma_1 \to \alpha_1/(1 + \alpha_1)$ when $\rho_1 \to 0$. These results are illustrated by Fig. 4.
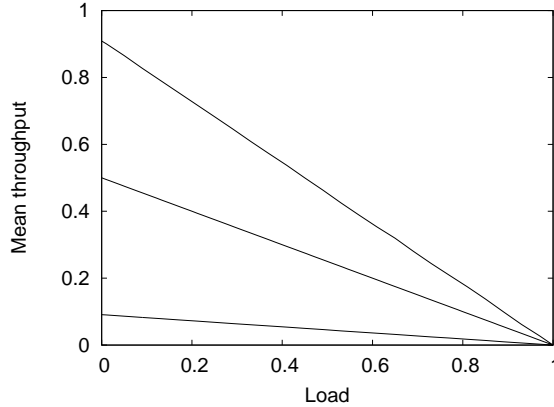

Figure 4: Impact of the mean backoff time on the mean throughput for a single link (ratio of mean packet transmission time to mean backoff time $\alpha_1 = 0.1, 1, 10$, from bottom to top).

Networks. In the following, we consider network scenarios and assume that the mean backoff time is the same for all flows and equal to the mean packet transmission time, so that $\alpha_k = 1$ for all links k. Flows have unit mean flow sizes. The traffic intensity is the same on all links, equal to ρ_1. We refer to the network load as the ratio of the per-link traffic intensity ρ_1 to its maximum value, given by the stability condition. Fig. 5 and 7 give the results obtained for the 3-link line of Fig. 2 and for the three 4-link networks of Fig. 6, for the same mean flow sizes . The results are obtained by the simulation of 10^7 jumps of the underlying Markov process, after a warm-up period of 10^5 jumps. We observe that the throughput decreases from its maximum value $1/2$ to 0 when the load grows from 0 to 1; it is lower on links that are in conflict with many other links, just like in wired networks, the mean throughput is lower on long routes, where flows go through many links [3].
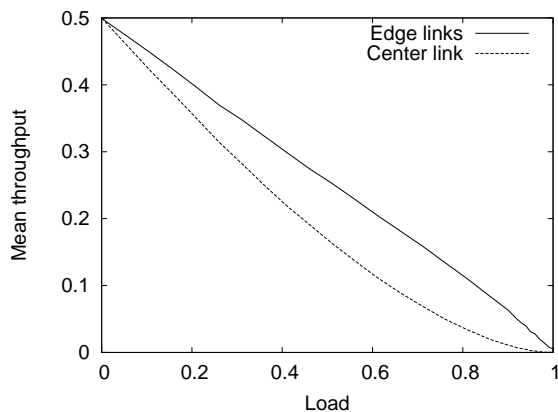
10

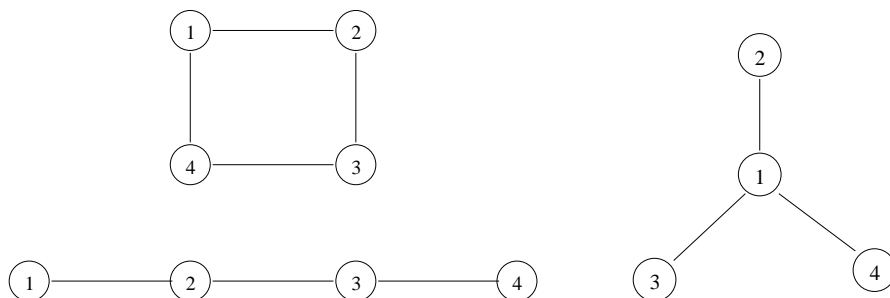Figure 5: Mean throughput in the 3-link line.

Figure 6: Conflict graphs of the square, the 4-link line and the 4-link star.

7. Conclusion

The standard CSMA algorithm is inherently unfair and inefficient. We have shown that the proposed flow-aware CSMA algorithm, where each *flow* (instead of each link) runs its own CSMA algorithm, is not only fair but efficient, in the sense that the network is stable whenever possible. To our knowledge, this is the first distributed algorithm that is provably optimal in terms of flow-level stability.

The considered packet-level model relies on a number of simplifying assumptions that we plan to relax in future work. These include the absence of collisions and hidden nodes. The interaction with the usual back-off mechanism of IEEE 802.11 should also be studied. One may also envisage different implementations of the proposed flow-aware CSMA algorithm where the attempt rate of each link is equal to some increasing function of the number of flows and the transmission opportunities are shared in a fair way between active flows, using a deficit round-robin scheduler for instance.

From a more theoretical perspective, it would be worth relaxing the assumption of exponential flow sizes and deriving bounds or approximations on the throughput performance of the algorithm.

