

A Packing Approach to Compare Slotted and Non-Slotted Medium Access Control

Mathilde Durvy

School of Computer and Communication Sciences
EPFL
CH-1015 Lausanne, Switzerland
mathilde.durvy@epfl.ch

Patrick Thiran

School of Computer and Communication Sciences
EPFL
CH-1015 Lausanne, Switzerland
patrick.thiran@epfl.ch

Abstract—In multi-hop ad hoc networks, the efficiency of a medium access control protocol under heavy traffic load depends mainly on its ability to schedule a large number of simultaneous non-interfering transmissions. However, as each node has only a local view of the network, it is difficult to globally synchronize transmission times over the whole network. How does the lack of global coordination affect spatial reuse in multi-hop wireless networks?

We show that in a de-centralized network the spatial reuse does not benefit from global clock synchronization. On the contrary, we demonstrate that non-slotted protocols using collision avoidance mechanisms can achieve a higher spatial reuse than the corresponding slotted protocols. By means of a simple backoff mechanism, one can thus favor the spontaneous emergence of spatially dense transmission schedules.

I. INTRODUCTION

The analysis and understanding of medium access control in single-hop wireless networks is now quite mature. However, most of the models developed for single-hop networks are hard to extend to multi-hop networks where nodes have only a local view of the network. As a consequence, the design of medium access control protocols for multi-hop ad hoc networks remains an essentially unguided quest.

In this paper, we do not analyze a specific Medium Access Control (MAC) protocol. Rather, we look at some of the fundamental choices that must be made in the design of MAC protocols. We evaluate the influence of these choices on the spatial reuse achieved by the corresponding MAC protocols. The spatial reuse measures the number of simultaneous successful transmissions per spatial unit. In multi-hop networks the spatial reuse is a good indicator of the protocol efficiency under heavy traffic loads.

Existing MAC protocols can be divided into slotted and non-slotted schemes. In slotted schemes, such as Time-Division Multiple Access (TDMA), time is divided into time slots. In each time slot, a subset of the network nodes is granted access to the channel. Provided the nodes can synchronize their clocks over the entire network, and acquire a global knowledge of its topology, it is then possible for a central authority to build

collision-free transmission schedules. However, in multi-hop ad hoc networks, the absence of a central authority makes it necessary to rely on sub-optimal distributed heuristics. In non-slotted schemes, transmissions from different nodes are not synchronized, and nodes typically contend for the medium in a decentralized manner. For example, in Carrier-Sense Multiple Access (CSMA), when a node has a packet to send, it first senses the medium. If the medium is idle, it transmits; whereas if it is busy, it backs off for a random amount of time.

In agreement with the well-known analysis of the ALOHA protocol [1], it is commonly believed that slotted protocols achieve a higher spatial reuse than their corresponding non-slotted version. Indeed, in the ALOHA protocol, a slotted scheme can greatly reduce the number of collisions between simultaneous transmissions and thus provide a higher spatial reuse. However, most of the current protocols use an additional carrier sensing mechanism that is often combined with a preliminary handshake before transmitting data packets. These mechanisms reduce considerably the number of collisions on data packets and thus the original need for synchronization. The contribution of this work is fourfold:

- 1) We show that a non-slotted protocol using an efficient collision avoidance mechanism can achieve a higher spatial reuse than its corresponding slotted version.
- 2) We demonstrate that a simple backoff mechanism creates some natural dependency between the successive transmission schedules that can favor the emergence of transmission schedules with a high spatial reuse.
- 3) We highlight the critical influence of the network topology on the ability to maintain a fair channel access at a high level of spatial reuse.
- 4) We propose a new approach to the performance analysis of MAC protocols in multi-hop networks based on a parallel between the MAC problem and the problem of packing geometrical shapes in a given area.

The paper is organized as follows. Section II introduces the system model and its assumptions. Section III describes the slotted and the corresponding non-slotted MAC protocols. Both protocols are distributed and only differ in the absence or presence of a global synchronization of the transmissions. In Section IV, we derive analytically the average spatial reuse

The work presented in this paper was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.




Fig. 1. Line topology. The transmission of connection i is successful if there is no other active node in its exclusion domain (the light gray area).

of the protocols on a simple line topology. We show that the non-slotted MAC protocol can achieve a higher average spatial reuse than the corresponding slotted MAC protocol and we analyze the reasons behind this somewhat surprising result. A common way to build a non-slotted protocol is to use a backoff mechanism. In Section V, we extend the results of the previous section to a larger class of non-slotted protocols. In particular, we analyze the influence of the backoff mechanism and packet size distribution on the spatial reuse. Finally, in Section VI, we give an aperçu of some of the trade-offs between fairness and spatial reuse. The findings of the paper are summarized in Section VII.

II. MODEL AND METRIC

A. Network Model

We assume a simple physical model where all active nodes emit at a fixed power and two nodes s and t can communicate if their distance is smaller than their *communication range* d . We call a direct *connection*, a pair of nodes that can communicate directly, and a *transmission schedule* the set of active connections at a given time.

A transmission between two connected nodes is successful if no other transmission takes place in its *exclusion domain*, defined as follows. Consider a connection i between node s_i and node t_i . The exclusion domain of a connection i is the area around this connection where no other node can be active in order to guarantee a successful transmission between nodes s_i and t_i . We assume that the exclusion domain of a connection i contains at least all the nodes in the communication range of either s_i or t_i . Any interfering transmission in the exclusion domain of an already active connection will cause a collision. In addition, we define the exclusion range R_E as the largest distance between a border point of the exclusion domain and the closest end-node of the corresponding connection. Figure 1 depicts a line network topology where nodes can communicate with their closest neighbors. The exclusion domain of connection i between node s_i and t_i is depicted in gray.

We consider bidirectional exchanges of duration T between connected nodes. A typical exchange consists of a handshake followed by a data packet and an acknowledgment. Consequently, a transmission from s_i to t_i or from t_i to s_i will result in the same exclusion domain since both nodes should be able to receive data.

Our assumptions are meant to create a framework similar to the IEEE 802.11 framework [2], where all nodes in the exclusion domain of a connection are silenced after the initial Request-To-Send (RTS) and Clear-To-Send (CTS) messages.




Fig. 2. Grid topology. The set of active connections depicted correspond to a transmission schedule with maximal spatial reuse. In such a schedule each transmission ‘occupies’ 4 space units (the dashed hexagons).

B. Spatial reuse

Our primary interest is in the design of MAC protocols for multi-hop ad hoc networks with high traffic loads - we assume saturated traffic conditions. In this setting, the throughput performance of a MAC protocol is dictated mainly by its ability to schedule a high number of simultaneous transmissions. Accordingly, we decided to take as our main metric the number of successful simultaneous transmissions per spatial unit, that is commonly called *spatial reuse* and that we denote by σ .

Unfortunately, the problem of maximizing the number of simultaneous successful transmissions in a network is NP-Complete (by equivalence with the maximal independent set problem [3]) and can only be solved for small topologies. In order to compare the spatial reuse achieved by a MAC protocol to the maximum achievable spatial reuse, one must thus resort to network topologies with specific symmetries. Consider, for example the line topology of Figure 1 and the grid topology of Figure 2. Let the distance between two neighboring nodes be equal to 1 space unit and assume that the communication range is also 1. In the line topology at most one connection out of three can be active simultaneously (in order to avoid collisions), hence the maximal spatial reuse is $1/3$. Similarly, in the grid topology, each active connection occupies at least four space units, the maximal spatial reuse is $1/4$.

