

# On the Fairness of Large CSMA Networks

Mathilde Durvy, Olivier Dousse, *Member, IEEE*, and Patrick Thiran, *Member, IEEE*

**Abstract**—We characterize the fairness of decentralized medium access control protocols based on CSMA/CA, in large multi-hop wireless networks. In particular, we show that the widely observed unfairness of these protocols in small network topologies does not always persist in large topologies.

In regular networks, this unfairness is essentially due to the unfair advantage of nodes at the border of the network, which have a restricted neighborhood and thus a higher probability to access the communication channel. In large 1D lattice networks these border effects do not propagate inside the network, and nodes sufficiently far away from the border have equal access to the channel; as a result the protocol is long-term fair. In 2D lattice networks, we observe a phase transition. If the access intensity of the protocol is small, the border effects remain local and the protocol behaves similarly as in one-dimensional networks. However, if the access intensity of the protocol is large enough, the border effects persist independently of the size of the network and the protocol is strongly unfair.

In irregular networks, the topology is inherently unfair. This unfairness increases with the access intensity of the protocol, but in a much smoother way than in regular two-dimensional networks.

Finally, in situations where the protocol is long-term fair, we provide a characterization of its short-term fairness.

**Index Terms**—Ad hoc Networks, Multi-hop, Medium Access Control, Performance Analysis

## I. INTRODUCTION

The goal of a medium access control protocol is to regulate the access of the network nodes to a shared resource, the communication channel. In large multi-hop networks, providing an efficient and fair access to the channel is a challenging task. Indeed the most commonly used medium access control protocol, IEEE 802.11, is known to be highly unfair. The tendency of IEEE 802.11 to starve some of the network links while giving a good access to the others has been observed experimentally on small testbed topologies [2]. It has also been investigated analytically in several papers [3], [4], [5], [6]. In particular, [3] and [6] show that even an *idealized CSMA/CA protocol* that retains only the essential features of IEEE 802.11 and does not suffer from different imperfections of the real protocol, such as overheads and collisions due to the RTS/CTS exchanges, exhibits this unfairness problem. We elaborate on the related work in Section II and detail the idealized CSMA/CA protocol in Section III.

Manuscript received 15 August 2008; revised 31 January 2009. A preliminary conference version of this paper appeared at Infocom'08 [1].

M. Durvy and P. Thiran are with EPFL, CH-1015 Lausanne, Switzerland (e-mail: mathilde.durvy@gmail.com; patrick.thiran@epfl.ch).

O. Dousse is with Nokia Research Center, CH-1015 Lausanne, Switzerland (e-mail: olivier.dousse@nokia.com).

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, and by the NetRefound FP6 FET Open Project of the European Union under grant number IST-034413.

In this paper we observe that in regular 1D (one-dimensional) or 2D (two-dimensional) networks, the unfairness of this idealized CSMA/CA protocol is due to *border effects*. The border effects are essentially local effects: the nodes at the boundaries of the network get an increased access to the channel, because they have fewer neighbors than nodes in the center. We expect therefore these border effects to rapidly become negligible in large network topologies. Contrary to our initial intuition, we will show that CSMA/CA protocols may cause these local effects to actually propagate throughout the whole network, creating a global effect. The network topology and the access intensity of the protocol play a key role in the propagation of these border effects. In Section IV we show that in large 1D lattice networks (*lines*) and for all finite access intensities, the border effects do not propagate inside the network. In contrast, we demonstrate that in large 2D lattice topologies (*grids*), a phase transition occurs. Under a certain access intensity, border effects fade away, whereas above a certain access intensity they propagate throughout the network, no matter how large it is. The proof techniques presented in this paper come from the area of statistical physics, and to the best of our knowledge, are for the first time applied in the context of multi-hop wireless networks.

Border effects play a key role in the fairness of the protocol. On the one hand, *when the border effects vanish*, the links inside the network get an equal access to the channel, and asymptotically, as the size of the network goes to infinity, *the protocol is long-term fair*. On the other hand, *when the border effects persist, the protocol is never long-term fair* and only a subset of the network links gets a good access to the channel. Settings where the protocol is long-term fair may not guarantee it to be short term fair, and we examine the latter properties in Section V. We evaluate by simulation the time needed for the idealized CSMA/CA protocol to reach a given (high) level of fairness. Taking the viewpoint of a given network link, we also compute how long on average the link waits to get access to the channel and how long it keeps the access to the channel once it has access.

## II. RELATED WORK

The work presented in this paper borrows tools from statistical physics. More precisely, it uses results from the theory of Markov random fields (for an introduction, see e.g. [7] or [8]). In particular, we apply two types of arguments: A percolation argument [9], [10] to prove that CSMA/CA based protocols can be long-term fair in large network topologies and the so-called *Peierls argument* to prove that this is not always the case. The simplest version of Peierls' argument can be found in [7]. Unfortunately, to characterize the behavior of

CSMA/CA protocols on large two-dimensional topologies, we have to resort to a more elaborate version of this argument. We adapt some of the techniques of [11], [12] to achieve this goal.

To model the network, we use the Markov chain formalism of [3] and [6]. This formalism was first proposed by [13], [14] to model the CSMA protocol and later extended to an idealized CSMA/CA protocol by [3]. The strength of this model is to preserve the dependence between nodes, which is typical of CSMA protocols. This model has also strong connections with Kelly's work [15] on loss networks as the stationary distribution of its Markov chain has the same form. Using this model, which is also known as the hard core model in statistical physics, [3] is probably the first to predict and explain the starvation phenomenon in multi-hop topologies. [6] extends the results of [3] and provides a closed form expression for the spatial reuse of CSMA/CA protocols in large one-dimensional networks. [16] extends the results further to consider the effect of the receiving and carrier sensing ranges on the model. Using Mean Field techniques, [17] also obtains asymptotically exact results in some scenarios with adaptive backoffs. Variations of the loss network model used in this paper have also been recently adopted to develop optimal distributed CSMA schemes making use solely of local information (such as the node backlogs, etc) [18], [19], [20].

Other recent models of CSMA/CA networks include [21], [22], [4], [23], and [5]. These models incorporate more features of the real protocols than the model used in this paper and typically lead to large systems of equations that must be numerically solved. As a consequence, these models offer limited insight into the behavior of CSMA/CA protocols in large multi-hop networks. However, some of these works do observe the starvation phenomenon. In particular, [4] attributes the unfairness of CSMA/CA protocols to a coordination problem inherent to carrier sensing protocols, while [5] concentrates on the role played by the minimum contention window in the starvation phenomenon.

Techniques from statistical physics have been applied first in [24] to analyze large-scale distributed resource sharing mechanisms. More recently, several works [25], [26], [27] showed the existence of a phase transition phenomenon in loss networks or hard core models. However, these papers only consider networks with relatively simple structures (i.e. tree or bipartite networks) and consequently their proof techniques cannot be used in our setting.

### III. MODEL

We consider an idealized CSMA/CA protocol in order to capture the essential features of CSMA/CA systems, and leave aside the effects due to the imperfections of the real protocol (we refer the reader to [28] for an overview of those effects in IEEE 802.11). This model has already been used in [3] as well as in [6] and [16].

We model the network using a graph where the vertices represent the nodes and the edges represent the links. A link can be in two states, active or idle. A link is active if there is a data transmission between its two end-nodes. Each link

has an *exclusion set* that contains all the links that cannot be active at the same time as itself. We assume that the relation of exclusion is symmetric, i.e., that if a link  $i$  has link  $j$  in its exclusion set, then link  $i$  is also in link  $j$ 's exclusion set. A link sees the communication channel as idle if and only if none of the links in its exclusion set is active.