11
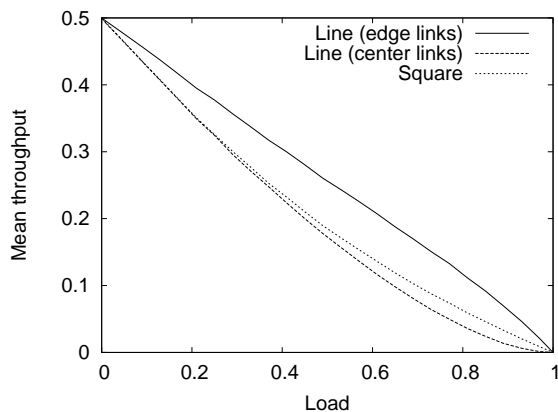
(a) Square and line
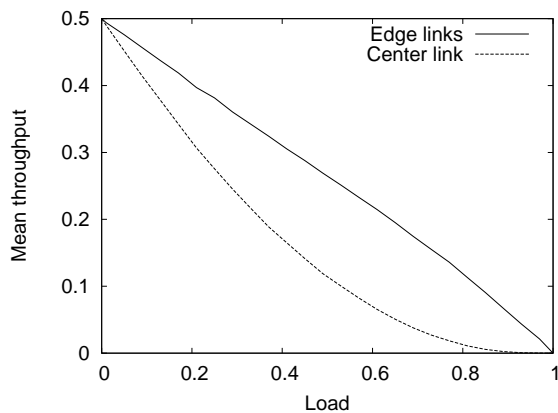
(b) Star

Figure 7: Mean throughput in 4-link networks.

Appendix

Proof of Proposition 1

This example is similar to the one studied in [17, p274]. We consider the fluid limits of the Markov process $X(t)$. Specifically, we define $X^{(n)}(t)$ as the Markov process $X(t)$ whose initial state is $X^{(n)}(0) = (\lfloor \beta_1 n \rfloor, \lfloor \beta_2 n \rfloor, \lfloor \beta_3 n \rfloor)$ for some non-negative real numbers $\beta_1, \beta_2, \beta_3$ such that $\beta_1 + \beta_2 + \beta_3 = 1$. We then define:

$$\bar{X}^{(n)}(t) = \frac{1}{n} X^{(n)}(nt).$$

The fluid limits of the Markov process $X(t)$, if they exist, are the limiting points of this set of processes when $n \to +\infty$. It is easy to check that the Markov process $X(t)$ belongs to the class $(\mathcal{C})$ defined in [17, p241] and that the associated Proposition 9.3 applies. In particular, the set $\{\bar{X}^{(n)}(t), n \in \mathbb{N}\}$ is tight and the fluid limits are continuous. The Markov process $X(t)$ is then positive recurrent if there exists some finite time after which all fluid limits are null, cf. [17, Theorem 9.7, p259]; it is transient if there exists some initial state $\beta_1, \beta_2, \beta_3$ such that,

12

after some finite time, some components of the fluid limits grow at least linearly to infinity [10].

We first calculate the fluid limit until the first time where one component reaches 0, if any, for all possible initial states. The three components of the process $X^{(n)}(t)$ behave as three coupled $M/M/1$ queues, with arrival rates $\lambda_1, \lambda_2, \lambda_3$ and state-dependent service rates. We denote by $\mu_k = 1/\sigma_k$ the maximum service rate of queue k, so that $\rho_k = \lambda_k/\mu_k$. The Markov process is positive recurrent if all queues empty in finite time in the limit and transient if, starting from some initial state, at least one queue grows linearly to infinity after some finite time.

We start with the case $\beta_1 > 0$, $\beta_2 > 0$, $\beta_3 > 0$. The three queues are then mutually independent, with respective service rates $\mu_1, 0, \mu_3$. The scaling property of the $M/M/1$ queue shows that the process $\bar{X}^{(n)}(t)$ weakly converges to the function:

$$(\beta_1 + (\lambda_1 - \mu_1)t, \beta_2 + \lambda_2 t, \beta_3 + (\lambda_3 - \mu_3)t),$$

until one of the components reaches 0, if any.

We now consider the case $\beta_1 = 0$, $\beta_2 > 0$, $\beta_3 > 0$. In view of (4), queue 1 has service rate μ_1 and is empty with probability $1 - \rho_1$. Queues 2 and 3 have service rates 0, μ_3 with probability ρ_1 and $\mu_2/2$, $\mu_3/2$ with probability $1 - \rho_1$. Proposition 9.14 of [17] applies and the process $\bar{X}^{(n)}(t)$ weakly converges to the function:

$$(0, \beta_2 + (\lambda_2 - \mu_2\frac{1 - \rho_1}{2})t, \beta_3 + (\lambda_3 - \mu_3\frac{1 + \rho_1}{2})t),$$

until one of the components reaches 0, if any.