III. MEDIUM ACCESS CONTROL PROTOCOLS

In this section we present a slotted and a corresponding non-slotted MAC protocol. We require both protocols to be distributed; no central authority can compute the transmission schedules. We assume a carrier sensing and collision avoidance mechanism that can accurately detect interfering transmissions in the exclusion domain of a connection. The protocols determine the state of a connection (active or idle) based exclusively on this local information. No global knowledge on the network is available.

A. Slotted MAC Protocols

In a slotted or synchronous MAC protocol, time is divided into slots of length T . In each time slot, a different set of direct connections is active.

Clearly, as already mentioned in Section I, if a central authority was able to compute the optimal assignment of the time slots between the different nodes, the resulting slotted MAC protocol would outperform the non-slotted protocols, but would not obey the requirement of being decentralized.




Fig. 3. Illustration of the two backoff dynamics. On the top (respectively, bottom), the non-frozen (resp., frozen) backoff dynamics performed by two neighboring connections 1 and 2. At the end of a transmission, the first connection to reach the end of a backoff interval (i.e., to have a backoff value equal to zero) emits. In the non-frozen backoff dynamic, connections constantly decrease their backoff value and pick a new backoff value as soon as their backoff value reaches zero. In the frozen backoff dynamic, connections decrease their backoff value only when they do not detect any activity in their exclusion domain. Moreover, an active connection waits until the end of its transmission to pick a new backoff value.

To obtain a decentralized slotted protocol, using the same local information as the non-slotted protocol and offering a high degree of spatial reuse, the set of active connections is built as follows. At the beginning of a time slot, all connections attempt to transmit in a random order. A connection becomes active if it does not detect any other already active connection in its exclusion domain. The set of active connections is complete when no other additional non-interfering connection can be added to the set. In other words, adding one more connection to the set would create a collision. We assume that this construction of the set of active connections takes a negligible time, that is included in T . At the end of a time slot, the building process is repeated, and another set of connections becomes active. In the slotted protocol, all transmissions taking place in a slot start at the same time and stop simultaneously as well; we speak of synchronous transmissions.

B. Non-slotted MAC Protocols

A standard way to construct a non-slotted or asynchronous MAC protocol is to use a backoff mechanism. Each direct connection is assigned a backoff value. Initially, the backoff value is chosen according to a *backoff distribution* (e.g. uniform in $[0, cw]$). We denote the mean of the backoff distribution by \bar{cw} . The backoff value of a connection is then decreased with time. When its backoff value reaches zero, a connection becomes active provided again that no connection is already active in its exclusion domain. Otherwise, a new backoff value is drawn randomly according to the backoff distribution, and the procedure is repeated. This *backoff dynamic* is illustrated on the top of Figure 3 for two neighboring connections. An alternative dynamic (followed for example by the IEEE 802.11 protocol) is to *freeze* the backoff of a connection when other

| Exchange time distribution | Backoff value distribution | backoff dynamics |
|---|---|-------------------------------|
| $C = \text{Constant}$ value T | $U = \text{Uniform in } [0, cw]$ with mean $\bar{cw} = cw/2$ | $F = \text{Frozen}$ |
| $E = \text{Exponential}$ with mean T | $E = \text{Exponential}$ with mean \bar{cw} | $\bar{F} = \text{not Frozen}$ |

TABLE I
ABBREVIATIONS FOR THE DIFFERENT BACKOFF MECHANISMS AND EXCHANGE TIME DISTRIBUTIONS

connections in its exclusion domain are active. In the context of single-hop networks, freezing the backoff of a connection when another connection is active provides a natural way to guarantee long-term fairness between connections [4]. The bottom of Figure 3 illustrates the frozen backoff dynamics: whenever a connection is active, its neighboring connection freezes its backoff. At the end of the transmission, the first connection picks a new backoff value, while the latter resumes the decremental process of its backoff value.

Connections remain active for the duration of an exchange time. We consider exchange times whose length is either a constant T , or is exponentially distributed, with mean T . Table I introduces the notation used to describe the different backoff mechanisms and exchange time distributions. In general, we identify a non-slotted protocol by the triplet (*exchange time distribution, backoff time distribution, backoff dynamics*). A triplet element equal to $*$ can take all possible values. In the non-slotted protocol, transmissions do not necessarily start or stop simultaneously; we speak of asynchronous transmissions.

C. Related Work

Surprisingly, relatively few papers analyze the performance of MAC protocols in the multi-hop setting. In the category of non-slotted protocols, most of the recent work is directed toward a multi-hop analysis of the IEEE 802.11 protocol. The first attempt to analyze the multi-hop behavior of the IEEE 802.11 protocol is, to the best of our knowledge, [5] who provides a lower bound on the throughput of the IEEE 802.11 protocol under the assumption that nodes transmit independently according to some Poisson process. The work by J.J. Garcia-Luna-Aceves and his coauthors is probably among the most advanced studies to date. [6] and [7] derive an analytical model to compute the throughput of IEEE 802.11 under saturated traffic condition. In [7], MAC protocols are approached as dynamical systems which take as input the successful transmission probabilities and output the scheduling rates. However, it is in general hard to establish the transfer function of such systems and assumptions need to be made.

In the category of slotted protocols, a lot of work has been done in the context of probabilistic MAC protocols. In probabilistic protocols a connection transmits a data packet in a given time slot with probability p . The SR-Aloha protocol [8] and the SEEDEx protocol [9] are two examples of distributed probabilistic protocols. They differ in their way to assign p




Fig. 4. Analogy between a medium access control problem and a packing problem. A transmission schedule corresponds to a set of non-overlapping intervals on the line.

to the different nodes. Probabilistic protocols are often easier to analyze since they explicitly assign a probability of being active to each node or connection. However, contrary to the slotted protocol described in this paper, they do not guarantee collision free transmission schedules and their spatial reuse can be severely affected by collisions.

IV. INFLUENCE OF SYNCHRONIZATION

The goal of this section is to compare the spatial reuse achieved by the slotted and non-slotted MAC protocols. We proceed in two steps. First, we compute analytically the average spatial reuse of the protocols on a line network topology. Second, we generalize by simulation the one-dimensional results to more general two-dimensional topologies.

A. Line Topology

To compute the average spatial reuse, σ , achieved by the slotted and non-slotted protocols we use a novel approach where we identify each of the protocol to a packing process that achieves on average the same spatial reuse. The study of packing process has a long history and resurfaced in the late fifties with, for example, the well known Rényi's parking problem [10]. Unfortunately, most of the exact results are only available in one dimension. We thus limit our analysis to a line network topology where nodes are placed at integer coordinate $\{0, 1, \dots, L\}$ and can communicate with their closest neighbors, i.e., the communication range $d = 1$. This simplistic setting is necessary to derive formal results and will be relaxed in the next sections.