In addition, we assume that each link maintains a *backoff timer*, which is initialized to a random value chosen according to an exponential *backoff distribution* with mean  $\lambda^{-1}$ . The timer of a link runs when the link sees the channel as idle; when the link senses the channel busy, it temporarily freezes its timer. A link becomes active when its backoff timer reaches zero, which is only possible if all the links in its exclusion set are idle. It stays active for an exponentially distributed time with mean  $\mu^{-1}$ . After each transmission, the link resets its backoff timer to a new random value.

This mechanism slightly differs from IEEE 802.11. In IEEE 802.11 each node maintains a unique backoff timer for all its links instead of a separate backoff timer per link. In addition, we make two assumptions. First, we assume that nodes can correctly assess the channel state. Second, we assume a continuous backoff distribution, instead of the discrete distribution implemented in the actual protocol, so that two timers have a zero probability to expire at the same time. These two assumptions imply that the idealized CSMA/CA protocol is collision-free.

We call *transmission pattern* the set of active links in the network at a given time. A transmission pattern must respect the constraint specified by the exclusion set of active links. We say that a transmission pattern is *maximal* (for a given network) if there is no other transmission pattern with a higher number of active links.

The dynamics of this idealized protocol can be described by a continuous time Markov chain, whose state space is the set of all possible transmission patterns in the network [6]. This Markov chain has a unique stationary distribution given by

$$\pi(x) = \frac{\rho^{n(x)}}{Z} \quad (1)$$

where  $x$  denotes a transmission pattern,  $n(x)$  is the number of active links in  $x$ ,  $\rho = \lambda/\mu$  is the access intensity, and  $Z = \sum_y \rho^{n(y)}$  is a normalizing constant.

In this paper, we consider essentially two different network settings: a line network composed of a sequence of equally spaced (by one space unit) nodes on a line where each pair of neighboring nodes is connected by a link, and a grid network where nodes are located on the vertices of a square lattice. In both networks, we assume that the exclusion set of a link contains all the links having an end-node connected to at least one of the link's end-nodes. A more convenient way to represent the exclusion set constraints is to associate to each active link a domain, *the exclusion domain*, that corresponds to the space occupied by an active link in the center of a maximal pattern. In the line network (Figure 1), the exclusion domain of a link is a line segment of length  $l = 3$  while in the grid network (Figure 2) it is a hexagon of area four. A transmission pattern on the grid (resp., line) network is thus





Fig. 1. Maximal pattern on the line network, the exclusion domain of a link is a line segment of length three.



(a) A maximal pattern with horizontal active links (b) A maximal pattern with vertical active links

Fig. 2. Maximal pattern on the grid network, the exclusion domain of a link is an hexagon of area four.

equivalent to a collection of non-overlapping hexagons (resp., line segments) on the 2D (resp., 1D) lattice.

#### IV. LONG-TERM FAIRNESS

In this section we look at the long-term fairness properties of CSMA/CA protocols in large networks. The long-term fairness of a protocol is the level of fairness achieved by the protocol after having run for a very long (in theory infinite) time. To assess this fairness, as the chain is ergodic, we look at the marginal stationary probabilities of individual links to be active, and we check whether they are all equal.

To do so, we need to study the structure of the stationary measure  $\pi$ . We use the concept of *contention graph*, whose vertices represent the links of the network and edges mutual exclusion of links. In other words, there is an edge between two vertices if their corresponding links are in one another's exclusion set. Note that this graph is different from the graph representing the network itself. We say that a vertex is in *state 1* if the corresponding link is active, and *state 0* otherwise. In this new graph, the exclusion rule implies that two adjacent vertices cannot be in state 1 simultaneously. A (maximal) transmission pattern on the original network topology is called a (maximal) independent set on the corresponding contention graph.

A central property of the stationary measure  $\pi$  is that it is a Markov field<sup>1</sup> on this graph. We describe briefly the properties of Markov fields which we need in the sections below. For more details, we refer the reader to [7]. The main property of Markov fields is that for each finite subset  $B$  of vertices, the conditional distribution of the states inside  $B$  given the state of all vertices outside  $B$  only depends on the state of the vertices in  $\partial B$ , where  $\partial B$  denotes all the vertices outside  $B$  that are adjacent to a vertex in  $B$ . More specifically, if  $x_B \in \{0, 1\}^{|B|}$  is a pattern in  $B$  and  $x_{\partial B} \in \{0, 1\}^{|\partial B|}$  a pattern on  $\partial B$ , where  $x_B$  and  $x_{\partial B}$  do not violate the exclusion rule, one can easily

<sup>1</sup>The term ‘‘Markov’’ here refers to the spatial structure of the measure. It has nothing to do with the fact that  $\pi$  is originally the stationary measure of a Markov chain (over time).





Fig. 3. Border effects on line networks,  $\rho = 155$ . The line network is represented on the x-axis. To each of its links, we associate a vertical bar that corresponds to the link access probability in the simulation.

verify from (1) that the measure conditioned on the state of vertices in  $\partial B$  is

$$\forall \text{ finite } B, \quad \pi(x_B | x_{\partial B}) = \frac{\rho^{n(x_B)} I(x_B, x_{\partial B})}{Z'}, \quad (2)$$

where  $Z'$  is an appropriate normalizing constant, and where  $I(x_B, x_{\partial B}) = 1$  if no adjacent vertices in  $x_B$  and  $x_{\partial B}$  are both in state 1 and  $I(x_B, x_{\partial B}) = 0$  otherwise. The advantage of (2) over (1) is that it can be used to define a measure over an infinite network. Indeed, (2) only refers to finite subsets  $B$ , whereas (1) addresses the whole graph, leading to an infinite normalization constant  $Z$  if the graph is infinite.

Let us now consider a measure  $\pi$  over the space of all possible transmission patterns that fulfills (2). A standard problem in statistical physics is to tell whether such a measure is unique. Essentially if there are several measures complying with (2), it means that the network can stabilize in several distinct stationary regimes, potentially with unfair distribution of link activities. This phenomenon corresponds to a phase transition of the Markov random field. On the other hand, if the measure is unique, as the equations in (2) are translation and reflection invariant, this measure cannot favor any link over the others. In this case, the network is long-term fair.

In the next two subsections, we look into the particular cases of the line and grid networks, as defined in Section III. In addition, we give a few general results on irregular networks at the end of the section.

##### A. The line network

We start with a few simulation results<sup>2</sup>, to give some insight into the behavior of the network. Figure 3(a) illustrates the stationary marginal probabilities of the links in a small four link topology. We observe that border effects can severely impact the fairness of the protocol. Yet in large networks border effects do not propagate inside the network, as shown in Figure 3(b).

*Theorem 1:* Consider an  $L$ -link line network. When  $L \rightarrow \infty$  and for all finite values of  $\rho$ , the probabilities of links in the center of the network to be active are asymptotically all equal and independent of the border condition.

*Proof:* We start by adding  $l - 1$  additional links at each extremity of our  $L$ -link network. These links form the *border* of our network. In the following, a *border condition* is the

<sup>2</sup>All the simulations in this paper were performed using a dedicated event driven simulator for the idealized CSMA/CA protocol (code available at [29])




Fig. 4. Illustration of the proof of Theorem 1 for  $l = 3$ . The  $L$  links are represented by black dots, and white dots the additional border links. Two possible border conditions are illustrated below: The probability of link  $i$  under the first one is clearly equal to the probability of link  $i + 1$  under the second one (solid ellipse). And yet, the probability of link  $i$  under any border condition remains asymptotically the same when  $1 \ll i \ll L$  (dashed ellipse). Thus, the probabilities of links  $i$  and  $i + 1$  are asymptotically equal.

activity pattern prescribed to the  $2(l - 1)$  border links. Setting all the border links to idle is equivalent to not adding border links at all; this corresponds to the usual border condition.

