



Traffic Capacity of Multi-Cell WLANs

Thomas Bonald, Ali Ibrahim, James Roberts

► **To cite this version:**

Thomas Bonald, Ali Ibrahim, James Roberts. Traffic Capacity of Multi-Cell WLANs. Sigmetrics, 2008, Annapolis, United States. hal-01244789

HAL Id: hal-01244789

<https://hal.archives-ouvertes.fr/hal-01244789>

Submitted on 16 Dec 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Traffic Capacity of Multi-Cell WLANs

Thomas Bonald, Ali Ibrahim, James Roberts
Orange Labs
38-40 rue general Leclerc
Issy-les-Moulineaux, France
{thomas.bonald,ali.ibrahim,james.roberts}@orange-ftgroup.com

ABSTRACT

Performance of WLANs has been extensively studied during the past few years. While the focus has mostly been on isolated cells, the coverage of WLANs is in practice most often realised through several cells. Cells using the same frequency channel typically interact through the exclusion region enforced by the RTS/CTS mechanism prior to the transmission of any packet.

In this paper, we investigate the impact of this interaction on the overall network capacity under realistic dynamic traffic conditions. Specifically, we represent each cell as a queue and derive the stability condition of the corresponding coupled queuing system. This condition is then used to calculate the network capacity. To gain insight into the particular nature of interference in multi-cell WLANs, we apply our model to a number of simple network topologies and explicitly derive the capacity in several cases. The results notably show that the capacity gain obtained by using M frequency channels can grow significantly faster than M , the rate one might intuitively expect. In addition to stability results, we present an approximate model to derive the impact of network load on the mean transfer rate seen by the users.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication* ;

C.4 [Computer Systems]: Performance of Systems—*Modeling techniques*

General Terms

Performance

Keywords

Multi-cell WLAN, IEEE 802.11, flow-level model, stability, capacity.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS'08, June 2–6, 2008, Annapolis, Maryland, USA.
Copyright 2008 ACM 978-1-60558-005-0/08/06 ...\$5.00.

1. INTRODUCTION

Our objective in the present paper is to investigate the downlink capacity of multiple interfering IEEE 802.11 Access Points (AP) using the Distributed Coordination Function (DCF) medium access control. Use of this technology to access the Internet in homes, enterprises and WiFi hotspots is growing rapidly leading to an increasingly dense implantation of APs. Interference between these APs and their respective users reduces overall traffic capacity. We consider a generic multiple AP network that we refer to as a multi-cell WLAN. The cell is the set of positions from which users will associate with a given AP. While the performance of single cell networks is well understood, there has as yet been relatively little evaluation of the impact of interference on the capacity of multi-cell networks.

We assume users situated within the coverage area of the WLAN download files from the Internet via their nearest AP. The rate at which downloads are generated times the average flow size defines a traffic intensity in bits per second. The traffic capacity of the WLAN is the maximum intensity it can support without the number of simultaneous downloads growing to infinity. The capacity depends somewhat on the statistical characteristics of demand as well as its spatial distribution over the WLAN coverage area. Our main focus, however, is on how capacity depends on the placement of the APs, their transmission range and the assignment of frequency channels to neighbouring cells.

The ultimate goal is to provide guidelines for planning multi-cell WLANs, to maximize their capacity or to minimize their cost, for example. However, our present ambition is limited to exploring the issues using relatively simple models that capture the impact of inter-cell interference in toy network configurations.

Interference is manifested in IEEE 802.11 networks by the CSMA collision avoidance protocol [9]. Stations must listen to the channel and defer transmission when it is sensed busy. When the channel is idle, stations compete for access using the exponential backoff procedure. In a multi-cell WLAN the efficiency of this procedure is compromised due to the hidden node phenomenon necessitating the use of channel reservation. This is performed through the exchange of short frames called RTS (Request-To-Send) and CTS (Clear-To-Send) prior to the transmission of a data packet.

Suppose, for example, that two users associated with different APs are located within their common transmission range while the two APs are not able to hear each other, as depicted in Figure 1(a). The transmission from AP₁ to user u_1 will not be detected by AP₂ leading to a collision




Figure 1: Impact of hidden nodes.

on reception if it begins to transmit to u_2 . The RTS/CTS mechanisms prevent this. The CTS frame sent by u_1 in response to an RTS from AP_1 is heard by AP_2 which refrains from transmitting until the channel is idle. In Figure 1(b), where the two downlink transmissions *can* coexist, collisions can still occur between data packets in one cell and acknowledgements in the other. With a prior RTS/CTS exchange, the second transmission would not take place since, after receiving the CTS from u_1 , u_2 would not respond to an RTS sent later by AP_2 . The area where channel access is inhibited during an ongoing transmission is thus extended to all user and AP positions within the transmission range of both sender and receiver. We term this the *exclusion region*.

The process by which stations gain channel access is very complex. Consider an AP with a packet to transmit to a particular user in its cell. It must first wait for any ongoing transmission in its exclusion region to complete. It may then continue to be blocked by some other concurrent transmission within this exclusion region but compatible with the first transmission. When it can eventually try to access the channel, it will be in competition with a variable number of other AP-user pairs. Depending on the outcome of this competition, it may be necessary to make several successive attempts before successful transmission. It seems impossible to precisely model such a complex stochastic system where all cells are inter-dependent through their overlapping exclusion zones. On the other hand, simulation is hardly useful to provide the general insights we seek. We therefore make some quite bold assumptions and simplifications, as discussed in the paper, in order to derive analytical results.

In this simplified setting, we are able to derive the stability condition and deduce the traffic capacity of multi-cell WLANs. When demand is less than capacity, we can estimate the mean time to perform a download for any given AP and user position. These general results are evaluated numerically for some toy network topologies. The results show how capacity changes abruptly depending on whether APs interfere directly or only through the users to which they transmit. This discontinuity has a significant impact on the potential gain in capacity brought by the use of different frequency channels in neighbouring cells. At best, this gain is significant, with capacity increasing up to four times as fast as the number of channels. Arguably, the broad characteristics of this behaviour do not depend on the simplifications introduced in our modelling and therefore provide valuable insight into the performance of real multi-cell WLANs.

In the remainder of the paper we introduce our model and

present the analysis that leads to closed form expressions for the traffic capacity and a procedure to evaluate mean flow download times as a function of load. The results are applied in the next sections to simple network configurations, with one or several channels, in one and two dimensional space. We evaluate flow throughput for the simplest two-AP network configuration and validate these results by means of packet-level simulations. We first present a brief review of related literature.

2. RELATED WORK

There is a huge amount of literature on the performance of IEEE 802.11 under DCF access control. However, much of this is focused on the performance of isolated cells with a static user population. In the seminal work of Bianchi [2] and its generalizations, by Gupta and Kumar [8] and by Kumar et al [11], for instance, the capacity of a single cell with N saturated stations uploading packets to the AP is determined as the solution to a fixed point equation. This is not directly related to our work, however, where we focus on downlink throughput and evaluate the impact of interference between multiple APs under dynamic traffic.

The performance of multi-hop wireless networks calls for an evaluation of interference effects similar to those occurring in our multi-cell WLAN. For example, Tassiulas and Ephremides [18] characterize the maximal capacity region of such networks and the scheduling algorithm that achieves this. More recent work has focused on determining scheduling algorithms that realize given performance objectives, e.g., [10, 13]. This work is useful in illustrating the inherent difficulty of accounting for the impact of interference. Our context is somewhat simpler in that connections are all single hop and we do not seek to define a scheduling algorithm. However, our focus is different in that we evaluate capacity in a dynamic scenario and how it depends on the relative positions of the APs.