1) *Spatial reuse of the slotted protocol:* In the slotted protocol, the process for building the transmission schedule is identically distributed at each time slot, and independent from time slot to time slot. The average spatial reuse at steady state is thus equal to the average spatial reuse observed in any of the time slots. Consider the process of building a transmission schedule for a given time slot. To each active connection we can associate an interval of length l on the line, which corresponds to the portion of the line 'occupied' by an active connection (Figure 4). According to the protocol, intervals of length l arrive in a random order and are accepted as long as they do not overlap with intervals already accepted. The construction of a transmission schedule for a given time slot is then equivalent to the *filling* of the line with as many non-overlapping intervals as possible. However, because of the randomness of the filling process, a part of the line might remain uncovered by the intervals. In 1962, Mackenzie [11]

used a recursive approach to compute the average fraction of the line left vacant when the length of the line network goes to infinity. Indeed, each accepted interval, reduces the initial problem to two independent, smaller instances of the same problem.

Exploiting Mackenzie's result, the average spatial reuse σ_f achieved by the slotted protocol when $L \rightarrow \infty$ can be recast as

$$\sigma_f \sim \exp[-2F(1)] \int_0^1 \exp[2F(u)] du, \quad (1)$$

where

$$F(u) = \int_0^u \frac{1 - y^{l-1}}{1 - y} dy = \sum_{i=1}^{l-1} \frac{u^i}{i}.$$

For $l = 3$, $\sigma_f = \int_0^1 \exp(u^2 + 2u - 3) du \simeq 0.275$ so that we are approximately 18% away from the maximal spatial reuse of $1/3$. For larger l , the spatial reuse decreases: the larger the exclusion domain, the more inefficient the slotted protocol.


2) *Spatial reuse of the non-slotted protocol:* In the non-slotted protocol, a given one-hop connection attempts to seize the channel at rate $1/\bar{c}\bar{w}$ and releases the channel at rate $1/T$. Like in the slotted protocol, a connection becomes active only if it does not collide with already active connections. Consequently, the set of active connections at a given time also corresponds to a set of non-overlapping intervals (of length l) on the line. In the packing formalism the non-slotted protocol is analogous to a *filling and depletion* process on the line. This kind of process has been studied in detail by Kelly, Zachary, Ziedins and other researchers [12], [13], [14] in the context of circuit-switched loss networks and was recently revisited by Baryshnikov et al. in [15].

The non-slotted protocol with exponentially distributed exchange times and backoff times can be modeled as a continuous-time Markov chain whose states correspond to the collision-free transmission schedules, or equivalently for the line network, to the sets of non-overlapping intervals on the line. The transition rate between a given collision-free transmission schedule and the same transmission schedule plus (respectively, minus) one transmission is $1/\bar{c}\bar{w}$ (resp., $1/T$). Figure 5 illustrates this Markov chain model on a toy line topology of 8 nodes. The left part of the figure depicts the valid collision-free transmission schedules, while the right part of the figure represents the corresponding continuous-time Markov chain. The stationary distribution of the Markov chain can then be computed using the global balance equations and the stationary probability of a *specific* transmission schedule with i transmissions can be expressed as

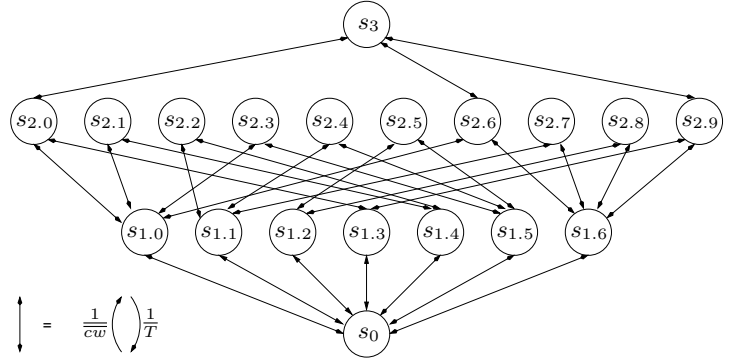
$$\pi(i) = (T/\bar{c}\bar{w})^i \pi(0), \quad (2)$$

where $\pi(0)$ is the stationary probability of the empty transmission schedule. Denoting by $N(i)$ the number of collision-free transmission schedules with i transmissions, we have

$$\pi(0) = \left(\sum_k N(k) (T/\bar{c}\bar{w})^k \right)^{-1}. \quad (3)$$



(a) Collision-free transmission schedules. There are 19 schedules $s_{i(\cdot)}$ with a number of transmissions $i \in \{0, 1, 2, 3\}$.



(b) State transition diagram of the Markov chain.

Fig. 5. Continuous-time Markov chain model of the non-slotted protocol for a small 8 node line topology. Each collision-free transmission schedule corresponds to a state of the Markov chain. Transition between states are symmetric, with rate $1/c\bar{w}$ upwards and $1/T$ downwards.

Equations (2) and (3) are not specific to the line network topology. Unfortunately, it is in general not possible to obtain a closed form expression for $N(i)$. However, in the case of a line topology we have

$$N(i) = \binom{i+v}{i},$$

where v is the length of the vacant space on the line once i intervals of length l are accepted. The moment generating function of the number of transmissions per transmission schedule can then be written as

$$Z(x) = \sum_i \left(N(i) (T/c\bar{w})^i \pi(0) \right) x^i.$$

Taking the first order derivative at $x = 1$ and dividing by the length of the line gives the average spatial reuse of the non-slotted protocol. Performing this computation is however non trivial and we refer the reader to [15] for a complete derivation. The resulting average spatial reuse obtained for the non-slotted protocol when $L \rightarrow \infty$ is

$$\sigma_{fd} \sim \frac{l(T/c\bar{w}) y_1^{l-1}}{1 + l(T/c\bar{w}) y_1^{l-1}} \quad (4)$$

where y_1 is the real root of $1 - y - (T/c\bar{w}) y^l$ closest to the origin.

Throughout the analysis we assumed a non-slotted protocol with exponentially distributed exchange times and backoff times. The goal of this assumption was to make the Markov model relatively straightforward. However, it is known [16] that (2) is insensitive to the exchange time distribution. Moreover, the backoff dynamic does not matter as long as the backoff distribution is exponential because of the memoryless properties of the exponential distribution. The average spatial reuse (4) is thus valid for all the non-slotted protocols with an exponential backoff mechanism.




Fig. 6. The clustered topology. The four shaded areas and the five transparent areas represent zones of high node densities (1.625 nodes per space unit) respectively low node density (0.5 nodes per space unit).

B. 2D Topologies

We now extend numerically the one-dimensional results to two-dimensional topologies.