The contention graph of the line network consists in  $L + 2(l - 1)$  vertices representing the links (including the border links). An edge connects vertices  $i$  and  $j$  whenever  $|i - j| \leq l - 1$ . Let us define a site percolation process on this graph by declaring each vertex “open” with probability  $\rho/(1 + \rho)$  and “closed” with probability  $1/(1 + \rho)$ , independently of each other. An *open path* is a sequence of connected open vertices that are distinct. Let us consider two open vertices  $i$  and  $j$ , and divide the interval between them into  $\lfloor |i - j|/l \rfloor$  segments of length  $l$  (ignoring the remainder of the division). The probability that there is an open path between  $i$  and  $j$  is less than the probability that each of these segments contains at least one open vertex. We have thus


$$P(i \leftrightarrow j) \leq \left[ 1 - \left( \frac{1}{1 + \rho} \right)^l \right]^{\lfloor |i-j|/l \rfloor}$$

This probability tends to zero when the distance between  $i$  and  $j$  increases.


Consider now two different border conditions and look at the probability that a link is active under each border condition. Corollary 1 in [10] states that the difference between the two probabilities is less than the probability of finding an open path from the given link (or more exactly the corresponding vertex in the contention graph) to a border link. As we have shown in the previous paragraph, this latter probability tends to zero as the size of the network goes to infinity.

We have thus proved that asymptotically, the probability that a link be active, does not depend on the border condition. To finish the proof, we need to prove that these probabilities are also asymptotically equal to each other. Figure 4 illustrates this last part of the proof. Consider two border conditions, i.e. two activity patterns prescribed to the  $2(l - 1)$  border links, the second one being equal to the first one shifted by one unit. According to the above result, the probability of a central link  $i$  to be active under both border conditions are asymptotically equal. However, by translation, the probability of link  $i$  to be active under the first border condition is exactly equal to the probability of link  $i + 1$  to be active under the second border condition. Therefore we can conclude that for any  $1 \ll i \ll L$ , links  $i$  and  $i + 1$  have asymptotically equal probabilities under any border condition. ■

Theorem 1 implies that if we move sufficiently far away from the border of the network, border effects ultimately disappear, as illustrated by Figure 3(b). Consequently, in the



(a) Border effects. Fraction of links with a given channel access probability in two simulations, one with  $\rho = 50$  and one with  $\rho = 155$ . For  $\rho = 50$  all links have similar access probabilities; the protocol has a high level of fairness. However, for  $\rho = 155$ , a small fraction of the links monopolize the access to the channel; the protocol has a low level of fairness.



(b) Long-term fairness of the idealized CSMA/CA protocol. Observe the sharp drop in the fairness of the protocol for values of  $\rho$  between 60 and 90, from values close to 1 to values slightly above  $1/8$ .

Fig. 5. Fairness on a  $34 \times 34$  node grid network

limit when the number of links goes to infinity, *the idealized CSMA/CA protocol is long-term fair and all links have the same probability to access the channel.*

### B. The grid network

We start again with illustrative simulation results: in large grids we observe that for low values of  $\rho$  the protocol provides a fair access to the links inside the network, but not for large values of  $\rho$ .

Figure 5(a) shows the histogram of the link access probabilities for two values of  $\rho$ . At  $\rho = 50$  all links have a very similar access probability to the channel. In contrast, at  $\rho = 155$ , two categories of links clearly appear: approximately  $1/8$  of the links have a very high channel access probability, whereas the remaining  $7/8$  of the links have almost no access to the channel.

In this section we show that in a grid network, the idealized CSMA/CA protocol exhibits a phase transition. Indeed we show that for low values of  $\rho$ , the protocol becomes long-term fair as the number of links goes to infinity. It has therefore a

similar behavior as on the line network. In contrast, for large values of  $\rho$ , the effects of the border always propagate inside the network independently of its size, and the protocol is not long-term fair.

To assess the fairness of the protocol we use *Jain's Fairness Index* [30].

**Definition 1:** [Jain's Fairness Index] Denote by  $p(j)$  the probability that a link  $j$  is active. The link fairness index FI of the protocol is

$$\text{FI} = \frac{\left(\sum_j p(j)\right)^2}{L \sum_j p(j)^2}. \quad (3)$$

The maximum fairness index is  $\text{FI} = 1$ . It corresponds to a network where all links access the channel equally. Yet, if only  $k$  links have an equal access to the channel and the remaining links have no access to the channel, the fairness index is  $k/L$ .

Figure 5(b) illustrates the phase transition using Jain's fairness index.

#### 1) Sub-critical regime:

**Theorem 2:** Consider an  $L$ -link grid network. There is some  $c_1 > 0$  such that when  $\rho < c_1$  and  $L \rightarrow \infty$ , the probabilities of links in the center of the network to be active are asymptotically all equal and independent of the border condition.

*Proof:* This theorem is proved in the same way as Theorem 1. We denote by  $G$  the contention graph of the network and define a site percolation process by declaring each vertex open with probability  $\rho/(1+\rho)$  independently. This site percolation process is less trivial than in one dimension: By Theorem 1.33 in [31], we know that it admits a percolation threshold  $p_c(G)$  that is finite and strictly positive. Thus, if  $\rho < p_c(G)/(1-p_c(G)) := c_1$ , the probability to find an open path connecting a given vertex to the border of the network tends to zero when the border moves away from the vertex. Therefore the same argument as in Theorem 1 can be used to conclude that if  $\rho < c_1$ , the probability of a link to be active is asymptotically independent of the border condition.

In a second step, we show that adjacent links have asymptotically equal probabilities. To do this, as in Theorem 1, we consider different border conditions that are equivalent through a transformation (shift or rotation) of the network. This allows to show that links that are equivalent through the same transformation have asymptotically equal probabilities. The details of this step are omitted for lack of space. ■

As a consequence of this theorem, if  $\rho < c_1$ , the fairness index of the network tends to 1 as  $L \rightarrow \infty$ .

#### 2) Super-critical regime:

**Theorem 3:** Consider an  $L$ -link grid network. For any  $L$ , any pair of adjacent links  $i$  and  $j$ , and any  $\delta < 1$ , there exists  $c_2 < \infty$  such that if  $\rho > c_2$ , the following two statements hold:

- there exist two distinct border conditions such that the difference of the probabilities of link  $i$  to be active under each border condition is larger than  $\delta$ ,
- there exists a border condition under which the difference between the probabilities of link  $i$  and link  $j$  to be active is larger than  $\delta$ .

In particular, this theorem implies that for sufficiently large values of  $\rho$ , the fairness index remains uniformly bounded away from 1 when  $L$  tends to infinity. Loosely speaking, the theorem means that the border effects do *not* vanish in the center of the network.

To prove the theorem, we use Peierls' argument [7]. Peierls introduced this argument in 1936 to show that the two-dimensional Ising model exhibits a phase transition. It is now a classical argument in statistical physics. The remainder of this section is devoted to this proof, and is structured as follows: we start with a few definitions; next, we present the outline of Peierls' argument; finally, we present the main steps necessary to apply the method of [12] to our setting.

#### i) Contours and their properties

In an infinite grid network, there are eight maximal transmission patterns (or equivalently eight maximal independent sets in the contention graph), four with horizontal active links and four with vertical active links. The four horizontal (resp. vertical) transmission patterns can be obtained from each other by translation. Each hexagon or link<sup>3</sup> belongs to exactly one of these eight patterns, hereafter called the *phase* of the hexagon.

Peierls' argument relies on the notion of *contour*. For the following definition it is helpful to consider the example of Figure 6(a).

**Definition 2:** Consider an arbitrary transmission pattern (i.e., a set of non-overlapping hexagons). An element (or contour) is either an area not covered by hexagons, or a common edge of two hexagons of two different phases. Two elements are connected if they share at least one vertex. A contour is a set of connected elements, which is maximal.