The work by Panda et al. [16] is, at first sight, particularly relevant. The authors consider the same configuration of interfering WLAN APs. However, they consider uplink traffic only, generalizing the model of Bianchi [2] for a network of two APs with so-called critical inter-AP spacing. We consider downlink traffic with two or more interfering APs.

Traffic capacity and throughput performance, in the flow level sense considered here, was evaluated by Lebeugle and Proutière [12] and Litjens et al. [14] for a single AP cell. Their results are a useful justification for our assumption in the next sections that the DCF protocols leads to approximate fair sharing of a constant cell capacity.

The present analysis is also closely related to previous work on the impact of inter-cell interference on the performance of cellular networks. Bonald et al. [3] evaluate bounds on download throughput performance under the flow level traffic model considered here. Although the network is also represented as a coupled queuing system, the nature of the interference phenomenon is different. A key feature of our model of multi-cell WLANs is that the service rate of each queue does not only depend on the activity state of the other queues, that is on the presence or absence of active users in the corresponding cells, but on the *locations* of these active users, that may collide or not with ongoing data transfers in the considered cell.

3. MODEL

In this section, we describe the model used to analyse the impact of interference on the traffic capacity of multi-cell WLANs. Each cell is represented as a queue that interacts with other queues through coupled service rates. This coupling, that captures the interference between neighbouring cells, is derived from a simple packet-level model representing the impact of RTS/CTS mechanisms.

3.1 Network structure

Consider a WLAN that consists of N access points (AP). We focus on the downstream traffic from the APs to randomly located users in the coverage region of the network. We assume this traffic consists of elastic data transfers, typically under the control of TCP.

There is an arbitrary finite set \mathcal{U} of user classes, corresponding to homogeneous regions in terms of radio characteristics. Specifically, two users of the same class are associated with the same AP and interact with other users and APs in the same way. We denote by $\mathcal{U}_i \subset \mathcal{U}$ the set of user classes associated with AP i . This corresponds to the cell served by AP i .

Users interfere through the RTS/CTS mechanisms, as described in Section 1. We model this interference through a function χ from $\mathcal{U} \times \mathcal{U}$ to $\{0, 1\}$, such that $\chi(j, l) = 1$ if and only if the transmissions to a class- j user and a class- l user cannot occur simultaneously. In particular, we have $\chi(j, l) = 1$ if user classes j, l are associated with the same AP. For all $i = 1, \dots, N$ and $j \in \mathcal{U}_i$, we refer to the set $\mathcal{E}_j^i = \{l \in \mathcal{U} : \chi(j, l) = 1\}$ as the *exclusion region* of user class j associated with AP i .




Figure 2: Notion of exclusion region.

Though we assume that both the number of APs and the set of classes are finite in the analysis, the results naturally extend to infinite networks with a continuous set of classes. For 2D networks for instance, each class may represent a specific location in the network, so that $\mathcal{U} = \mathbb{R}^2$. A simple interference model in this context consists in considering any transmission successful if there is no other transmitting or receiving station within a distance R from the source and from the receiver, the distance R corresponding the transmission range of RTS/CTS signals for both users and APs.

The interference function between two users located in

$u_1, u_2 \in \mathbb{R}^2$ is then given by:

$$\chi(u_1, u_2) = 0 \iff \begin{cases} d(u_1, u_2) > R, & d(u_1, v_2) > R \\ d(v_1, u_2) > R, & d(v_1, v_2) > R \end{cases}$$

where v_1, v_2 are the respective locations of the associated APs and $d(u, v)$ is the distance between u and v . Thus the exclusion region of any transmission typically consists of two disks, one centered on the AP and the other on the user, as illustrated by Figure 2(a). When some APs are located in the region formed by these two disks, the whole cells served by these APs must be added to the exclusion region. This is illustrated by Figure 2(b).

3.2 Packet-level model

Packet scheduling is assumed to be FIFO at each AP. When the transmission of a packet is scheduled, the AP performs DCF access control competing with other transmissions in its exclusion region until the packet is eventually delivered. Thus the APs interact in a very complicated way that depends on the location of active users in the corresponding cells.

To get some insight into the impact of this interaction on the overall network behaviour, we consider the following simple packet-level model. All packets have the same size and the network operates synchronously in a slotted fashion: transmissions occur only at the beginning of a time slot and the transmission of each packet takes one slot. The slot duration may be seen as the average transmission time of a packet for an isolated AP, including all overheads like the times to access the channel and to transmit control messages like RTS/CTS signals and MAC/TCP acknowledgements. Thus the throughput of an isolated AP is equal to 1 packet/slot, which is the throughput unit used in the rest of the paper. Note that we neglect the impact of variable radio conditions, that may result in different physical throughputs depending on the location in the cell. In our model, the throughput of a user depends on her/his location in a given cell through her/his interaction with other cells only.

Now let ξ_j^i be the probability that AP i attempts to serve a class- j user after any successful transmission. We assume that the mean number of slots an AP remains idle before transmitting a packet is equal to the number of other active users in the corresponding exclusion region, when each AP i independently selects a class- j user with probability ξ_j^i . Thus the mean time for a user to access the channel is proportional to the number of active transfers in her/his exclusion region, which is a reasonable assumption. For all $k \neq i$ and $j \in \mathcal{U}_i$, the probability that the user selected by AP k is in the exclusion region \mathcal{E}_j^i of class- j users is given by:

$$\sum_{l \in \mathcal{U}_k} \xi_l^k \chi(j, l).$$

We deduce the average number of slots δ_j^i needed by AP i to successfully transmit the packet of a class- j user:

$$\delta_j^i = 1 + \sum_{k \neq i} \sum_{l \in \mathcal{U}_k} \xi_l^k \chi(j, l). \quad (1)$$

The average number of slots needed by AP i to successfully transmit any packet is then given by:

$$\delta_i = \sum_{j \in \mathcal{U}_i} \xi_j^i \delta_j^i. \quad (2)$$

We conclude that the throughput of AP i is equal to $1/\delta_i$. This throughput is equal to 1 for an isolated AP and is less than 1 in the presence of inter-cell interference. A fraction ξ_j^i of the throughput of AP i is allocated to class- j users. Thus the throughput allocation is entirely determined by the interference function χ and the probabilities ξ_j^i , for all $i = 1, \dots, N$ and $j \in \mathcal{U}_i$.

3.3 Flow-level model

In the following, we refer to a class- j flow as a data transfer to a class- j user. We assume class- j flows arrive at AP i as a Poisson process of intensity λ_j^i . Flow sizes are i.i.d. exponential with parameter μ .

Let x_j^i be the number of class- j flows at AP i . We denote by x_i the total number of flows at AP i :

$$x_i = \sum_{j \in \mathcal{U}_i} x_j^i.$$

We are interested in the evolution of the network state x that describes the number of ongoing flows of each class at each AP. This depends on the throughput allocation in state x . In order to apply the above packet-level model, it remains to determine the probabilities ξ_j^i in state x , for all $i = 1, \dots, N$ and $j \in \mathcal{U}_i$.