1) *Simulation Settings:* In addition to a line topology of 50 nodes, we consider three types of two-dimensional topologies. All 2D topologies have on average 100 nodes that are located in a 10×10 square area. In the grid topology, nodes are placed on a grid at a distance of one space unit of each other. In the irregular Poisson topology, nodes are randomly distributed as a Poisson process of intensity one. Finally, in the clustered topology (Figure 6) the original square area is divided into nine zones, four with high node density and five with low node density. The line and the grid topologies are connected topologies that typically arise in situations where the node positions can be chosen. The Poisson and the clustered topologies are random topologies that might emerge if one cannot control the individual positioning of the nodes. We pick a communication range $d = \sqrt{4/\pi}$ so that, for both

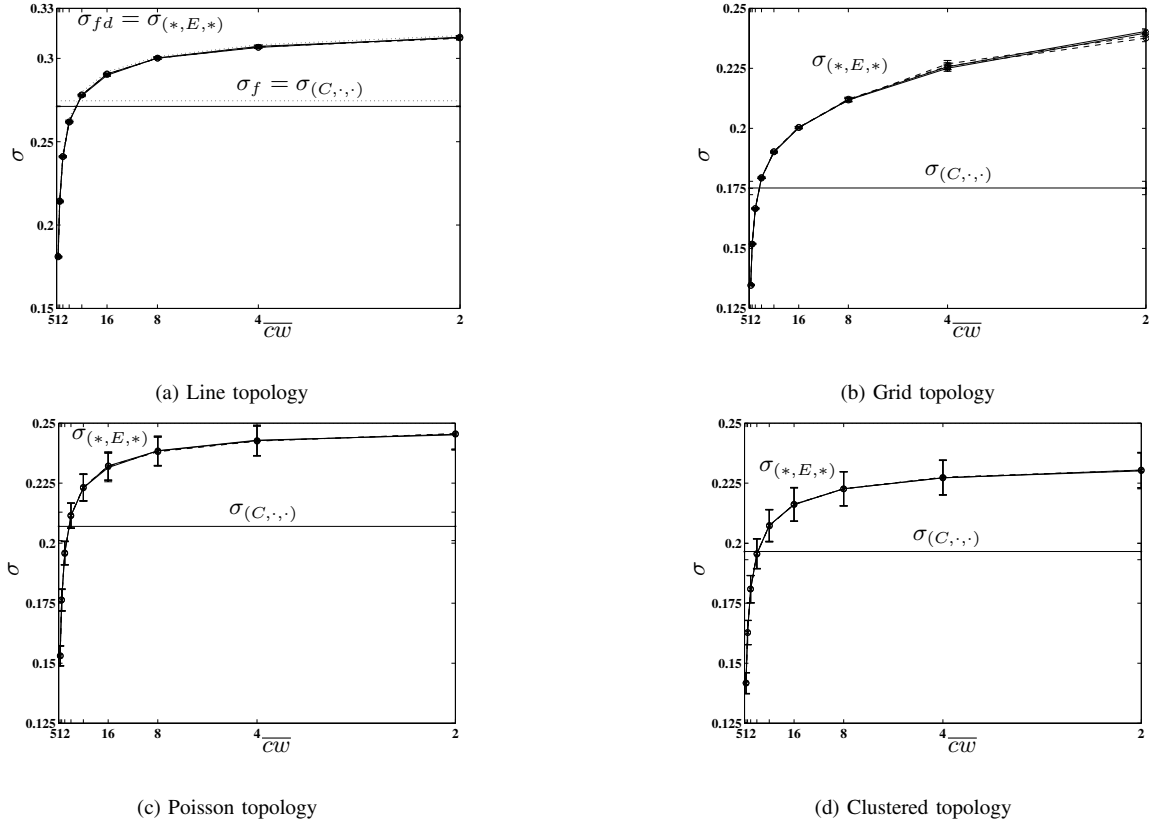


Fig. 7. The average spatial reuse of the slotted and the non-slotted protocols as a function of T/\overline{cw} . As the value of T is fixed to 420, we label the x axis by decreasing values of \overline{cw} . The non-slotted protocols with exponential backoff distribution are denoted by $(*, E, *)$ while the slotted protocol is denoted by (C, \cdot, \cdot) . The spatial reuse of all non-slotted protocols converges towards the maximal spatial reuse ($1/3$ in the line topology, $1/4$ in the grid topology) as the average backoff time \overline{cw} decreases.

the Poisson and the grid topology, each node has on average 4 direct connections (5.25 in the clustered topology). The exclusion range is set to the communication range. Taking the exclusion range larger than the communication range scales down the numerical results obtained but does not change the behavior of the different protocols.

For completeness we now introduce the simulation parameters. To make the simulations more realistic we aligned most of the parameter values on the IEEE 802.11 specifications [2]. The average backoff \overline{cw} takes value in $\{2, 4, 8, \dots, 512\} \cdot (20\mu s)$. The average exchange time, T , is fixed to $420 \cdot (20\mu s)$, the time of a RTS-CTS-DATA(1000bytes)-ACK exchange in the IEEE 802.11 protocol. All simulations assume saturated traffic conditions (i.e., nodes have always something to send). To provide accurate results, each simulation is repeated 50 times (using different random seeds). The simulations run for 50s, yet, in order to remove the transient behavior of the protocols, only the 10 last seconds of the simulations are taken into account. The simulation results were obtained with [17] and consist of the values averaged over the 50 experiments, and the 95% confidence intervals. All figures follow the same format, we plot the average spatial reuse as a function of T/\overline{cw} . However, as T is fixed, we label the x axis with the value of \overline{cw} .

2) *Simulation Results:* According to the notations of Table I, the non-slotted protocol with exponential backoff distribution can be identified by the triplet $(*, E, *)$. Similarly, the slotted protocol can be identified by the triplet (C, \cdot, \cdot) , i.e., the backoff parameters are inexistent and the exchange time distribution is by definition always constant. Figure 7(a) shows the value of σ_f (Eq. (1)) and σ_{fd} (Eq. (4)) for $l = 3$ on an infinite network topology. Despite the limited network size, the average spatial reuse obtained by simulating the respective MAC protocols (C, \cdot, \cdot) and $(*, E, *)$ on a line topology of 50 nodes agrees very well with the analytical results. Figures 7(b), 7(c) and 7(d) show the average spatial reuse achieved by the different protocols on respectively, the grid, the Poisson, and the clustered topology.

For all network topologies we observe that the average spatial reuse of the non-slotted protocols increases as their average backoff time \overline{cw} decreases. Indeed, for small \overline{cw} , the average spatial reuse achieved by the non-slotted protocols approaches the optimal value of $1/3$ for the line topology and $1/4$ for the grid topology. Consequently, if the average backoff value is low enough, i.e., if the network nodes try to access the channel often enough, it becomes advantageous to use a non-slotted protocol instead of our reference slotted protocol. At first sight this result is rather counter-intuitive; in