Let  $C$  be a contour.  $C$  is delimited by an outer polygon and possibly several inner polygons (for example, in Figure 6(a) there are two inner polygons). All edges of the outer polygon necessarily belong to hexagons of the same phase, called the *phase of the contour*  $C$ . The interior of  $C$  is the area contained inside its outer polygon. This area can be tiled with hexagons of the phase of  $C$  (Figure 6(b)). If the interior of such an hexagon intersects a contour element, we say that this hexagon *belongs* to the contour or that the contour *contains* the hexagon. The length of a contour is defined as the number of hexagons belonging to the contour and is denoted by  $L(C)$ .

**Lemma 1:** The number of contours of length  $k$  containing a given hexagon is upper bounded by  $q^k$ , with  $q = 15 \cdot 2^{16}$ .

*Proof:* See [1] ■

Before we begin with Peierls' argument we need one last notion, the notion of  $C$ -inner contour.


**Definition 3:** A  $C$ -inner contour is a contour inside  $C$  but not inside any other contour.

Figure 7(a) gives an example of a transmission pattern with a contour  $C$  and two  $C$ -inner contours.

#### ii) Outline of Peierls' argument

Consider two adjacent links  $i$  and  $j$  and denote by  $A$  the phase of link  $j$ . As adjacent links cannot be in the same phase, link  $i$  is not in phase  $A$ . Now assume that the border of the network is fully packed with hexagons of phase  $A$ . Under this border

<sup>3</sup>As each link corresponds to one hexagon and vice-versa, we use interchangeably link and hexagon



(a) Transmission pattern with a single contour  $C$ . The contour  $C$  corresponds to the white area on the figure but also contains the two encircled edges.

(b) The interior of the contour is tiled with hexagons in the phase of  $C$ . The contour contains  $L(C) = 33$  of these hexagons.

(c) Contour removal operation. The number of additional hexagons (compared to the original transmission pattern in (a)) is 16.

Fig. 6. Removal of a single contour  $C$ .

condition, any pattern where link  $i$  is active contains at least one contour around  $i$ . An upper bound to the probability that  $i$  is active is thus the probability of finding an outer contour (i.e., a contour that is not inside another contour and thus necessarily in phase  $A$ ) around link  $i$ .

Let  $y$  denote an arbitrary transmission pattern on the network that is compatible with the border condition. The probability of a given contour  $C$  is

$$P(C) = \frac{\sum_{y \ni C} \rho^{n(y)}}{\sum_y \rho^{n(y)}}, \quad (4)$$

where  $y \ni C$  denotes a pattern  $y$  that contains contour  $C$ . Assume that we can find an upper bound on  $P(C)$  of the form

$$P(C) < \rho^{-\varepsilon L(C)} \quad (5)$$

where  $\varepsilon > 0$ . Denoting by  $p(i|A)$  the probability that link  $i$  is active under our border condition in phase  $A$ , we can then write

$$p(i|A) = \sum_{C \ni i} P(C) < \sum_{C \ni i} \rho^{-\varepsilon L(C)} < \sum_l l^2 q^l \rho^{-\varepsilon l} := f(\rho)$$

where  $C \ni i$  is a contour that surrounds  $i$ , where  $l^2$  is an upper bound on the number of possible starting hexagons for a contour of length  $l$ , and hence where  $l^2 q^l$  is an upper bound on the number of contours  $C$  of length  $l$  surrounding link  $i$ . We observe that  $f(\rho)$  decreases to 0 when  $\rho \rightarrow \infty$ .

This upper bound actually holds for the probability of any link that is *not* in phase  $A$ . In particular, we can apply it to the 20 links that contend with link  $j$  (link  $i$  being one of them). Applying a union bound, we find that the probability that they are all inactive is at least  $1 - 20f(\rho)$ , and thus

$$p(j|A) \geq (1 - 20f(\rho)) \frac{\rho}{1 + \rho}.$$

The above lower bound is an increasing function of  $\rho$  and tends to one when  $\rho$  tends to infinity. Therefore, there is a value  $c_2$  such that the upper bound on  $p(i|A)$  and the lower bound on  $p(j|A)$  are separated by more than  $\delta$ , which proves the second

claim of the theorem. The first claim can be verified similarly by considering the border condition where all hexagons are in the same phase as link  $i$ , say  $B$ : we get a lower bound  $p(i|B) \geq 1 - f(\rho)$  and the first claim follows.

iii) *Upper bound on  $P(C)$*

The main difficulty of the proof is to obtain the upper bound (5) on the probability  $P(C)$  of a given outer contour  $C$ . In principle, the idea is quite simple. Let  $\bar{y}$  denote an arbitrary transmission pattern with contour  $C$ , Equation (4) can be rewritten as

$$P(C) = \frac{\sum_{\bar{y}} \rho^{n(\bar{y})}}{\sum_y \rho^{n(y)}}.$$

If to each  $\bar{y}$  we can map another transmission pattern  $y'$  with  $\varepsilon L(C)$  more hexagons and if this mapping is one to one, then, to each term in the numerator we have associated a term in the denominator such that  $n(y') - n(\bar{y}) = \varepsilon L(C)$ . Since reducing the number of terms at the denominator can only increase the fraction, we have

$$P(C) < \frac{\sum_{\bar{y}} \rho^{n(\bar{y})}}{\sum_{y'} \rho^{n(y')}} = \frac{\sum_{\bar{y}} \rho^{n(\bar{y})}}{\sum_{\bar{y}} \rho^{n(\bar{y}) + \varepsilon L(C)}} = \rho^{-\varepsilon L(C)},$$

which is exactly the upper bound (5) with  $\varepsilon = \varepsilon$ .

a) *Contour removal and transformation*: The challenge is thus to find the one-to-one mapping  $\bar{y} \leftrightarrow y'$ . The mapping is the simplest when the original transmission pattern  $\bar{y}$  includes only the outer contour  $C$  and no  $C$ -inner contour. Figures 6(a) (for  $\bar{y}$ ) and 6(c) (for  $y'$ ) illustrate this mapping. To obtain the resulting pattern  $y'$ , we start from  $\bar{y}$ , and we fill the interior of  $C$  with hexagons of phase  $A$ , which is the phase of the outer contour  $C$ . In  $y'$ , the contour  $C$  is no longer visible; this is why the mapping between  $\bar{y}$  and  $y'$  is often called a *contour removal*. It is easy to see that the number of additional hexagons in  $y'$  is proportional to  $L(C)$  and is at least equal to  $\varepsilon L(C)$  where  $\varepsilon$  is a lower bound on the average uncovered area per unit of contour length.

In general, however, the contour  $C$  contains  $C$ -inner contours (Figure 7(a)), which can be divided in three categories