Assume all flows associated with same AP get the same throughput. This is a natural assumption under the assumed FIFO packet scheduling policy, since these flows experience the same packet delay and the same packet loss rate at the AP. The probability that AP i selects the packet of a class- j flow after a successful transmission is then proportional to x_j^i in all states x such that $x_i > 0$:

$$\xi_j^i = \frac{x_j^i}{x_i}.$$

In view of (1) and (2), the throughput of AP i is equal to $1/\delta_i$, with

$$\delta_i = \sum_{j \in \mathcal{U}_i} \frac{x_j^i}{x_i} \delta_j^i \quad (3)$$

and

$$\delta_j^i = 1 + \sum_{k \neq i: x_k > 0} \sum_{l \in \mathcal{U}_k} \frac{x_l^k}{x_k} \chi(j, l). \quad (4)$$

In addition, this throughput is evenly shared by all active flows served by AP i .

Thus the model corresponds to a network of N multi-class processor-sharing queues with state-dependent service rates. Class- j customers arrive at queue i as a Poisson process of intensity λ_j^i and have i.i.d. exponential service requirements with parameter μ . In view of (3)-(4), the service rate $1/\delta_i$ of queue i depends on the whole network state x that describes the number of customers of each class in each queue. It is this coupling of service rates that captures the impact of interference between neighbouring cells.

4. ANALYSIS

We first determine the stability region of the network, from which we derive the notion of traffic capacity. We then propose a fixed-point approximation for the flow-level performance when the network is stable.

4.1 Stability region

Let $\rho_j^i = \lambda_j^i/\mu$ be the traffic intensity of class- j flows at AP i . We denote by α_j^i the proportion of class- j traffic at AP i :

$$\alpha_j^i = \frac{\rho_j^i}{\rho_i} \quad \text{with} \quad \rho_i = \sum_{j \in \mathcal{U}_i} \rho_j^i.$$

We say that the network is stable if the Markov process that describes the evolution of the network state x is ergodic. We have the following key results, proved in Appendix A.

THEOREM 1. *The network is stable if for all $i = 1, \dots, N$,*

$$\rho_i \sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i < 1,$$

where

$$\beta_j^i = 1 + \sum_{k \neq i} \sum_{l \in \mathcal{U}_k} \alpha_l^k \chi(j, l). \quad (5)$$

Observe that, for isolated cells, the model reduces to mutually independent processor-sharing queues so that the stability condition becomes $\rho_i < 1$ for all $i = 1, \dots, N$. The impact of inter-cell interference on stability is captured by the multiplying factor $\sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i$ on the traffic intensity at AP i . In the following, we refer to the *load* of AP i as the product:

$$\rho_i \times \sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i.$$

THEOREM 2. *If the network is stable, there exists some $i \in \{1, \dots, N\}$ such that*

$$\rho_i \sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i \leq 1.$$

The inequalities of Theorems 1 and 2 coincide, up to the critical case, when all APs have the same load. This allows us to calculate in Sections 5 and 6 the traffic capacity of a number of practically interesting symmetric networks, following the approach described below. It proves very difficult to derive the necessary and sufficient stability condition for heterogeneous load distributions, as in most coupled queuing systems with more than two queues [1, 4, 6, 17].

4.2 Traffic capacity

We define the traffic capacity of a network as the maximum traffic intensity such that the network is stable, assuming a fixed traffic distribution among user classes. When all APs have the same load, the necessary and sufficient stability condition is, up to the critical case, that for all $i = 1, \dots, N$,

$$\rho_i \sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i < 1.$$

We deduce the traffic capacity of cell i :

$$C_i = \frac{1}{\sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i}.$$

This traffic capacity is equal to 1 if the cell is isolated (since $\beta_j^i = 1$ for all $j \in \mathcal{U}_i$) and is less than 1 in the presence of inter-cell interference.

In the continuous setting introduced in Section 3, assuming cells of equal area A and uniform traffic distribution throughout the network, we obtain:

$$C = \frac{A}{\int_{\mathcal{C}} \beta(u) du}, \quad (6)$$

where \mathcal{C} is the considered cell and, in view of (5),

$$\beta(u) = \frac{1}{A} \int_{\mathcal{U}} \chi(u, v) dv. \quad (7)$$

Note that $\beta(u)$ is equal to the area of the exclusion region associated with a user located in u normalized by the cell area. It may be simply interpreted as a measure of the *average interference* suffered at location u . We apply formulas (6) and (7) to a number of symmetric network topologies in Sections 5 and 6.

4.3 Mean transfer time

Finally, it is interesting to analyse how the throughput of each user depends on network load under the stability condition. In the following, we refer to the flow throughput of class j at AP i as the ratio of the mean flow size, $1/\mu$, to the mean duration W_j^i of class- j flows:

$$\gamma_j^i = \frac{1}{\mu W_j^i}.$$

By Little's law, the mean number of class- j flows at AP i is given by:

$$X_j^i = \lambda_j^i W_j^i. \quad (8)$$

Unfortunately, there is no simple expression for the stationary distribution of the network state. A useful approximation consists in decoupling queues by replacing the state-dependent parameters δ_j^i by constants ϕ_j^i . In view of (3), class- j flows are then served at rate:

$$\frac{x_j^i}{\sum_{l \in \mathcal{U}_i} x_l^i \phi_l^i}.$$

This is equivalent to a discriminatory processor-sharing queue of unit service rate with weights ϕ_j^i , where class- j customers require exponential services of parameter μ/ϕ_j^i . We assume that the queue is stable, that is:

$$\rho_i \sum_{j \in \mathcal{U}_i} \alpha_j^i \phi_j^i < 1.$$

We can then apply the results of Fayolle et. al. [7] to get the mean flow duration of each class:

$$W_j^i \left(1 - \frac{\rho_i}{2} \sum_{l \in \mathcal{U}_i} \alpha_l^i \phi_l^i \right) - \frac{\rho_i}{2} \sum_{l \in \mathcal{U}_i} \alpha_l^i \phi_l^i W_l^i = \frac{\phi_j^i}{\mu}. \quad (9)$$

It remains to estimate the weights ϕ_j^i , for which we use the following fixed point equations:

$$\phi_j^i = 1 + \sum_{k \neq i} \left(\sum_{l \in \mathcal{U}_k} \phi_l^k \rho_l^k \right) \frac{\sum_{l \in \mathcal{U}_k} X_l^k \chi(j, l)}{\sum_{l \in \mathcal{U}_k} X_l^k}. \quad (10)$$

These equations follow from (3) by replacing the condition $x_k > 0$ by the probability that $x_k > 0$ and the random variables x_l^k by their mean X_l^k . This approximation turns out to be very accurate, as shown in Section 7.

5. 1D NETWORKS

In this section, we calculate the traffic capacity of 1D networks, assuming users are randomly located on a line. The set of user classes is $\mathcal{U} = \mathbb{R}$. The traffic distribution is assumed to be uniform on the coverage region of the network. We consider a simple radio model where users associate with the nearest AP and any transmission is successful if and only if there is no other transmitting station within a distance R from the source and from the receiver. We take $R = 1$ in the rest of the paper so that the unit length corresponds to the transmission range of both APs and users.

We first consider the simple case of two APs. This provides useful insight into the impact of interference and constitutes a basic building block for understanding the behaviour of more complex network topologies. We then analyse the more practically interesting case of an infinite number of APs regularly located on the line and sharing one or more frequency channels.

5.1 Case of two APs

Consider two APs separated by a distance d : AP 1 is located at $-\frac{d}{2}$ and AP 2 at $\frac{d}{2}$. When $d > 2$, each cell is a segment of length 2 centered at the AP; when $d < 2$, the cells are the segments $[-1 - \frac{d}{2}, 0]$ and $[0, \frac{d}{2} + 1]$. Now consider one user in each cell, located at u_1 and u_2 , respectively. The mutual interference region of these two users is shown by the shaded areas of Figure 4, where $r_1 = u_1 + \frac{d}{2}$ and $r_2 = u_2 - \frac{d}{2}$ correspond to the relative location of these users in the corresponding cells, as illustrated by Figure 3.