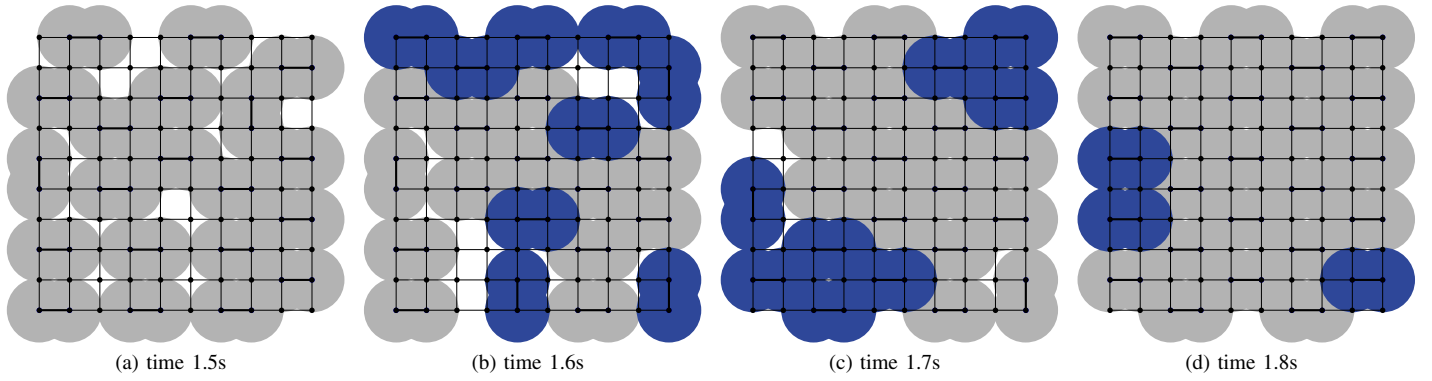
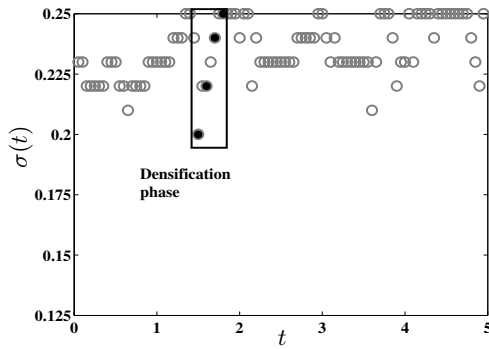



Fig. 8. Densification of the transmission schedule under the non-slotted protocol (E, E, \bar{F}) ($\bar{c}\bar{w} = 2$). Nodes are positioned on a grid and each node can communicate with its four closest neighbors. Active connections are represented by a thick black line. We depict the exclusion domain of the recently active connections by two dark gray disks (one around each end-node of the connection), the exclusion domain of connections that were already active in the previous transmission schedule are depicted in light gray.



(a) Spatial reuse at time t observed in one specific simulation instance of the non-slotted protocol.



(b) Spatial reuse at time t averaged over 50 simulations. The average spatial reuse of the non-slotted protocol takes a few seconds to converge to its final value. As expected, the average spatial reuse of the slotted protocol remains constant over time.

Fig. 9. Evolution of the spatial reuse of the non-slotted protocol (E, E, \bar{F}) ($\bar{c}\bar{w} = 2$) and slotted protocol over time (on the grid network).

the non-slotted protocols, connections spend some additional time idle due to the backoff mechanism. They should therefore be active less frequently and the non-slotted protocols should exhibit a lower spatial reuse. Why is that not the case? Why does the non-slotted protocols perform so well?

3) *Deciphering of the Results:* We now show that the non-slotted protocols using a small average backoff time *automatically* favor transmission schedules with a high number of simultaneous transmissions. Combining Equations (2) and (3), we can write

$$\pi(i) = \frac{(T/\bar{c}\bar{w})^i}{\sum_k N(k) (T/\bar{c}\bar{w})^k} \quad (5)$$

where $\pi(i)$ is the steady state probability, under the non-slotted protocols, of an individual, collision-free transmission schedule with i transmissions. For $\bar{c}\bar{w} < T$, $\pi(i)$ increases with the value of i . Consequently, the transmission schedules with a high number of active connections have an increased probability to appear, compared to those with only a few active connections. *In the limit* $T/\bar{c}\bar{w} \rightarrow \infty$, *only the transmission*

schedules that maximize the spatial reuse have a non-trivial probability. It is thus not very surprising that the non-slotted protocols should achieve a very high level of spatial reuse at low average backoff times.

To better illustrate the operation of the non-slotted protocols, consider the case where $\bar{c}\bar{w} \ll T$. The non-slotted protocols first fill the space with active transmissions until no other additional non-interfering transmission can be added. At the end of this first phase the spatial reuse is equal (on average) to $\sigma_{(C, \dots)}$ (the spatial reuse under the slotted protocol). Thereafter, the transmission schedule slowly densifies as an ending active connection is almost immediately replaced by one or *more* new active connections. Figure 8 shows a densification phase observed in one of the 50 simulations of the non-slotted protocol (E, E, \bar{F}) ($\bar{c}\bar{w} = 2$) on the grid network. The number of simultaneous transmissions increases from 20 ($t = 1.5s$, Figure 8(a)) to 22 ($t = 1.6s$, Figure 8(b)) to 24 ($t = 1.7s$, Figure 8(c)) and finally to 25 ($t = 1.8s$, Figure 8(d)). The corresponding levels of spatial reuse $\sigma(t)$ are shown in Figure 9(a). Due to the constant arrivals (i.e., starts) and departures (i.e., ends) of transmissions, the spatial

reuse fluctuates a lot. However, the average over the 50 simulations presented in Figure 9(b) shows a clear, although progressive, increase of the spatial reuse achieved by the non-slotted protocol in the first five seconds of simulation. Indeed, once a transmission schedule that maximizes the spatial reuse is reached, the probability of going back to a transmission schedule with a low spatial reuse is small as it would imply the end of several transmissions between two consecutive transmission attempts. This event is unlikely when the average exchange time is much larger than the average backoff time. In the case of the slotted protocol, the set of transmissions in a time slot is independent of the set of transmissions in the other time slots; the spatial reuse remains constant over time.

The combination of the collision avoidance and random backoff mechanisms of the non-slotted protocol creates a local coupling between the connections, which is strong enough to spatially organize the transmissions in quasi optimal global schedules. A similar densification phenomenon occurs in the context of the packing of granular materials [18]. In essence, the additional randomness introduced by the backoff mechanism of the non-slotted protocol produces the same effect as shaking a bucket of sand to achieve a better packing of the sand inside.

V. INFLUENCE OF THE BACKOFF AND EXCHANGE TIME DISTRIBUTIONS

In this section we study how the backoff and the exchange time distributions can affect the spatial reuse achieved by non-slotted MAC protocols. In addition to the exponential backoff distributions introduced in the mathematical analysis, we investigate uniform backoff distributions, which are common in practice. We consider both constant and exponentially distributed exchange times.

We proceed in two steps. First we complete our analysis of MAC protocols using a perfect collision avoidance mechanism, and second we move to the more general class of MAC protocols with a vulnerable collision avoidance mechanism.

A. Perfect Collision Avoidance Mechanism

By perfect collision avoidance mechanism, we mean that all nodes can determine accurately and instantaneously whether the channel is busy or not. This is of course not the case in most protocols, but it can be achieved (see for example the floor acquisition multiple access (FAMA) protocol [19] or busy-tone protocols [20]). In this setting all transmissions complete successfully: there is no collision.

Figure 10 shows that independently from the backoff dynamic and exchange time distribution, the average spatial reuse σ achieved by the non-slotted protocols with uniform backoff distribution increases as the average backoff time $\bar{c}\bar{w}$ decreases. In particular, all the non-slotted protocols exhibit a higher spatial reuse than the slotted protocol. The densification phenomenon described in the previous section is thus not limited to protocols with an exponential backoff distribution.