Next, we consider the case $\beta_1 = \beta_2 = 0$, $\beta_3 > 0$. In view of (4), queue 1 has service rate μ_1. Queue 2 has service rate $\mu_2/2$ if queue 1 is empty and 0 otherwise. This queue is stable if $\rho_2 < (1 - \rho_1)/2$, which we assume. Queue 2 then remains empty in the limit, and the service rate of queue 3 is equal to μ_3 with probability $\rho_1 + (1 - \rho_1)\pi_{2,1}$ and to $\mu_3/2$ otherwise. We deduce that the process $\bar{X}^{(n)}(t)$ weakly converges to the function:

$$(0, 0, \beta_3 + (\lambda_3 - \mu_3(\frac{\rho_1 + 1}{2} - \frac{1 - \rho_1}{2}\pi_{2,1}))t),$$

whenever component 3 is positive.

Finally, we consider the case $\beta_1 = \beta_3 = 0$, $\beta_2 > 0$. In view of (4), the service rates of queues 1 and 3 are equal to μ_1 and μ_3 when both are non-empty and to $\mu_1/2$ and $\mu_3/2$ otherwise. This system is stable if $\rho_1 < (1 + \rho_3)/2$ and $\rho_3 < (1 + \rho_1)/2$, which we assume. Queues 1 and 3 then remain empty in the limit. The service rate of queue 3 is equal to μ_2 with probability π_0 and to $\mu_2/2$ with probability $\pi_{1,3}$. The process $\bar{X}^{(n)}(t)$ weakly converges to the function:

$$(0, \beta_2 + (\lambda_2 - \mu_2(\pi_0 - \frac{\pi_{1,3}}{2}))t, 0),$$

whenever component 2 is positive.

To conclude the proof, we consider the evolution of the fluid limit in the following five cases (the others follow by symmetry):

1. Assume $\rho_1 < (1 + \rho_3)/2$ and $\rho_3 < (1 + \rho_1)/2$. Note that this implies $\rho_1 < 1$ and $\rho_3 < 1$. Queue 1 and 3 empty in finite time, independently of queue 2. Queue 2 then empties in finite time if $\rho_2 < \pi_0 + \pi_{1,3}/2$; it grows linearly to infinity if $\rho_2 > \pi_0 + \pi_{1,3}/2$.

2. Assume $\rho_1 < (1 + \rho_3)/2$ and $\rho_3 > (1 + \rho_1)/2$. If $\rho_1 \geq 1$ then $\rho_3 > 1$ and queue 3 grows linearly to infinity. We now assume $\rho_1 < 1$. If $\rho_2 > (1 - \rho_1)/2$ then queue 2 grows linearly to infinity. If $\rho_2 = (1 - \rho_1)/2$ then starting from a state where $\beta_1 = 0$, $\beta_2 > 0$ and $\beta_3 > 0$, queue 1 stays empty, queue 2 is constant and queue 3 grows linearly to infinity. We assume that $\rho_1 < 1$ and $\rho_2 < (1 - \rho_1)/2$. Starting from the initial state $\beta_1 = \beta_2 = 0$, $\beta_3 > 0$, queue 3 grows linearly to infinity if $\rho_3 > (1 + \rho_1)/2 + \pi_{2,1}(1 - \rho_1)/2$. We assume that $\rho_3 < (1 + \rho_1)/2 + \pi_{2,1}(1 - \rho_1)/2$. Starting from the initial state $\beta_1 = \beta_2 = 0$, $\beta_3 > 0$, queue 3 then empties in finite time. It remains to prove that, starting from any initial state, queues 1 and 2 empty in finite time. We first note that, since $\rho_1 < 1$ and $\rho_3 < 1$, queue 1 or queue 3 empties in finite time. Moreover, if both queues 1 and 3 are empty but not queue 2, then queue 3 grows linearly. Thus we can assume that queue 1 empties before queue 3. We know that queue 2 empties in finite time in this case.

3. Assume $\rho_1 < (1 + \rho_3)/2$ and $\rho_3 = (1 + \rho_1)/2$. Note that $\rho_1 < 1$ and $\rho_3 < 1$ in this case. Moreover, we have $\pi_0 = 0$ and $\pi_{1,3} = 1 - \rho_1$, so that the inequality $\rho_2 < \pi_0 + \pi_{1,3}/2$ is equivalent to $\rho_2 < (1 - \rho_1)/2$. If the latter is satisfied, then if queue 1 is non-empty then queue 2 empties in finite time independently of queue 3. We just have to consider the case where $\beta_1 = \beta_2 = 0$ and $\beta_3 > 0$. Because $\rho_3 = (1 + \rho_1)/2 < (1 + \rho_1)/2 + \pi_{2,1}(1 - \rho_1)/2$, queue 3 empties in finite time. If $\rho_2 > (1 - \rho_1)/2$, we choose an initial state such that queue 1 empties before 3. When queue 1 is empty, queue 3 is constant and queue 2 grows linearly to infinity.