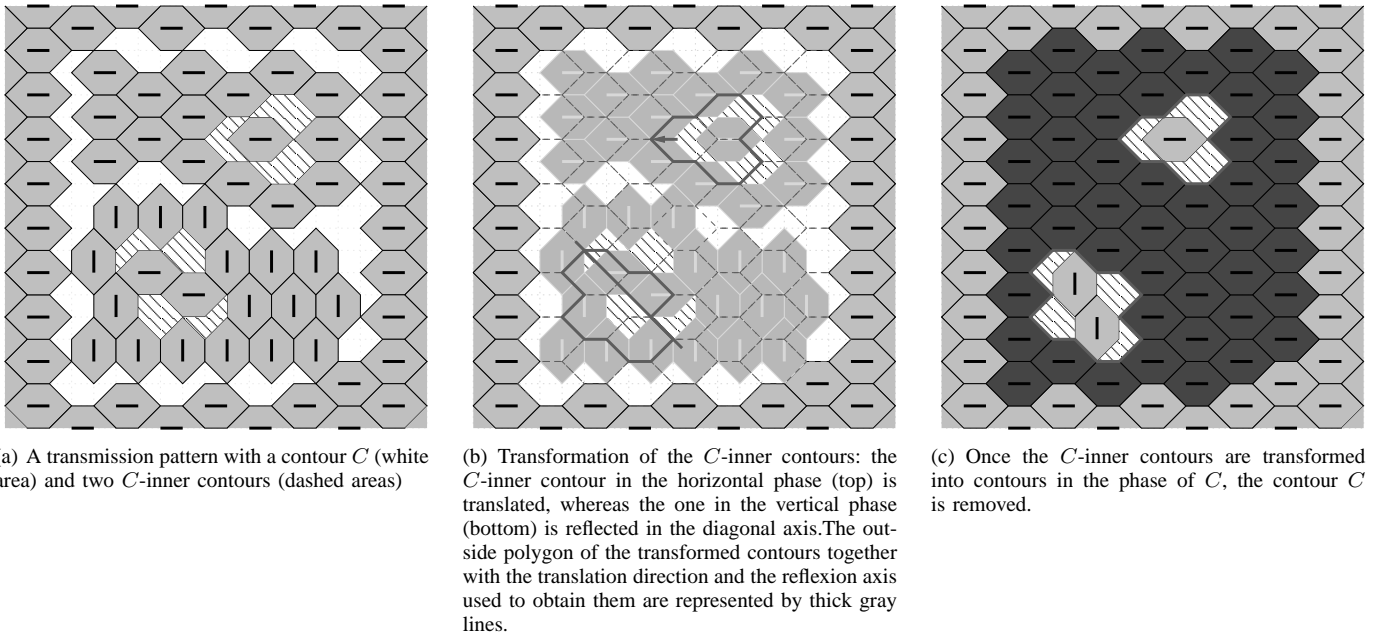


Fig. 7. Contour  $C$  with two  $C$ -inner contours. The two  $C$ -inner contours are transformed and the contour  $C$  is removed.

(we assume without loss of generality that the phase of  $C$  is horizontal):

- 1)  $C$ -inner contours in the same phase  $A$  as  $C$
- 2)  $C$ -inner contours in a horizontal phase different from the phase  $A$  of  $C$
- 3)  $C$ -inner contours in a vertical phase.

To obtain  $y'$  from  $\bar{y}$ , we first *transform* the  $C$ -inner contours into contours of phase  $A$ . The  $C$ -inner contours in the first category are already in phase  $A$  and thus do not need to be transformed. The  $C$ -inner contours in the second category are translated to become contours of phase  $A$ .<sup>4</sup> Finally, to transform the  $C$ -inner contours in the third category into contours of phase  $A$  we reflect them in a diagonal axis of the grid.<sup>5</sup> Figure 7(b) illustrates the details of this mapping on a small example.

Unfortunately, the reflected contours might touch<sup>6</sup> the contour  $C$ . The contours touching  $C$  are removed together with  $C$ . Figure 8 gives an example of such a situation. Moreover, the reflected contours might touch each other, in which case they are reflected in a common axis, but their joint reflection might also touch  $C$ , in which case the contours need to be removed, and so on.

Heilmann [12] describes an iterative process to determine which contours should be removed and which contours should be transformed: the same process can be applied in our setting. The result of this process is a set  $K$  of contours to transform and a set  $T$  of contours to remove. To obtain  $y'$  from  $\bar{y}$ , we first transform all the contours in  $K$  and we then remove all the contours in  $T$ . In  $y'$ , the interior of  $C$  (as well as the interior

of all the other contours in  $T$ ) is thus filled with hexagons of phase  $A$ , whereas the interior of the transformed contours is left unchanged (see Figure 7(c)).

*b) Enumeration of the contours to be removed:* In the case where the set  $T$  of contours to remove contains at least one  $C$ -inner contour, the mapping between  $\bar{y}$  and  $y'$  may no longer be one to one, unless we specify the set  $T$ . Indeed, from  $y'$  it is possible to reverse the mapping and recover the  $C$ -inner contours of  $\bar{y}$  that have been transformed but not the ones that have been removed. We have therefore

$$P(C) < \sum_{T \ni C} \rho^{-\epsilon L(T)} \quad (6)$$

where  $L(T)$  is the sum of the lengths of the contours in set  $T$ .


In the iterative process of [12], the only contour initially in the set  $T$  is the outer contour  $C$ .  $C$ -inner contours are then added to  $T$  if and only if their reflection (or their joint reflection) touches a contour already in  $T$ . To obtain an upper bound of the type (5), we need to enumerate all possible sets  $T$ . We do not go into the details of this lengthy and rather complicated enumeration which can be found in [12]. Here we restrict ourselves to the enumeration of the sets  $\tilde{T} \subseteq T$  with two contours: the outer contour  $C$  and a  $C$ -inner contour of length  $k$  whose reflection touches  $C$  at a specific contact point. This is sufficient to put our setting in gear with the recursive procedure in [12]: Using the contribution of the set  $\tilde{T}$  to the sum in (6) as an initial seed for this recursive procedure will allow us to directly apply the bound of [12] for  $P(C)$ .

Before we start to enumerate the sets  $\tilde{T}$ , we need to make the notion of contact point precise. Let us call  $D$  the  $C$ -inner contour in  $\tilde{T}$ . The contour  $C$  (resp.,  $D$ ) is contained in  $L(C)$  (resp., in  $L(D) = k$ ) hexagons in the phase of  $C$  (resp., of  $D$ ). A necessary condition for contour  $C$  and the reflection of  $D$  to touch each other is that two of their hexagons touch each

<sup>4</sup>We always choose the shortest possible translation, if there are two possibilities, we favor moving up

<sup>5</sup>The reflection axis is the (top-left-bottom-right) diagonal that splits the interior of the contour in two areas as equal as possible (if several reflection axes satisfy the above condition, we favor the upper most one).

<sup>6</sup>We say that two contours touch each other if two of their hexagons (one belonging to each contour) overlap or share a common edge.



(a) A transmission pattern with a contour  $C$  (white area) and two  $C$ -inner contours (dashed areas). The  $C$ -inner contour in the vertical phase contains one more hexagon than in Figure 7(a).

(b) Transformation of the  $C$ -inner contours. The reflection of the vertical  $C$ -inner touches the contour  $C$ .

(c) The resulting transmission pattern. The horizontal  $C$ -inner contour has been transformed and the vertical  $C$ -inner contour has been removed together with the outer contour  $C$ .

Fig. 8. Contour  $C$  with two  $C$ -inner contours. One  $C$ -inner contour is transformed while the other one is removed.

other (i.e., overlap or share a common edge). A contact point between a contour  $C$  and a contour  $D$  can thus be specified as two hexagons (one from  $C$  and one from the reflection of  $D$ ) that touch each other. Given the first hexagon, there are exactly 33 possible relative positions for a second hexagon. Therefore, per hexagon of  $C$ , or equivalently per unit of length of  $C$ , there are 33 possible points of contact with another contour. The enumeration method in [12] requires that there is only one contact point per unit of contour length. To satisfy this constraint, we redefine the length of a contour as 33 times the number of its hexagons. We use a symbol prime to denote this new length, e.g.,  $L'(C) = 33L(C)$ .

Using Lemma 1, we find that an upper bound on the number of contours of length  $k'$  touching  $C$  at a specific contact point is  $q^{k'/33}$  and their contribution to  $P(C)$  is less than  $(q\rho^{-\epsilon})^{k'/33}$ .

We then use the same recursion as in [12] to compute the contribution of all sets  $T$ . Adapting relations (3.27) and (3.28) of [12], we finally obtain that for  $n = \frac{1}{33}(\epsilon \log \rho - \log q) > f_0$  where  $f_0$  is a constant

$$P(C) < \rho^{-L(C)(\epsilon - 33B(\log \rho)^{-1}66^{-n})}$$

with  $B = \frac{1}{256}(1 + \frac{66}{f_0 - 1})$ . For  $\rho$  large enough,  $\epsilon = \epsilon - 33B(\log \rho)^{-1}66^{-n}$  is positive. Consequently, we can use this upper bound on  $P(C)$  to conclude Peierls' argument as presented in the outline.