Non-slotted protocols using an exponential backoff distribution are insensitive to the exchange time distribution and to the

| | | |
|---|----------|--|
| Uniform in $[0, c\bar{w}]$ | \equiv | Uniform in $\{1, 2, \dots, c\bar{w} - 1\}$ |
| Exponential, with mean $\bar{c}\bar{w}$ | \equiv | Geometric, with mean $\bar{c}\bar{w}$ |

TABLE II

CORRESPONDING CONTINUOUS AND DISCRETE TIME DISTRIBUTION.

backoff dynamics. As we have already seen, this insensitivity stems from the very special memoryless property of the exponential variable. Indeed, although non-slotted protocols using a uniform backoff distribution remain insensitive to the exchange time distribution, they are sensitive to the backoff dynamics. More specifically, protocols of type $(*, U, \bar{F})$ achieve on average a higher spatial reuse than protocols of the type $(*, U, F)$, that have the same level of spatial reuse as $(*, E, *)$ protocols. In the next section we show that the price to pay for this increase in spatial reuse is a lower fairness.

Although the exchange time distribution has no influence on the average level of spatial reuse achieved by the protocols in the last 10s of the simulations, it impacts the transient behavior of the protocols. Protocols with constant exchange times take longer to converge to their final level of spatial reuse than protocols with exponential exchange times. For example, we have seen in Figure 9(b) that a protocol of type (E, E, \bar{F}) , with $\bar{c}\bar{w} = 2$, takes only a couple of seconds to converge to its final level of spatial reuse in the grid topology. We found that a protocol of type (C, E, \bar{F}) , with the same value of $\bar{c}\bar{w}$, takes between 10s and 15s to converge. The additional randomness introduced by the exponential exchange times can thus help a protocol to converge more rapidly toward spatially efficient transmission schedules. Constant exchange times tend to synchronize the attempts to seize the wireless medium, which slows down the densification process.

B. Vulnerable Collision Avoidance Mechanism

Despite the use of carrier sensing and collision avoidance mechanisms, most of the existing non-slotted MAC protocols cannot completely avoid collisions. In particular, it is not rare that two nodes attempting to transmit at very close time instants cannot detect each other and that the resulting transmissions collide. For example in the IEEE 802.11 protocol, the backoff counter takes values that are integer multiples of a backoff slot of $20\mu s$. Two neighboring nodes attempting to transmit in the same backoff slot collide. An easy way to extend our study to MAC protocols with vulnerable collision avoidance mechanism is therefore to discretize the backoff and exchange time distributions. Table II displays the continuous time distributions and their discrete time equivalent. The spatial reuse achieved by the non-slotted protocols with a uniform backoff distribution and a vulnerable collision avoidance mechanism is displayed in Figure 11. These protocols have a very similar behavior in the Poisson and clustered topology. To avoid redundancy in the figures and save some space we only show the results for the Poisson topology. The level of spatial reuse achieved by the slotted protocol with perfect collision avoidance mechanisms is shown for reference.

Collisions have a strong negative effect on the spatial

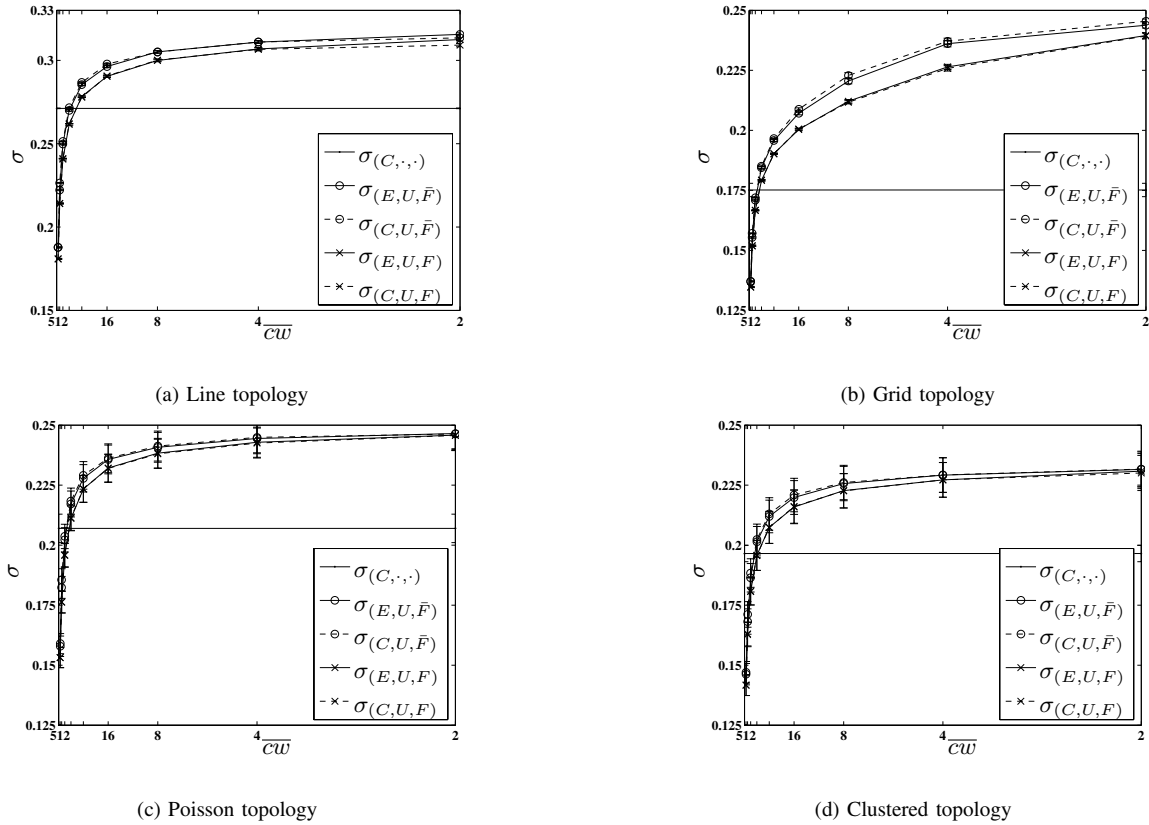


Fig. 10. The average spatial reuse σ achieved by the non-slotted protocols with *uniform backoff distribution* and *perfect* collision avoidance mechanism. Protocols with an uniform backoff distribution achieve the same high level of spatial reuse as protocols with an exponential backoff distribution.

reuse achieved by all non-slotted protocols. At large average backoff values, protocols with vulnerable, respectively, perfect collision avoidance mechanisms, achieve a very similar level of spatial reuse. However, using protocols with vulnerable collision avoidance mechanisms makes it impossible to arbitrarily lower the average backoff time $\bar{c}w$ to achieve a higher spatial reuse. In particular, in the Poisson and the clustered topologies, the non-slotted protocols outperform the slotted protocol on a very restricted set of $\bar{c}w$. Although the comparison between the two protocols is now biased by the fact that we still assume a perfect collision avoidance mechanism for the slotted protocol but not for the non-slotted one, it reveals that the average backoff value of non-slotted protocols with vulnerable collision avoidance mechanism must be carefully selected. This result is in agreement with [21] that already reported, in single-hop networks, the sensitivity of the IEEE 802.11 protocol to the average backoff value.