Figure 3: 1D network with two APs.

Applying (6), we obtain the traffic capacity C of each cell as a function of d . We distinguish four cases:

(a) *Total coupling*, $d \in (0, 1]$: The two APs can never transmit simultaneously and thus behave as a single cell of unit service rate. We have:

$$C(d) = \frac{1}{2}.$$

(b) *Partial coupling, overlapping transmission areas*, $d \in (1, 2]$: The APs can only interfere indirectly through their destination. We obtain:

$$C(d) = \frac{\left(1 + \frac{d}{2}\right)^2}{\frac{5}{2} + d}.$$

(c) *Partial coupling, non-overlapping transmission areas*, $d \in (2, 3]$: The two transmission areas no longer overlap, but the cells still interfere. We have:

$$C(d) = \frac{8}{17 - 6d + d^2}.$$

(d) *Independent cells*, $d \in (3, \infty)$: Each cell behaves as if it were alone yielding:

$$C(d) = 1.$$

The results are illustrated by Figure 5. We observe that capacity has irregular behaviour at the three critical distances $d = 1, 2$ and 3 . At $d = 1$ capacity jumps discontinuously from $\frac{1}{2}$ to $\frac{9}{14}$ as the APs cease to interfere directly.




Figure 4: Mutual interference region of two users.




Figure 5: Traffic capacity vs. inter-AP distance.

The form of the shaded interference region changes abruptly as we move from case (a) to case (b) in Figure 4. At $d = 2$ and $d = 3$ capacity is continuous but its slope is discontinuous. The reason is again the change in form of the exclusion region as we move from (b) to (c) and from (c) to (d).

5.2 Infinite linear network in 1D space

Consider an infinite 1D linear network where APs are located respectively at nd , $n \in \mathbb{Z}$, as shown in Figure 6.




Figure 6: Infinite linear network in 1D.

We first assume all APs use the same channel and thus use the subscript 1. We distinguish two cases:

(a) *Overlapping transmission areas*, $d \in (0, 2]$. The cell associated with AP i is the segment $[(i - 1/2)d, (i + 1/2)d]$. Consider a particular user located at $u \in (0, \frac{d}{2}]$. Its exclusion region is equal to the union of the segment $[-1, u + 1]$ and of all cells whose AP is located on this segment. In view of (7), the level of interference $\beta(u)$ is equal to the area of this exclusion region normalized by the cell area, d . It may be

expressed as:

$$\beta(u) = 1 + e_1 \left(\frac{u+1}{d} \right) + e_1 \left(\frac{-1}{d} \right),$$

where $e_1(s)$ is the total area of cells whose AP is located on the segment $(0, sd]$. This function is illustrated in Figure 7. The steps at $s = 1, 2, \dots$ occur as an AP is included in the segment ($e_1(s) = s$ for $s = 1, 2, \dots$); e_1 then remains constant during the first half interval and increases linearly with slope 1 during the next half; it jumps again as the next AP is included.




Figure 7: Functions e_M , $M = 1, 2, 3$.

We deduce from (6) the capacity C_1 of each cell as a function of the inter-AP distance:

$$C_1(d) = \left(1 + e_1 \left(\frac{1}{d} \right) + 2 \int_{\frac{1}{d}}^{\frac{1}{d} + \frac{1}{2}} e_1(s) ds \right)^{-1}. \quad (11)$$

(b) *Non-overlapping transmission areas*, $d \in (2, 3]$. In this case, the transmission areas do not overlap as the separation distance d is greater than 2. A given cell interferes only with its immediate neighbours since the separation distance with other cells is greater than 4. The total interference seen from the cell is therefore equal to twice that of two interfering APs derived in the previous section and the cell capacity is:

$$C_1(d) = \frac{4}{13 - 6d + d^2}.$$

Figure 8 illustrates the cell capacity as a function of the inter-AP distance d . We observe that capacity presents exactly the same types of irregularity discussed in the previous example but this time the set of irregular points is discrete and infinite.

The graph can best be understood by observing how cell capacity decreases from 1 to 0 as the distance between access points shrinks from $+\infty$ to 0. The first type of irregularity is discontinuity occurring at points $d = \frac{1}{n}$, where capacity decreases sharply to $\frac{1}{2n+1}$. Each of these points corresponds to the introduction of a new tier of interfering APs into the transmission range of the considered AP, say 0. The n -th interfering tier consisting of APs n and $-n$ enters the transmission range of AP 0 when $nd = 1$. At this point interference with tier $n+1$ is null and cell capacity is $\frac{1}{2n+1}$, $2n+1$ being the number of APs located within the transmission range. As the inter-AP distance decreases further, interference with tier $n+1$ begins and proceeds as in the case of two cells. The result is slope discontinuity occurring when $(n+1/2)d = 1$. Capacity at this point can also be computed easily by observing that each cell of tier $n+1$




Figure 8: Capacity, infinite network in 1D.




Figure 9: Density, infinite network in 1D.

increases the mean slot time by $\frac{1}{8}$ and all previous tiers are within the transmission range. The capacity is then $\frac{4}{8n+5}$. Between each couple of irregular points, capacity behaves like that of two interfering cells.

Figure 9 shows the maximal traffic density normalised to the maximal traffic density of an isolated cell. By definition, the maximal traffic density is the ratio of cell capacity to the cell area:

$$\varrho_1(d) = \frac{C_1(d)}{\min(d, 2)}.$$

It demonstrates that a maximal density of 1 is possible only when $d \approx 0$ or $d > 3$. In the latter case the network no longer has complete coverage. The former is obtained at the cost of very high AP density.

5.3 Use of multiple channels

We generalise the results of the previous section to a network with $M > 1$ non-interfering channels. The APs operating on channel k , $k = 0, \dots, M-1$ are placed at $(nM + k)d$, $n \in \mathbb{Z}$. As in the case of a single channel, we distinguish two cases: *overlapping* ($d \leq 2$) and *non-overlapping* ($d > 2$) transmission areas. The case of *non-overlapping* is trivial, capacity C_M is always equal to 1 because the distance between any two interfering APs is greater than 3. The case of *overlapping* is obtained from equation (11) by replacing e_1 with e_M :

$$C_M(d) = \left(1 + e_M\left(\frac{1}{d}\right) + 2 \int_{\frac{1}{d}}^{\frac{1}{d} + \frac{1}{2}} e_M(s) ds\right)^{-1},$$

where $e_M(s)$ is the total area of cells whose AP is located on the segment $(0, sd]$ and uses the same channel as AP 0. The

function is illustrated in Figure 7 for $M = 2, 3$. Note that e_M steps by 1 at sM , $s \in \mathbb{Z}$ as a new interfering AP is included in the segment; it remains constant for $\frac{1}{2} + (M-1)$ and then increases linearly with slope 1 until it reaches $(s+1)M$; it jumps again at this point and so on.

Figure 8 illustrates the cell capacity for $M = 2, 3$. Capacity increases as expected with the number of channels while the qualitative behaviour remains the same. We have jumps at points $d = \frac{1}{nM}$ and slope discontinuity at $d = \frac{1}{nM - \frac{1}{2}}$ for $n = 1, 2, \dots$.