As a result, in the *super-critical regime the idealized CSMA/CA protocol is not long-term fair even if the size of the network goes to infinity*. The effects of the border do not disappear, and a positive fraction of the links get better access to the channel.

### C. Irregular networks

Up to now we only considered regular topologies. In practice, these topologies prevail when the position of the network nodes can be controlled and is fixed. This is typically the case for indoor networks or more generally for sensor networks. However, in situations where a careful deployment of the




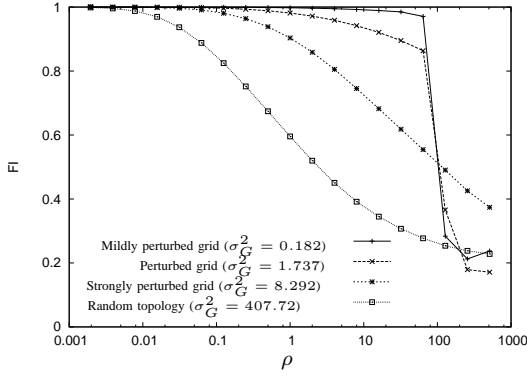
Fig. 9. Example of a grid topology with perturbations

network nodes is not possible, the network topology is often irregular.

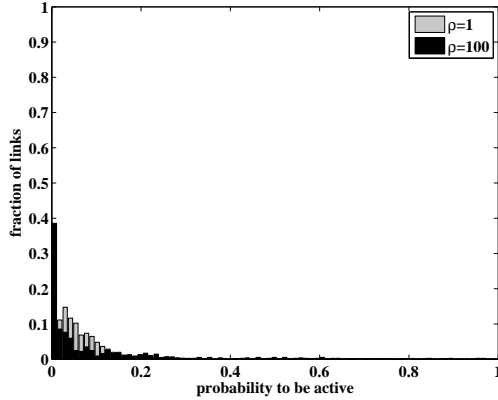
Figure 10(a) shows simulation results on the fairness index of various irregular topologies. The topology of the three first simulations have been obtained by perturbing the alignment of the nodes in a grid network (see Figure 9). The last topology was obtained by scattering 1000 nodes uniformly on a square area. For the simulations, we considered only the largest connected component of nodes (which contains about 600 nodes). Figure 10(b) illustrates the histogram of the link access probabilities in the last case.

In contrast to the grid, irregular topologies are intrinsically unfair, as the number of contending links varies not only at the boundary but everywhere in the network. Let us first evaluate how the irregularity of the topology affects the fairness index FI (3) when  $\rho$  is very small. More precisely, let us consider a connected contention graph  $G$  of  $L$  nodes (corresponding





(a) Fairness Index computed at the end of a very long simulation of the idealized CSMA/CA protocol in irregular networks.



(b) This figure represents the fraction of links with a given channel access probability in two simulations on the random topology, one with  $\rho = 1$  and one with  $\rho = 100$ .

Fig. 10. Fairness of the idealized CSMA/CA protocol in irregular networks.

therefore to a topology of  $L$  links as before). Let  $\delta_j$  denote the degree of Node  $j$  on this graph (i.e., the number of links in the exclusion domain of link  $j$ ), and let

$$\sigma_G^2 = \frac{1}{L} \sum_{j=1}^L \delta_j^2 - \left( \frac{1}{L} \sum_{j=1}^L \delta_j \right)^2.$$

We note that  $\sigma_G^2$  is the sample variance of the node degree of the contention graph  $G$ . The following lemma shows how FI depends on  $\sigma_G^2$ , and thus on the irregularity of the topology, for very small values of  $\rho$ .

*Lemma 2:* If the contention graph  $G$  is countable, locally finite and connected, then

$$FI(\rho) = 1 - \sigma_G^2 \rho^2 + O(\rho^3). \quad (7)$$

*Proof:* Let  $N(i, j)$  be the number of transmission patterns (i.e., independent sets on  $G$ ), which contain exactly  $i$  active links (node of  $G$ ), among which link  $j$  (node  $j$  in  $G$ ). The probability  $p(j)$  that link  $j$  is active is simply given by summing (1) on all patterns  $x$  that contain  $i$  as active links, which we can write as

$$p(j) = \frac{1}{Z} \sum_{i=1}^L N(i, j) \rho^i = \frac{1}{Z} \left( \rho + \sum_{i=2}^L N(i, j) \rho^i \right) \quad (8)$$

because  $N(1, j) = 1$  as there is only one transmission pattern having only  $j$  as active link.

Plugging (8) in (3) yields, after a few manipulations and keeping only terms up to the second order in  $\rho$ ,

$$FI(\rho) = 1 - \frac{\rho^2}{L^2} \left[ L \sum_{j=1}^L N^2(2, j) - \left( \sum_{j=1}^L N(2, j) \right)^2 \right] + O(\rho^3).$$

Now,  $N(2, j) = L - 1 - \delta_j$ , because any pattern containing  $j$  as active link can only have another link among the  $L - 1 - \delta_j$  links that are not in the exclusion domain of  $j$ . Replacing  $N(2, j)$  by  $L - 1 - \delta_j$  in this expression leads to (7). ■

Lemma 2 shows that for small values of  $\rho$ , the variability of the node degree of the contention graph is responsible for a “topological” unfairness. When  $\rho$  increases, the coupling effect between contending nodes adds another source of unfairness, like in regular topologies, but the effect is smoother because of the irregularity of the graph. The maximal independent sets are clearly much more difficult to compute in irregular topologies, and they will no longer have the same isomorphic properties as in the 2D regular lattice. It is therefore difficult to dissociate the unfairness due to the topology, from the additional unfairness enforced by the combined effects of backoff and carrier sense mechanisms. In addition, phase transitions on hard core models in general topologies may be quite complex and non monotonic with respect to parameter  $\rho$  [32].


We can however compute a range of low values of the access intensity that ensure the network to be in the sub-critical regime. Indeed, we can observe that the first part of the proof of Theorem 2 only requires the contention graph  $G$  to be countable, locally finite and connected. Therefore, for any network whose contention graph fulfills these properties, we can conclude that if  $\rho < p_c(G)/(1 - p_c(G))$ , where  $p_c(G)$  is the independent site percolation probability of the graph  $G$ , the link probabilities are asymptotically independent of the border conditions.

We note that the “topological” unfairness, which is absent from regular topologies, depends on a local feature of the contention graph  $G$  (the variance  $\sigma_G^2$  of the node degree) at very small values of  $\rho$ , whereas the range of access intensities that safely avoid the onset of the long range unfair effects due the “coupling” mechanisms, depends on a global feature of  $G$  (the site percolation threshold  $p_c(G)$ ).


## V. SHORT-TERM FAIRNESS

In the previous section, we saw that in many cases, CSMA/CA protocols give an equal share of the channel to each link in the center of the network on the long term. However, this does not guarantee that certain links are not starved for long periods (and then keep the channel for an equally long period, preventing the neighboring links to be activated).

In this section we attempt to characterize the short-term fairness of these protocols in settings where they are long-term fair. Short-term fairness is important for delay sensitive applications such as real-time audio and video. Several studies



(a) 2000 node line network.



(b) 34x34 node grid network. The value of the short-term fairness horizon explodes just before the phase transition threshold

Fig. 11. Short-term fairness horizon of the idealized CSMA/CA protocol. We measure the average time needed by the protocol to reach a fairness index of 0.9. The average exchange time  $\mu^{-1}$  is equal to 1 time unit (and thus  $\rho = \lambda$ ).

(see for example [33]) also show that the lack of short-term fairness at the MAC layer can severely impact the performance of a reliable transport protocol such as TCP.

We consider two measures of short-term fairness. A network-wide measure, where we evaluate the time necessary for the protocol to reach a high level of fairness, and a link-centric measure, where we compute how long on average a link has to wait to get access to the channel.