Protocols with vulnerable carrier sensing exhibit a sensitivity to the exchange time distribution that was inexistent in the protocols with perfect carrier sensing. Consider two colliding transmissions. In protocols with constant exchange times they start and end at the same time; the collision length is equal to the exchange time T . In protocols with exponential exchange times, the two transmissions also start at the same time but end at different times; the collision length is equal to the *maximum* of the two exchange lengths,

which is equal on average to $3T/2$. For an equal number of collisions, protocols using constant exchange times spend thus less time in collision and can, as a result, send more packets. This capability is beneficial as long as the probability of sending successfully a packet remains reasonably high. At low average backoff value, this is obviously not the case and the spatial reuse achieved by constant exchange length protocols systematically drops under the spatial reuse achieved by their exponential exchange length counterparts. The combination of constant exchange times with the frozen backoff dynamic is particularly bad as it tends to synchronize the transmissions, resulting in repeated collisions. In practice, it is common to double the contention window after a collision and reset it after a successful transmission (the IEEE 802.11 protocol implements such a mechanism). This exponential increase of the contention window can potentially reduce the number of collisions and thus increase the spatial reuse.

VI. SPATIAL REUSE VERSUS FAIRNESS

In a multi-hop ad hoc network with high traffic load, the average spatial reuse achieved by a protocol determines the average amount of data that can be successfully transmitted by the protocol in the whole network. However, it does not give any guarantee in terms of per connection data rate.

To assess the link layer fairness of the protocols we use the Jain's Fairness Index (FI) [22]. Consider a network with




Fig. 11. The average spatial reuse σ achieved by the non-slotted protocols with *uniform backoff distribution* and *vulnerable* collision avoidance mechanism. Due to the strong negative impact of collisions on the densification process one cannot arbitrarily lower the average backoff value \overline{cW} in order to obtain a higher level of spatial reuse.




Fig. 12. The level of fairness achieved by the non-slotted protocols with uniform backoff distribution and exponential sending times. The protocols using a perfect (vulnerable) collision avoidance mechanism are represented by a solid (dashed) line. The marker type identifies the backoff dynamic of the protocol.

n one-hop connections and denote by (x_1, x_2, \dots, x_n) , the number of packets sent per connection under a given MAC protocol. The fairness index of the protocol is

$$FI = \frac{(\sum_i x_i)^2}{n \sum_i x_i^2}.$$


This definition of the fairness index has the advantage to give very intuitive values. For example, in a network where all connections access the channel equally the fairness index is 1. Yet, if only k of the n connections have an equal access to the channel and the remaining connections have no access to the channel, the fairness index is k/n .

Figure 12 shows the short-term fairness achieved by the non-slotted protocols with a uniform backoff distribution. The short-term fairness is computed based on the packets sent during the last 10 seconds of the simulations. For a better accuracy (due to the faster convergence), and to limit the number of graphs, we plotted the results for the exponential exchange time distribution, the constant exchange time distribution yields similar results.


At large average backoff times, the fairness index of the protocols merely reflects the (un)fairness of the underlying

network topology. On the one hand, topologies with a regular placement of the nodes such as the grid or the line topology are inherently fair (Figures 13(a) and 14(a)). The only source of unfairness is due to the boundary effect: the connections at the border, having fewer nodes in their exclusion domain, have more often the opportunity to send a packet. In such topologies, the non-slotted protocols and the slotted protocol can reach an equally high level of fairness. On the other hand, topologies with a non-regular positioning of the nodes such as the Poisson or the clustered topology are inherently unfair due to zones of low (and high) contention inside the network. In such topologies, a backoff mechanism can help to improve fairness and the non-slotted protocols achieve a level of fairness higher than the slotted protocol.


At lower average backoff times, the non-slotted protocols start to organize the active connections in transmission schedules of higher spatial reuse. The downside of this self-organization property is that it might create unfairness, as connections that do not belong to such schedules are penalized and given less frequently access to the channel than the others. In general, the non-slotted protocols using a frozen backoff mechanism and a perfect collision avoidance



(a) Line topology, $\overline{c_w} = 512$. Except for the border effect, no organization is visible; the fairness is good (FI = 0.98) but the spatial reuse is low (the average number of packets sent per connection is 224 which corresponds to a spatial reuse $\sigma = 0.18$).



(b) Line topology, $\overline{c_w} = 2$. The organization is visible deep in the interior of the line topology and connections in the schedule of maximal spatial reuse emit more than others; the fairness is reduced (FI = 0.73) but the spatial reuse is high (the average number of packets sent is 387 which corresponds to a spatial reuse $\sigma = 0.32$).




(c) Circle topology, $\overline{c_w} = 2$. All connections belong to a transmission schedule of maximal spatial reuse. The densification process does not create any unfairness; the fairness is good (FI = 0.95) and the spatial reuse is high (the average number of packets sent per connection is 373 which corresponds to a spatial reuse $\sigma = 0.31$)


Fig. 13. Line and circle topologies. Average number of packets sent per connection by the (E, U, F) protocol. In the circle topology the last one-hop connection is between the right and left most nodes.

mechanism maintain the highest level of fairness. In multi-hop networks (like in single-hop networks) a frozen backoff dynamic improves the fairness of the non-slotted protocols with a uniform backoff mechanism. The use of a perfect collision avoidance mechanism is especially beneficial in the Poisson and the clustered topologies, where collisions tend to concentrate in zones of high node density creating some additional unfairness.

In the limit $\overline{c_w} \rightarrow 0$, or more precisely $T/\overline{c_w} \rightarrow \infty$, only the transmission schedules that maximize the spatial reuse have a non-trivial stationary probability (this has been shown for a subset of the non-slotted protocols in Section IV-B.3). In the case of the line topology there is only one transmission schedule that maximizes the spatial reuse and approximately one third of the connections belong to this schedule; the long-term fairness $FI \simeq 1/3$. In the grid topology, there are grosso modo four possible transmission schedules that maximize the spatial reuse and a little more than half the connections belong to these schedules; the long-term fairness $FI \simeq 1/2$. In the Poisson and the clustered topology the transmission



(a) $\overline{c_w} = 512$. Border connections send more than the others that have a similar throughput (FI = 0.89, $\sigma = 0.13$).



(b) $\overline{c_w} = 2$. Only the connections in schedules of maximal spatial reuse send a significant number of packets (FI = 0.29, $\sigma = 0.24$).

Fig. 14. Grid topology. Average number of packets sent per connection by the (E, U, F) protocol. The thicker and the darker a line representing a connection the higher the number of packet sent between its two end-nodes.

schedules that maximize the spatial reuse are not known and unfortunately we cannot gain any insight on the long-term fairness of the non-slotted protocols on these two topologies. It is interesting to compare the long-term fairness for $T/\overline{c_w} \rightarrow \infty$ with the short-term fairness obtained experimentally for the highest value of $T/\overline{c_w}$, namely for $\overline{c_w} = 2$. On the one hand, in the line topology, the short-term fairness value is well above the long-term fairness value. Figure 13(b) shows that at $\overline{c_w} = 2$ the transmission schedule that maximizes the spatial reuse on the line topology still coexists with other transmission schedules, leading to a better fairness than expected. On the other hand, in the grid topology, the short-term fairness value is well under the long-term fairness value. Figure 14(b) shows that in the grid topology a non-slotted protocol using an average backoff time of $\overline{c_w} = 2$ schedules almost exclusively the transmissions that are part of the optimal schedules. Once such a schedule is reached, the high level of coupling between its transmissions makes it difficult to switch to one of the other schedules of maximal spatial reuse. In fact, the results in terms of short-term fairness show that among the four schedules of

maximal spatial reuse, only two are present on average in the last 10s of the simulation.