The maximal traffic densities ϱ_M for $M = 2, 3$ are shown in Figure 9. The figure shows that traffic density has an infinite number of (local maximum, local minimum) pairs. This occurs because the capacity slope decreases from 0 to $-\infty$ as we move from $\frac{1}{(n-1)M}$ to $\frac{1}{nM}$. The local maximum is attained in segment $[\frac{1}{nM - \frac{1}{2}}, \frac{1}{nM - 1}]$ and the local minimum is attained at $\frac{1}{nM}$. The global maximum always occurs during interference with the first tier only and is therefore located between $\frac{1}{M - \frac{1}{2}}$ and $\frac{1}{M - 1}$ with:

$$\varrho_M(d) = \frac{\frac{2}{d}}{1 + \left(\frac{1}{d} + 1 - M\right)^2}.$$

The maximum of this expression is $\sqrt{M^2 - 2M + 2} + M - 1$ occurring at $d_M^* = 1/\sqrt{M^2 - 2M + 2}$. For example, if only two channels are used, the optimal inter-AP distance is $d = 1/\sqrt{2}$ and the network can accept up to 2.4 times as much traffic as a single channel network. When $M = 3$ (as in 802.11b/g), capacity increases 4.23 fold when $d = 1/\sqrt{5}$.

6. 2D NETWORKS

We now investigate the traffic capacity of 2D networks. The set of user classes is $\mathcal{U} = \mathbb{R}^2$. Again, the traffic distribution is assumed to be uniform on the coverage region of the network. We successively consider four network topologies. We only outline results derived in a similar way to those for 1D networks. To obtain analytic expressions we use the infinity norm distance, $d_\infty(u, v) = \max(|u_1 - v_1|, |u_2 - v_2|)$ for all $u, v \in \mathbb{R}^2$, instead of the 2-norm distance, $d_2(u, v) = \sqrt{(u_1 - v_1)^2 + (u_2 - v_2)^2}$ for all $u, v \in \mathbb{R}^2$. The transmission area of each AP is thus no longer a circle of radius 1 but a square with sides of length 2. Of course, the approach applies to the usual euclidian distance as well but leads to more complex expressions.

6.1 Two and four APs

APs in the two node network are separated by a distance d . In the 4 node network they are at the vertices of a square with sides of length d . Formulas for the capacity can be derived by generalizing the method used in §5.1. Figure 5 shows the respective network capacities. The capacity of 2APs in 2D is slightly higher than in 1D. This is because the fraction of user positions included in an exclusion zone is somewhat smaller in 2D. On the other hand, interference in the 4 AP network is greater leading to significantly reduced capacity.

6.2 Infinite linear network in 2D space

We generalise to 2D the network of Figure 6. Each cell is now a rectangle of height 2 and width d centred at the AP if $d \leq 2$ and a square with side length 2 if $d > 2$. M frequency

M	d_M^*	ϱ_M^*	$\frac{1}{M-1}$	$\varrho_M\left(\frac{1}{M-1}\right)$
4	0.3156	6.1688	0.3333	6
9	0.1240	16.0632	0.1250	16
25	0.0416	48.0209	0.0417	48
100	0.0101	198.0051	0.0101	198

Table 1: Infinite linear network in 2D space.

channels are allocated to the APs as in §5.2. In this and the next sections, we evaluate the optimal traffic density ϱ_M^* . Proceeding similarly as in the previous sections, we derive the interference β at some fixed point (u, v) of the reference cell:

$$\beta(u, v) = 1 + e_M\left(\frac{-1}{d}\right) + e_M\left(\frac{1}{d}\right) + \left(e_M\left(\frac{u+1}{d}\right) - e_M\left(\frac{1}{d}\right)\right)\left(\frac{2-|v|}{2}\right).$$

The factor $(2-|v|)/2$ accounts for the fact that the receiver is now situated in 2D. It is less than 1 so that capacity in 2D is greater than in 1D. We have:

$$C_M(d) = \left(1 + 2e_M\left(\frac{1}{d}\right) + \left(2 \int_{\frac{1}{d}}^{\frac{1}{d}+\frac{1}{2}} e_M(s) ds - e\left(\frac{1}{d}\right)\right)\left(1 - \frac{d}{8}\right)\right)^{-1}.$$

Capacity for $M = 1, 4$ is illustrated in Figure 10. Maximal traffic density is given by

$$\varrho_M(d) = \frac{C_M(d)}{2 \min(d, 2)}.$$

It is optimal when interference occurs only with tier 1 for d between $\frac{1}{M-\frac{1}{2}}$ and $\frac{1}{M-1}$. In this range, ϱ_M has the following expression:

$$\varrho_M(d) = \frac{\frac{2}{d}}{1 + \left(1 - \frac{d}{8}\right)\left(\frac{1}{d} - M + 1\right)}.$$

The optimal density can be shown to be greater than $2(M-1)$, the value obtained for $d = 1/(M-1)$. Table 1 gives the exact optimal density derived numerically and shows that the bound is a good approximation for $M \geq 4$. The scaling behaviour is the same for the linear network in 1D and 2D.




Figure 10: Capacity, infinite network in 2D.




Figure 11: Density, infinite network in 2D.

6.3 Infinite grid network in 2D space

In the 2D grid network, the APs are at the vertices of a regular grid where the distance between adjacent vertices is d . Each cell is a square with sides of length d if $d < 2$ and 2 if $d > 2$. The number of channels is M which we assume to be the square of an integer. For the sake of completeness we give the expression of $\beta(u, v)$ and C_M :

$$\beta(u, v) = \left(2e_{\sqrt{M}}\left(\frac{1}{d}\right) + 1\right)^2 + \frac{2}{d} \left(e_{\sqrt{M}}\left(\frac{u+1}{d}\right) - e_{\sqrt{M}}\left(\frac{1}{d}\right)\right) + \frac{2}{d} \left(e_{\sqrt{M}}\left(\frac{v+1}{d}\right) - e_{\sqrt{M}}\left(\frac{1}{d}\right)\right) - \left(e_{\sqrt{M}}\left(\frac{u+1}{d}\right) - e_{\sqrt{M}}\left(\frac{1}{d}\right)\right) \left(e_{\sqrt{M}}\left(\frac{v+1}{d}\right) - e_{\sqrt{M}}\left(\frac{1}{d}\right)\right),$$

$$C_M(d) = \left(1 + 3e_{\sqrt{M}}^2\left(\frac{1}{d}\right) + 4e_{\sqrt{M}}\left(\frac{1}{d}\right) - 4\frac{1}{d}e_{\sqrt{M}}\left(\frac{1}{d}\right) + 4\left(\frac{2}{d} + e_{\sqrt{M}}\left(\frac{1}{d}\right)\right) \int_{\frac{1}{d}}^{\frac{1}{d}+\frac{1}{2}} e_{\sqrt{M}}(s) ds - \left(2 \int_{\frac{1}{d}}^{\frac{1}{d}+\frac{1}{2}} e_{\sqrt{M}}(s) ds\right)^2\right)^{-1}.$$

The maximal traffic density is now given by

$$\varrho_M(d) = \frac{C_M(d)}{\left(\min(d, 2)\right)^2}.$$

The capacity and maximal traffic density are plotted in Figures 10 and 11, respectively, for $M = 1, 4$. The optimal traffic density occurs in interval $\left[\frac{1}{\sqrt{M}-\frac{1}{2}}, \frac{1}{\sqrt{M}-1}\right]$ where ϱ_M is given by

$$\varrho_M(d) = \frac{\frac{4}{d^2}}{1 + \frac{4}{d} \left(\frac{1}{d} + 1 - \sqrt{M}\right)^2 - \left(\frac{1}{d} + 1 - \sqrt{M}\right)^4}.$$