#### A. Short-term Fairness Horizon

We begin our study of CSMA/CA short-term fairness by a few illustrative simulation results. We consider only topologies where the idealized protocol is long-term fair. Figure 11(a) shows how long on average we need to operate the idealized CSMA/CA protocol to reach a fairness index of 0.9 in a large line network. As expected, this time, which is also called the *short-term fairness horizon* [33] of the protocol, increases progressively with the value of  $\rho$ . Figure 11(b) shows that in the grid network the short-term fairness horizon of the protocol increases progressively with  $\rho$  and suddenly diverges at a finite value of  $\rho$ , which marks the phase transition described in the previous section. Indeed, above the critical value of  $\rho$ , the protocol is not fair anymore. Close below the threshold, the




Fig. 12. Transmission pattern on the line network with a hole. The succession of the 3 figures represent the jump of a hole by  $l = 3$  space units (event (ii)).

protocol is still long-term fair but one has to wait a very long time to have a fair channel access among the network links.

#### B. Average link access and waiting times

We now characterize the short-term fairness of CSMA/CA systems from a link-centric viewpoint. Our goal is to answer two questions: (i) How long on average does a link wait to get access to the channel, and (ii) How long on average does a link keep the access to the channel? Note that we say that a link has access to the channel if it is active or if it does not see the channel as busy.

In the case of the line network, when the access intensity  $\rho$  is finite but sufficiently large, the link access and waiting time can be estimated by describing the dynamics of *holes*, where a hole is defined as a unit space not covered by a line segment (of length  $l$ ) corresponding to an active link (see Figure 12). Holes separate links in different phases and are the exact one-dimensional equivalents of contours in the grid network. When  $L \rightarrow \infty$  the density of holes is [6]:

$$\eta(\rho) \sim \frac{1}{1 + l\rho y_1^{l-1}} \quad (9)$$

where  $y_1$  is the real root of  $1 - y - \rho y^l$  closest to the origin.

The following proposition gives analytical expressions for the average link access and waiting time based on the dynamic motion of holes along the line. Although these expressions rely on several approximations, Figure 13 shows an extremely good fit with the simulation results.

*Proposition 1:* Consider an  $L$ -link line network with  $\rho$  large. When  $L \rightarrow \infty$ , a link waits to get access to the channel for on average

$$\frac{1}{\mu} \left( \frac{l}{1 - \eta(\rho)} - 1 \right) \left( 1 + \frac{1 - \eta(\rho)}{l\eta(\rho)} \right) \quad (10)$$

time units, and keeps the access to the channel for on average

$$\frac{1}{\mu} \left( 1 + \frac{1 - \eta(\rho)}{l\eta(\rho)} \right) \quad (11)$$

time units.

*Derivation of formulae (10) and (11):* In this computation, we assume that  $\rho$  is large so that the density of holes in the network is small. Moreover, we assume a fixed density of holes equal to  $\eta(\rho)$ . We can describe the dynamics of the network by the motion of the holes along the line.

A link switches from active to idle state at rate  $\mu$ . When a link next to a hole becomes idle (Figure 12(b)), two equally likely events may occur: (i) the same link becomes active again

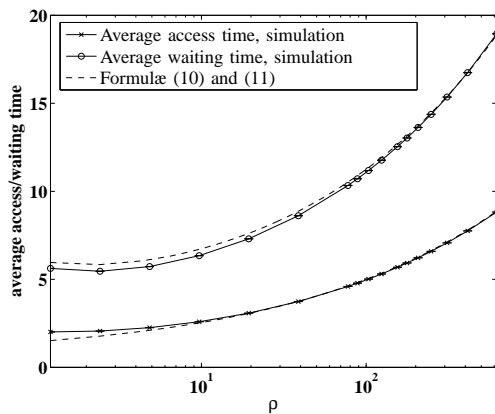


Fig. 13. Average time a link waits to get access (resp. keeps the access) to the channel in a 2000 node line network with  $\mu = 1$ . Comparison between the simulation results obtained for the idealized CSMA/CA protocol and the values predicted by formulae (10) and (11).

(we are back to the pattern in Figure 12(a)) or (ii) another link becomes active (Figure 12(c)). In terms of holes, event (i) means that the hole does not move, whereas event (ii) means that the hole jumps by  $l$  space units; we call the latter a *useful jump*. A hole makes therefore a useful jump at rate  $\mu/2$ , where the factor  $1/2$  is the probability of event (ii).

Denote by  $n = L(1 - \eta(\rho))/l$  the number of active links and by  $k = L\eta(\rho)$  the number of holes in the network. The small number of holes makes it reasonable to assume that they move independently of each other at rate  $\mu$  (i.e.,  $\mu/2$  in a given direction). In particular, a hole which is next to another hole might jump above it. We call this one space unit jump event (iii). Event (iii), similarly to event (i), does not change the transmission pattern.

The  $k$  holes make a jump at a total rate of  $k\mu$ . However, only events of type (ii), holes which make a useful jump over an active link, modify the transmission pattern, as we just saw. The probability of a useful jump is the ratio  $p = n/(k-1+n)$ ,<sup>7</sup> and consequently useful jumps happen at rate  $k\mu p$ .

Given that we have  $n$  active links in the network, the rate of useful jump per active link is thus  $k\mu p/n$ . This implies that a link gets access to the channel for on average  $n/(k\mu p)$  time units. Replacing  $n$  by  $L(1 - \eta(\rho))/l$  and  $k$  by  $L\eta(\rho)$  gives (11).

Similarly, the rate of useful jumps per inactive link is  $(k\mu p)/(k + n(l - 1))$ , from which we easily get (10).  $\square$

This reasoning, however, does not apply to the grid network: as we have seen in the previous section, phases are delimited by contours, not by holes. The dynamics of contours is too complicated to allow a simple method as the one above to give reasonable approximations of the waiting times. In any case, such a characterization would not make sense in the super-critical regime where we have two categories of links with very different access and waiting times.

<sup>7</sup>This ratio stems from the fact that a useful jump happens when the hole jumps over an active link and *not* over another hole, since you have  $n$  active links and  $k - 1$  other holes you obtain  $n/(k - 1 + n)$ .

## VI. CONCLUSION

The combination of the random backoff and carrier sensing mechanisms, two key features of CSMA/CA protocols, creates a coupling between nodes, which prompts a maximal transmission pattern to spontaneously emerge in the network, but at the cost of the starvation of links that do not belong to this maximal pattern. In regular topologies, only nodes at the border are favored and yet their advantage can affect the nodes inside the network, to an extent that depends on the parameters of the protocol (which determine the strength of the coupling) and on the dimension of the network. In some cases (grid when the access intensity is large enough), we have shown that these border effects propagate arbitrarily far inside the network without vanishing; this leads to the complete starvation of some links in the network. In other cases (line; grid when the access intensity is small enough), the protocol is long-term fair far from the border, but the short-term fairness horizon strongly depends on the channel access intensity.

In irregular topologies, the topological inequalities are not restricted to the border of the network, and the level of competition to access the channel varies from link to link. This adds a topological component to the unfairness already present in regular topologies. As a result, CSMA/CA protocols are typically very unfair in irregular topologies at all values of the access intensity, even though their unfairness is more pronounced at high values.

## ACKNOWLEDGMENT

We would like to thank Nicolas Macris (EPFL) for helpful discussions on Peierls' argument.