4. Assume $\rho_1 \geq (1 + \rho_3)/2$ and $\rho_3 > (1 + \rho_1)/2$. Then $\rho_1 > 1$ and $\rho_3 > 1$ so that queues 1 and 3 grow linearly to infinity.

5. Assume $\rho_1 = (1 + \rho_3)/2$ and $\rho_3 = (1 + \rho_1)/2$. Then $\rho_1 = \rho_3 = 1$ and $\pi_0 = \pi_{1,3} = 0$. If $\rho_2 = 0$, the vector ρ lies on the boundary of the stability region. If $\rho_2 > 0$, queue 2 grows linearly to infinity.

References

References

[1] Ben Fredj, S., Bonald, T., Proutière, A., Régnié, G., Roberts, J. W., 2001. Statistical bandwidth sharing: a study of congestion at flow level. In: Proceedings of ACM SIGCOMM. pp. 111–122.

[2] Bonald, T., 2007. Insensitive traffic models for communication networks. Discrete Event Dynamic Systems 17 (3), 405–421.

[3] Bonald, T., Massoulié, L., 2001. Impact of fairness on Internet performance. In: Proceedings of ACM SIGMETRICS/Performance. pp. 82–91.

[4] Dimakis, A., Walrand, J., 2006. Sufficient conditions for stability of longest-queue-first scheduling: second-order properties using fluid limits. Adv. in Appl. Probab. 38 (2), 505–521.

[5] Fayolle, G., Iasnogorodski, R., 1979. Two coupled processors: The reduction to a riemann-hilbert problem. Probability Theory and Related Fields 47 (3), 325–351.

[6] Gupta, A., Lin, X., Srikant, R., 2009. Low-complexity distributed scheduling algorithms for wireless networks. IEEE/ACM Trans. Netw. 17 (6), 1846–1859.

[7] Jiang, L., Walrand, J., 2008. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. In: the 46th Annual Allerton Conference on Communication, Control, and Computing.

[8] Joo, C., Lin, X., Shroff, N. B., 2009. Understanding the capacity region of the greedy maximal scheduling algorithm in multihop wireless networks. IEEE/ACM Trans. Netw. 17 (4), 1132–1145.

[9] Leconte, M., Ni, J., Srikant, R., 2009. Improved bounds on the throughput efficiency of greedy maximal scheduling in wireless networks. In: MobiHoc'09. ACM, pp. 165–174.

[10] Meyn, S., 1995. Transience of multiclass queueing networks via fluid limit models. Annals of Applied Probability 5, 946–957.

[11] Modiano, E., Shah, D., Zussman, G., 2006. Maximizing throughput in wireless networks via gossiping. SIGMETRICS Perform. Eval. Rev. 34 (1), 27–38.

[12] Ni, J., Tan, B., Srikant, R., 2010. Q-CSMA: Queue-length based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks. In: IEEE INFOCOM.

[13] Proutière, A., Yi, Y., Chiang, M., 2008. Throughput of random access without message passing. In: the 46th Annual Allerton Conference on Communication, Control, and Computing.

[14] Proutière, A., Yi, Y., Lan, T., Chiang, M., 2010. Resource allocation over network dynamics without timescale separation. In: IEEE INFOCOM.

[15] Rajagopalan, S., Shah, D., 2008. Distributed algorithm and reversible network. In: CISS. pp. 498–502.

[16] Rajagopalan, S., Shah, D., Shin, J., 2009. Network adiabatic theorem: an efficient randomized protocol for contention resolution. In: SIGMETRICS'09. ACM, pp. 133–144.

[17] Robert, P., 2003. Stochastic Networks and Queues. Stochastic Modeling and Applied Probability Series. Springer-Verlag, New York.

[18] Roberts, J. W., Massoulié, L., 2000. Bandwidth sharing and admission control for elastic traffic. Telecommunication Systems 15, 185–201.

[19] Tassiulas, L., Ephremides, A., 1992. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. IEEE Transactions on Automatic Control 37, 1936–1948.

[20] van de Ven, P., Borst, S., Shneer, S., 2009. Instability of MaxWeight scheduling algorithms. In: IEEE INFOCOM.

[21] Wu, X., Srikant, R., Perkins, J. R., 2007. Scheduling efficiency of distributed greedy scheduling algorithms in wireless networks. IEEE Transactions on Mobile Computing 6 (6), 595–605.