As above we can lower bound the optimal density by its value at $d = 1/(\sqrt{M}-1)$, $4(\sqrt{M}-1)^2$. Table 2 confirms that this bound is a good approximation for $M \geq 4$. We observe that maximal traffic density increases 4 times as fast as the number of frequency channels.

$$\varrho_M^* > \varrho_M\left(\frac{1}{\sqrt{M}-1}\right) = 4(\sqrt{M}-1)^2.$$

Again, for the purpose of comparison, we provide in Table 2 the exact and approximate values for d_M^* and ϱ_M^* . The table shows that for large M , ϱ_M^* scales approximately as $4(\sqrt{M}-1)^2$. Thus, even if the cells are subject to more

M	d_M^*	ϱ_M^*	$\frac{1}{\sqrt{M-1}}$	$\varrho_M\left(\frac{1}{\sqrt{M-1}}\right)$
4	0.8376	4.8395	1	4
9	0.4855	16.4856	0.5	16
25	0.2490	64.2490	0.2500	64
100	0.1111	324.1111	0.1111	324

Table 2: Infinite grid network in 2D space.

interference in 2D, the maximal traffic density scales as four times the number of used frequency channels.

7. FLOW-LEVEL PERFORMANCE

In this section we evaluate the mean flow rate as a function of user position and inter-AP distance d for the simple two AP network of §5.1. In the following ϱ is traffic density per unit of length.

7.1 Impact of traffic intensity

We take $d = 1.2$, which corresponds to the case of partial coupling and overlapping transmission areas, cf. Figure 4. The cell length is equal to $1 + d/2 = 1.6$. In Figure 12 we plot the mean flow rate as a function of traffic intensity per cell, $\rho = \varrho \times 1.6$, for two user positions. One user position is at the edge of the right hand cell ($u = 1.6$), the other midway between the two APs ($u = 0$).

Fig. 12(a) is derived from the approximate analytical model of §4.3 while Fig. 12(b) plots results from simulation of the underlying state-dependent processor sharing system of §3.3. The approximation agrees closely with the simulation results. The figures show behaviour typical of a discriminatory processor sharing system. The mean rate (for all user positions) decreases approximately linearly from 1 at zero load (any flow receives all the cell capacity) to 0 as traffic attains the capacity C evaluated in §5.1 for $d = 1.2$ ($C \approx 0.7$). Users at the far edge gain higher throughput on average than users situated between the APs but the difference is slight and disappears at the extremes. Users at $u = 1.6$ gain higher throughput since they do not interfere with the other cell while users at $u = 0$ always interfere. The difference is not large, however, since all active users in a cell in any occupancy state receive exactly the same rate.

Figures 12(c) and 12(d) correspond to NS-2 simulations with UDP traffic. The mean flow rate and traffic intensity in both scenarios are normalised to the capacity of an isolated AP. The RTS/CTS signals are assumed to be instantaneously exchanged in the Fig. 12(c) while 12(d) corresponds to the actual IEEE 802.11 standard. The analytic results agree with simulations in Fig. 12(c) but differ somewhat in Fig. 12(d). The flow rate goes to zero as traffic intensity approaches 0.6 which is less than the computed capacity $C \approx 0.7$. This is because the non-negligible RTS/CTS transmission times reduce capacity through hidden-node collisions, as follows.

In 802.11b, the RTS transmission time is approximately equal to 10 backoff slots while the minimal contention window length is 32 slots. For $d = 1.2$ the APs are not able to hear each other; if one AP transmits an RTS to a user located midway between the two APs, the other AP will provoke a collision at the user station if it also attempts a




Figure 12: Impact of traffic density.

transmission during the 10 slot period. On the other hand, if the APs can hear each other, the RTS/CTS scheme succeeds if and only if the backoff counters do not have the same value.

7.2 Impact of distance

Figure 13 plots the mean flow rate against the inter-AP distance d for two traffic densities, $\varrho = 0.25$ and $\varrho = 0.4$. The flow rate variations are irregular since capacity evolves differently with d in the four regions identified in §5.1. Note that flow rate is zero at load $\varrho = 0.4$ for $d \in [0.55, 1]$ since traffic intensity exceeds capacity. Flow rate increases between $d = 2$ and $d = 3$ but, of course, this does not take into account users situated out of range of either AP.

Simulation results in Fig. 13(b) and Fig. 13(c) confirm the accuracy the model when RTS/CTS times are neglected. However, the plots in Fig. 13(d) show some discrepancies. The flow rate is lower in the range $1 < d < 2$. This is due to the hidden-node problem explained above, the impact being particularly significant at high traffic load when the flow rate goes to zero because of collisions. On the other hand, the flow rate is slightly larger when $0 < d < 1$. This is because the aggregate capacity of two APs that can hear each other is slightly larger than the capacity of a single isolated AP. While the backoff overhead for an isolated AP is the mean of a random variable uniformly distributed over the contention window, it is approximately equal to the mean of the minimum of two uniformly distributed random variables when two APs compete.

7.3 Impact of user position

Finally, we derive analytical results for the mean flow transfer time $W(r)$ of a user as a function of her/his relative location in the cell, r . By symmetry, we can focus on cell 1. We denote by $f(r)$ the position of that user of cell 2 located at the edge of the exclusion region of the user of relative location r in cell 1.

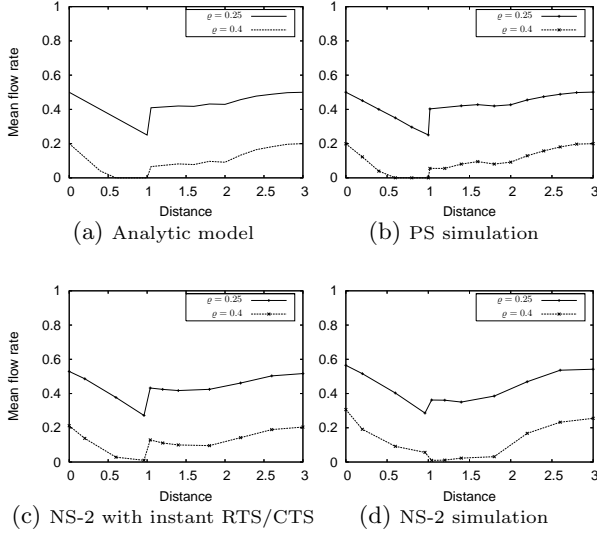


Figure 13: Impact of normalised inter-AP distance.

PROPOSITION 1. Assuming inter-cell interference is constant, as in §4.3, the mean transfer time satisfies the differential equation:

$$W'(r) = \theta f'(r)W(-f(r)), \quad (12)$$

where θ is a constant that depends on inter-AP distance d and traffic density ρ only.

The proof is in Appendix B. The function f is the upper limit of the shaded exclusion regions illustrated in Figures 4(a) to 4(d). Equation (12) can thus be solved piecewise giving the following expressions that depend on the extent of cell coupling.