So does spatial reuse always come at the price of fairness? Fortunately the answer to this critical question is no. Consider the line topology and join its two end points to make it a circle. Each connection then belongs to the same number of transmission schedules with one active connection, the same number of transmission schedules with two active connections, etc. Given that each transmission schedule with the same number of active connections has the same probability, it follows that at steady state all connections have the same probability to be active. Figure 13(c) confirms that in the circle topology the densification process does not create any unfairness (all connections have an equally high access to the channel). The network topology determines the transmission schedules of maximal spatial reuse that are selected by the non-slotted protocols at low average backoff values. Consequently, in a topology where each connection belongs to a schedule of maximal spatial reuse (such as the circle topology) a non-slotted protocol using a perfect collision avoidance mechanism can achieve both fairness and spatial reuse. However, in a topology where some connections do not belong to any transmission schedule of maximal spatial reuse (such as the line or the grid topologies) the non-slotted protocols cannot simultaneously achieve fairness and spatial reuse.

VII. CONCLUSION

In the design of MAC protocols for multi-hop ad hoc networks, the difficulty of maintaining an accurate synchronization of the node clocks across the network is often resented as a drawback. In this paper, we do not address the issue of clock synchronization. Rather, we show that in a decentralized network, global synchronization is not necessary to achieve a high level of spatial reuse. On the contrary, we demonstrate that a non-slotted (i.e., asynchronous) protocol can reach a higher spatial reuse than the corresponding slotted (i.e., synchronous) protocol, by using a simple backoff mechanism and a local but efficient collision avoidance mechanism. These two features introduce a sufficient degree of coupling between neighboring connections to progressively organize the transmissions in the network and eventually maximize the spatial reuse. We identify two conditions for this coupling to occur. First, the average backoff time must be significantly lower than the average exchange time to maintain a suitable level of correlation between the successive transmission schedules. Second, the level of collisions should remain low to allow the network to fully benefit from the self-organization properties of the non-slotted protocols. The well-known IEEE 802.11 is a good example of a non-slotted protocol based on a backoff mechanism which is unable to achieve a high level of spatial reuse due to a defective collision avoidance mechanism. We are currently investigating possible changes in the IEEE 802.11 protocol to improve its ability to organize transmissions in multi-hop ad hoc networks.

Finally, we observe that the price to pay for a high level of spatial reuse is often, but not always, a poor fairness.

In a topology where each one-hop connection belongs to a schedule of maximal spatial reuse, the non-slotted protocols can simultaneously achieve high levels of spatial reuse and fairness. However, in a topology where only a subset of the connections belongs to the schedules of maximal spatial reuse, a high level of spatial reuse can only be achieved at the cost of fairness. The network topology thus plays a major role in the fairness achieved by the non-slotted protocols and we suspect that a low level of mobility might help avoid completely unfair situations. In the future, we plan to further investigate the behavior of the non-slotted protocols under variable topologies and traffic loads.

REFERENCES

- [1] D. Bertsekas and R. Gallager, *Data networks*. Prentice-Hall, Inc., 1987.
- [2] *IEEE802.11, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std., Aug 1999.
- [3] M. Garey and D. Johnson, "Computers and intractability: A guide to the theory of NP-completeness," in *Freeman*, New York, 1983.
- [4] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *IEEE INFOCOM'03, San Francisco, CA*, April 2003.
- [5] H. S. Chhaya and S. Gupta, "Performance modeling of asynchronous data transfer methods of IEEE 802.11 mac protocol," *Wirel. Netw.*, vol. 3, no. 3, pp. 217–234, 1997.
- [6] Y. Wang and J. Garcia-Luna-Aceves, "Performance of collision avoidance protocols in single channel ad hoc networks," in *IEEE ICNP, (Paris, France)*, November 2002, pp. 68–78.
- [7] M. M. Carvalho and J. J. Garcia-Luna-Aceves, "A scalable model for channel access protocols in multihop ad hoc networks," in *MobiCom '04*. ACM Press, 2004, pp. 330–344.
- [8] F. Baccelli, B. Blaszczyzyn, and P. Muhlethaler, "A Spatial Reuse Aloha MAC Protocol for Multihop Wireless Mobile Networks," in *Allerton conference on communication, control, and computing*, 2003.
- [9] R. Rozovsky and P. Kumar, "SEEDEX: A MAC protocol for ad hoc networks," in *Proceedings of MobiHoc*, 2001.
- [10] E. W. Weisstein, "Rényi's Parking Constants." [Online]. Available: <http://mathworld.wolfram.com/RenyisParkingConstants.html>
- [11] J. Mackenzie, "Sequential filling of a line by intervals placed at random and its applications to linear adsorption," in *J. Chem. Phys.*, vol. 37, no. 4, 1962, pp. 723–728.
- [12] F. P. Kelly, "Loss networks," in *The Annals of Applied Probability*, vol. 1, no. 3, 1991, pp. 319–378.
- [13] K. Ramanan, A. Sengupta, I. Ziedins, and P. Mitra, "Markov random field models of multicasting in tree networks," *Adv. in App. Prob.*, vol. 34, no. 1, pp. 58–84, 2002.
- [14] S. Zachary and I. Ziedins, "Loss networks and markov random fields," *Journal of Applied Probability*, vol. 36, pp. 403–414, 1999.
- [15] Y. Baryshnikov, E. G. J. Coffman, and P. Jelenkovic, "Space filling and depletion," in *J. Appl. Probab.*, vol. 42, no. 3, 2004, pp. 691–702.
- [16] D. Y. Burman, J. P. Lehoczky, and Y. Lim, "Insensitivity of blocking probabilities in a circuit-switching network," in *Journal of Applied Probability*, vol. 21, no. 4, 1984, pp. 850–859.
- [17] Simulator [Online]. Available: <http://icapeople.epfl.ch/mdurvy/research.html>.
- [18] J. Talbot, G. Tarjus, and P. Viot, "The adsorption-desorption model and its application to vibrated granular materials," *Phys. Rev.*, vol. E 61, no. 5429, 2000.
- [19] C. L. Fullmer and J. J. Garcia-Luna-Aceves, "Solutions to hidden terminal problems in wireless networks," in *SIGCOMM '97*. ACM Press, 1997, pp. 39–49.
- [20] J. Deng and Z. Haas, "Dual Busy Tone Multiple Access (DBTMA): A New Medium Access Control for Packet Radio Networks," in *IEEE ICUPC'98, Florence*, October 1998.
- [21] F. Cali, M. Conti, and E. Gregori, "Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit," in *IEEE/ACM Trans. on Networking*, 8(6):785–799, December 2000.
- [22] R. Jain, *The Art of Computer Systems Performance Analysis*. John Wiley and Sons, Inc., 1991.