## REFERENCES

- [1] M. Durvy, O. Dousse, and P. Thiran, "Border Effects, Fairness, and Phase Transition in Large Wireless Networks," in *INFOCOM, Phoenix*, 2008.
- [2] C. Chaudet, D. Dhoutaut, and I. G. Lassous, "Experiments of Some Performance Issues with IEEE 802.11b in Ad Hoc Networks," in *WONS, St Moritz*, 2005.
- [3] X. Wang and K. Kar, "Throughput Modelling and Fairness Issues in CSMA/CA Based Ad-Hoc Networks," in *INFOCOM, Miami*, 2005.
- [4] M. Garetto, T. Salonidis, and E. Knightly, "Modeling Per-flow Throughput And Capturing Starvation in CSMA Multi-hop Wireless Networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 4, pp. 864–877, 2008.
- [5] K. Medepalli and F. Tobagi, "Towards Performance Modeling of IEEE 802.11 based Wireless Networks: A Unified Framework and its Applications," in *INFOCOM, Barcelona*, 2006.
- [6] M. Durvy and P. Thiran, "A Packing Approach to Compare Slotted and Non-Slotted Medium Access Control," in *INFOCOM, Barcelona*, 2006.
- [7] R. Kindermann and J. L. Snell, *Markov Random Fields and their Applications*. American Mathematical Society, Series: Contemporary mathematics, 1980.
- [8] H. Georgii, O. Häggström, and C. Maes, "The random geometry of equilibrium phases," in *Phase Transitions and Critical Phenomena*, C. Domb and J. Lebowitz, Eds. Academic Press, 2001, vol. 18, pp. 1–142.
- [9] J. van den Berg, "A uniqueness condition for gibbs measures, with application to the 2-dimensional ising antiferromagnet," *Commun. Math. Phys.*, vol. 152, pp. 161–166, 1993.
- [10] J. van den Berg and C. Maes, "Disagreement percolation in the study of markov fields," *Ann. probab.*, vol. 22, no. 2, pp. 91–114, 1994.
- [11] O. J. Heilmann and E. Praestgaard, "Phase transition of hard hexagons on a triangular lattice," *J. Stat. Phys.*, vol. 9, no. 1, pp. 23–44, 1973.
- [12] O. J. Heilmann, "The use of reflection as symmetry operation in connection with the peierls' argument," *Commun. Math. Phys.*, vol. 36, pp. 91–114, 1974.

- [13] F. A. Tobagi, "Modeling and performance analysis of multihop packet radio networks," *Proc. IEEE*, vol. 75, no. 1, pp. 135–155, 1987.
- [14] R. Boorstyn, A. Kershenbaum, B. Maglaris, and V. Sahin, "Throughput analysis in multihop CSMA packet radio networks," *IEEE Trans. Commun.*, vol. 35, no. 3, pp. 267–274, 1987.
- [15] F. P. Kelly, "Loss networks," *Ann. Appl. Probab.*, vol. 1, no. 3, pp. 319–378, 1991.
- [16] M. Durvy, O. Dousse, and P. Thiran, "Self-Organization Properties of CSMA/CA Systems and Their Consequences on Fairness," *IEEE Transactions on Information Theory*, vol. 55, no. 3, pp. 931–943, 2009.
- [17] C. Bordenave, D. McDonald, and A. Proutiere, "Performance of random medium access control: An asymptotic approach," in *Proc. ACM Sigmetrics'08*, June 2008.
- [18] A. Proutiere, Y. Yi, and M. Chiang, "Throughput of random medium access control: An asymptotic approach," in *Proc. CISS'08*, March 2008.
- [19] L. Jiang and J. Walrand, "A distributed CSMA for throughput and utility maximization in wireless networks," in *Proc. Allerton*, September 2008.
- [20] P. Marbach, A. Eryilmaz, and A. Ozdaglar, "Achievable rate region of CSMA schedulers in wireless networks with primary interferences constraints," in *Proc. IEEE Conf. on Decision and Control'07*, December 2007.
- [21] M. M. Carvalho and J. J. Garcia-Luna-Aceves, "A scalable model for channel access protocols in multihop ad hoc networks," in *MobiCom*. ACM Press, 2004, pp. 330–344.
- [22] M. Garetto, J. Shi, and E. W. Knightly, "Modeling Media Access in Embedded Two-Flow Topologies of Multi-Hop Wireless Networks," in *MobiCom*, 2005.
- [23] A. Jindal and K. Psounis, "Characterizing the Achievable Rate Region of Wireless Multi-hop Networks with 802.11 Scheduling," *To appear in IEEE/ACM Trans. Netw.*
- [24] Y. Yemini, "A statistical mechanics of distributed resource sharing mechanisms," in *INFOCOM, San Diego*, 1983.
- [25] G. M. Louth, *Stochastic networks: complexity, dependence and routing*. (thesis) Cambridge Univ., 1990.
- [26] I. Z. K. Ramanan, A. Sengupta and P. Mitra, "Markov random field models of multicasting in tree loss networks," *Adv. Appl. Probab.*, vol. 34, pp. 58–84, 2002.
- [27] K. R. B. Luen and I. Ziedins, "Nonmonotonicity of phase transitions in a loss network with controls," *Ann. Appl. Probab.*, vol. 16, no. 3, pp. 1528–1562, 2006.
- [28] M. Durvy and P. Thiran, "Understanding the Gap between the IEEE 802.11 Protocol Performance and the Theoretical Limits," in *SECON, Reston*, 2006.
- [29] "Discrete event simulator for the idealized 802.11 protocol," Available at <http://icapeople.epfl.ch/mdurvy/research.html>.
- [30] R. Jain, *The Art of Computer Systems Performance Analysis*. John Wiley and Sons, Inc., 1991.
- [31] G. Grimmett, *Percolation*. Springer-Verlag, 1999.
- [32] G. Brightwell, O. Häggström, and P. Winkler, "Nonmonotonic behavior in hard-core and widow-rowlinson models," *Journal of Statistical Physics*, vol. 94, pp. 415–435, 1999.
- [33] C. E. Koksal, H. Kassab, and H. Balakrishnan, "An analysis of short-term fairness in wireless media access protocols," in *Measurement and Modeling of Computer Systems*, 2000, pp. 118–119. [Online]. Available: [citeseer.ist.psu.edu/koksal00analysis.html](http://citeseer.ist.psu.edu/koksal00analysis.html)



**Mathilde Durvy** received her M.Sc. degree in 2002 and her Ph.D. degree in 2007 in Communication Systems from EPFL, Lausanne, Switzerland. She is currently working for Cisco Systems and is one of the co-authors of  $\mu$ IPv6, the smallest, open-source, IPv6 certified stack. Her research interests are in the design and analysis of communication protocols for wireless networks. In 2005 she received an Infocom travel grant, in 2007 she was a finalist of the Google Anita Borg scholarship, and in 2008 she received the best poster award at SenSys.



**Olivier Dousse** (S'02-M'05) received his M.Sc. degree in Physics from the Swiss Federal Institute of Technology at Lausanne, Switzerland (EPFL) in 2000, and his Ph.D. degree in Communication Systems from the same institution in 2005. From 2006 to 2008, he was with Deutsche Telekom Laboratories in Berlin. He is currently a Principal Member of Research Staff at Nokia Research Center in Lausanne. His research interests are in stochastic models for communication and social networks. He received the honorable mention of the 2005 ACM Doctoral Dissertation Competition, and he was awarded the Top 3 Student Prize of the EPFL Graduate School in communication systems in 2001. He was runner-up for the IEEE-Infocom Best Paper Award in 2003.



**Patrick Thiran** (S'89-M'97) is an associate professor at EPFL. He received the electrical engineering degree from the Université Catholique de Louvain, Louvain-la-Neuve, Belgium, in 1989, the M.S. degree in electrical engineering from the University of California at Berkeley, USA, in 1990, and the PhD degree from EPFL, in 1996. He became an adjunct professor in 1998, an assistant professor in 2002 and an associate professor in 2006. From 2000 to 2001, he was with Sprint Advanced Technology Labs, Burlingame, CA. His research interests include communication networks, performance analysis, dynamical systems and stochastic models. He is currently active in the analysis and design of wireless multi-hop networks and in network monitoring. He served as an associate editor for the IEEE Transactions on Circuits and Systems in 1997-99, and he is currently an associate editor for the IEEE/ACM Transactions on Networking. He was the recipient of the 1996 EPFL Ph.D. award, and of the 2008 Crédit Suisse Teaching Award.