(a')- Total coupling, $d \in (0, 1]$:

$$W(r) = \frac{1}{\mu} \times \frac{1}{1 - (1 + \frac{d}{2})\rho} \quad r \in [-1, \frac{d}{2}].$$

(b')- Partial coupling, overlapped transmission areas, $d \in (1, 2]$:

$$W(r) = \begin{cases} W(-1) & r \in [-1, 0] \\ W(\frac{d-1}{2}) \times \sqrt{2} \sin\left(\theta(r - \frac{d-1}{2}) + \frac{\pi}{4}\right) & r \in [0, d-1] \\ W(1) & r \in [d-1, 1] \end{cases}$$

(c')- Partial coupling, non-overlapped transmission areas, $d \in (2, 3]$:

$$W(r) = \begin{cases} W(-1) & r \in [-1, d-2] \\ W(\frac{d-1}{2}) \times \sqrt{2} \sin\left(\theta(r - \frac{d-1}{2}) + \frac{\pi}{4}\right) & r \in [d-2, 1] \end{cases}$$

(d')- Independent cells, $d \in (3, \infty)$:

$$W(r) = \frac{1}{\mu} \times \frac{1}{1 - 2\rho} \quad r \in [-1, 1].$$

Figure 14(a) plots W as a function of r for $d = 1.8$ and $d = 2.2$, corresponding to cases (b') and (c') above, respectively. The results highlight the negative impact of interference for users situated close to the centre of the network and having a larger exclusion region. Simulation results in Figures 14(b)-14(c) confirm the accuracy of the analytical




Figure 14: Impact of user position.

approximation. However because of the hidden node problem mentioned above, we observe a larger transfer time in Fig. 14(d).

8. CONCLUSION

We have proposed a model to evaluate the downlink traffic capacity of a multi-cell WLAN. The capacity is defined as the limiting traffic intensity (flow arrival rate \times mean flow size) beyond which realized flow throughput tends to zero. The model has allowed us to evaluate the capacity of some toy network configurations providing insight into the impact of inter-cell interference.

We observe an abrupt increase in capacity as APs cease to interfere directly. The impact of interference via the users to whom they transmit is less significant and decreases as the inter-AP distance increases.

For multi-channel networks with a regular pattern of frequency re-use, the variation in capacity as a function of distance produces clear local maxima and minima in achievable traffic density. Optimal AP placement corresponds to spacing neighbouring APs using the same channel by slightly more than their transmission range. Capacity rapidly decreases to a global minimum when the inter-AP spacing is small enough to bring APs into direct conflict.

With optimal AP placement, the capacity of an M -channel 2D grid network is $4M$ times that of a single channel network. That the gain is amplified by the factor of 4 is due to a contraction of the cell coverage area as APs (of all channels) are more closely spaced. This phenomenon is not an artefact of the model and is a significant observation for the design of frequency re-use in real networks.

The model allows an evaluation of the expected throughput of a download as a function of cell load and user position. Results for the simplest 2-AP network show that the system behaves broadly like a discriminatory processor sharing system. Users experiencing lower inter-cell interference gain higher throughput but the difference between best and worst positions remains slight. Performance is mainly governed by the traffic capacity that fixes the point where

throughput goes to zero.

The model makes many simplifying assumptions whose significance we now briefly discuss. The traffic model assumes Poisson flow arrivals and exponential flow sizes. These assumptions are necessary for the proof of the key stability result but we expect a reasonable degree of insensitivity to carry over from the underlying processor sharing system. We have assumed equal sharing of cell capacity between concurrent flows. This would not occur if users had significantly different round trip times, due to TCP RTT bias, for instance. This discrepancy is unlikely to affect our broad conclusions, however. We ignore the impact of upstream traffic. This is clearly a bold simplification since even download traffic generates TCP ACK packets in the upstream. We suppose these ACKs can be accounted for by extending the packet transmission time. We have chosen a simple radio propagation model and derived numerical results only for toy symmetric networks with uniform traffic. These assumptions could be removed at the cost of added complexity.

The most contestable simplifications are in the way we model the impact of interference. As noted in the introduction, the channel access process is extremely complex and clearly beyond any precise stochastic modelling approach. Our assumptions, that time is slotted and that the time to transmit a packet is proportional to the number of programmed transmissions within its exclusion zone, effectively de-couple the complex inter-cell interference allowing the analytical developments. The assumption is clearly correct when there is no interference, and is reasonable when cells interfere completely. In an average sense, it is also intuitively plausible that transmission times increase linearly with the amount of interference.

The model ignores the impact of finite RTS/CTS transmission times. We have observed that this introduces discrepancies due to the additional collisions that occur leading to an overestimate of network capacity.

In future work we intend to relax some of the above assumptions. We also mean to verify and quantify the predicted phenomena in a more realistic network configuration. The ultimate objective is to derive practical guidelines for the design and operation of dense multi-cell, multi-channel WLAN access networks. Finally, it would be interesting to evaluate the impact on performance of possible modifications to network operation. In particular, the assumed FIFO queuing in APs might be replaced by a more opportunistic scheme where a packets are chosen for transmission depending on how much interference they cause.

9. REFERENCES

- [1] M. Armony and N. Bambos. Queueing networks with interacting service resources. In *Proc. 37th Annual Allerton Conf. Commun., Control, Comp.*, 1999.
- [2] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, 2000.
- [3] T. Bonald, S. Borst, N. Hegde, and A. Proutière. Wireless data performance in multi-cell scenarios. In *SIGMETRICS'04*, pages 378–380. ACM Press, 2004.

- [4] S. Borst, M. Jonckheere, and L. Leskela. Stability of parallel queueing systems with coupled rates. *to appear in Journal of Discrete Events and Dynamic Systems*, 2007.
- [5] J. Dai. On positive Harris recurrence of multiclass queueing networks: A unified approach via fluid limit models. *Annals of Appl. Probability*, 5(1):49–77, 1995.
- [6] G. Fayolle and R. Iasnogorodski. Two coupled processors: the reduction to a Riemann-Hilbert problem. *Z. Wahr. verw. Ge*, 47(3):325–351, 1979.
- [7] G. Fayolle, I. Mitrani, and R. Iasnogorodski. Sharing a processor among many job classes. *J. ACM*, 27(3):519–532, 1980.
- [8] N. Gupta and P. R. Kumar. A performance analysis of the 802.11 wireless LAN medium access control. *Communications in Information and Systems*, 3(4):279–304, 2004.
- [9] IEEE 802.11 Standard. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, 1999.
- [10] C. Joo and N. B. Shroff. Performance of random access scheduling schemes in multi-hop wireless networks. In *INFOCOM 2007*, 2007.
- [11] A. Kumar, E. Altman, D. Miorandi, and M. Goyal. New insights from a fixed point analysis of single cell IEEE 802.11 WLANs. In *INFOCOM 2005*, 2005.
- [12] F. Lebeugle and A. Proutiere. User-level performance in WLAN hotspots. In *ITC 19*, 2005.
- [13] X. Lin and S. B. Rasool. Constant-time distributed scheduling policies for ad hoc wireless networks. In *IEEE Conference on Decision and Control, 2006*, pages 1258–1263, 2006.
- [14] R. Litjens, F. Roijers, J. Van den Berg, R. Boucherie, and M. Fleuren. Performance analysis of wireless lans: An integrated packet/flow level approach. In *ITC 18*, 2003.
- [15] S. P. Meyn. Transience of multiclass queueing networks and their fluid models. *Annals of Appl. Probability*, 5:946–957, 1995.
- [16] M. K. Panda, A. Kumar, and S. H. Srinivasan. Saturation throughput analysis of a system of interfering IEEE 802.11 WLANs. In *WoWMoM 2005*, 2005.
- [17] R. R. Rao and A. Ephremides. On the stability of interacting queues in a multiple-access system. *IEEE Trans. on Information Theory*, 34:918–930, 1988.
- [18] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. on Automatic Control*, 37(12):1936–1948, 1992.

APPENDIX

A. STABILITY

Proof of Theorem 1.

The proof proceeds by applying the fluid limit approach of Dai [5]. Let $\bar{x}_j^i(t)$ be the class- j fluid volume at AP i at time t . This represents the number of class- j ongoing flows when the flow population and the time are scaled by the

same factor, growing to infinity.

We denote by $\bar{x}_i(t)$ the total fluid volume at AP i at time t and by $\bar{\xi}_j^i(t)$ the proportion of class- j fluid volume at AP i at time t , when $\bar{x}_i(t) > 0$:

$$\bar{\xi}_j^i(t) = \frac{\bar{x}_j^i(t)}{\bar{x}_i(t)} \quad \text{with} \quad \bar{x}_i(t) = \sum_{j \in \mathcal{U}_i} \bar{x}_j^i(t).$$

It follows from the strong law of large numbers that at any time t such that $\bar{x}_i(t) > 0$ for all $i = 1, \dots, N$:

$$\frac{d\bar{x}_j^i}{dt} = \lambda_j^i - \frac{\mu}{\bar{\delta}_i(t)} \frac{\bar{x}_j^i(t)}{\bar{x}_i(t)},$$

with

$$\bar{\delta}_i(t) = \sum_{j \in \mathcal{U}_i} \bar{\xi}_j^i(t) \bar{\delta}_j^i(t)$$

and

$$\bar{\delta}_j^i(t) = 1 + \sum_{k \neq i} \sum_{l \in \mathcal{U}_k} \bar{\xi}_l^k(t) \chi(j, l).$$

We first prove that $\bar{\xi}_j^i(t)$ tends to α_j^i , the proportion of class- j traffic at AP i , for all $j \in \mathcal{U}_i$. Note that:

$$\frac{d\bar{x}_j^i}{dt} = \lambda_j^i > 0 \quad \text{if} \quad \bar{x}_j^i(t) = 0 \quad \text{and} \quad \bar{x}_i(t) > 0.$$

Now for all $j, l \in \mathcal{U}_i$, we have at any time t such that $\bar{x}_j^i(t) > 0$ and $\bar{x}_l^i(t) > 0$:

$$\frac{d}{dt} \ln \left(\frac{\bar{x}_j^i(t)}{\bar{x}_l^i(t)} \right) = \frac{d\bar{x}_j^i}{dt} \frac{1}{\bar{x}_j^i(t)} - \frac{d\bar{x}_l^i}{dt} \frac{1}{\bar{x}_l^i(t)} = \frac{\lambda_j^i}{\bar{x}_j^i(t)} - \frac{\lambda_l^i}{\bar{x}_l^i(t)}.$$

Thus the ratio $\bar{x}_j^i(t)/\bar{x}_l^i(t)$ decreases if and only if it is larger than λ_j^i/λ_l^i . Since $\bar{\xi}_j^i(t)/\bar{\xi}_l^i(t) = \bar{x}_j^i(t)/\bar{x}_l^i(t)$ and $\alpha_j^i/\alpha_l^i = \lambda_j^i/\lambda_l^i$, the ratio $\bar{\xi}_j^i(t)/\bar{\xi}_l^i(t)$ decreases if and only if it is larger than α_j^i/α_l^i . Using the fact that:

$$\sum_{j \in \mathcal{U}_i} \bar{\xi}_j^i(t) = \sum_{j \in \mathcal{U}_i} \alpha_j^i = 1,$$

we deduce that $\bar{\xi}_j^i(t)$ tends to α_j^i for all $j \in \mathcal{U}_i$ when t tends to infinity.

In view of (4) and (5), this in turn implies that $\bar{\delta}_j^i(t)$ tends to β_j^i for all $j \in \mathcal{U}_i$ and that $\bar{\delta}_i(t)$ tends to $\sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i$ when t tends to infinity. Thus for all $\varepsilon > 0$, we have for sufficiently large t such that $\bar{x}_i(t) > 0$ for all $i = 1, \dots, N$:

$$\frac{d\bar{x}_j^i}{dt} \leq \lambda_j^i - \mu \frac{\alpha_j^i}{\sum_{l \in \mathcal{U}_i} \alpha_l^i \beta_l^i} (1 - \varepsilon).$$

Note that this inequality holds even if $\bar{x}_k(t) = 0$ for some $k \neq i$, since this can only increase the service rate of AP i .

Define the fluid workload of AP i at time t as:

$$\bar{w}_i(t) = \sum_{j \in \mathcal{U}_i} \frac{\bar{x}_j^i(t)}{\mu} \beta_j^i.$$

We have for sufficiently large t such that $\bar{w}_i(t) > 0$:

$$\frac{d\bar{w}_i}{dt} \leq \varrho_i \sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i - 1 + \varepsilon,$$

which is negative for sufficiently small ε . Thus $\bar{w}_i(t) = 0$ and $\bar{x}_j^i(t) = 0$ for all $j \in \mathcal{U}_i$ for sufficiently large t . This property holds for all APs: the fluid model is stable, which implies the ergodicity of the underlying Markov process [5]. \square

Proof of Theorem 2.

We show that if $\varrho_i \sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i > 1$ for all $i = 1, \dots, N$, the fluid model introduced in the proof of Theorem 1 is unstable. We prove in a similar way that $\bar{\xi}_j^i(t)$ tends to α_j^i and $\bar{\delta}_i(t)$ tends to $\sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i$ when t tends to infinity. Thus for all $\varepsilon > 0$, we have for sufficiently large t such that $\bar{x}_i(t) > 0$ for all $i = 1, \dots, N$:

$$\frac{d\bar{x}_j^i}{dt} \geq \lambda_j^i - \mu \frac{\alpha_j^i}{\sum_{l \in \mathcal{U}_i} \alpha_l^i \beta_l^i} (1 + \varepsilon)$$

and

$$\frac{d\bar{w}_i}{dt} \geq \varrho^i \sum_{j \in \mathcal{U}_i} \alpha_j^i \beta_j^i - 1 - \varepsilon,$$

which is positive for all $i = 1, \dots, N$ for sufficiently small ε . Thus $\bar{w}_i(t)$ increases at least linearly for all $i = 1, \dots, N$: the fluid model is unstable, which implies the transience of the underlying Markov process [15]. \square

B. PERFORMANCE

Proof of Proposition 1.

Cell 1 corresponds to the interval $[a, b]$ with $a = -\frac{d}{2} - 1$ and $b = \min(-\frac{d}{2} + 1, 0)$. Equations (8), (9) and (10) applied to cell 1 have the following continuous counterparts:

$$X(r) = \lambda W(r),$$

$$W(r) \left(1 - \frac{\varrho}{2} \int_a^b \phi(r) dr \right) - \frac{\varrho}{2} \int_a^b \phi(r) W(r) dr = \frac{\phi(r)}{\mu},$$

$$\phi(r) = 1 + \varrho \left(\int_a^b \phi(r) dr \right) \frac{\int_{-f(r)}^b X(r) dr}{\int_a^b X(r) dr},$$

where $\lambda = \varrho\mu$ is the flow arrival density and the last equality follows by symmetry. Let \bar{X} and $\bar{\phi}$ be the integrals of X and ϕ over $[a, b]$. Differentiating the above equations with respect to r , we obtain:

$$X'(r) = \lambda W'(r),$$

$$\left(1 - \frac{\varrho}{2} \bar{\phi} \right) W'(r) = \frac{\phi'(r)}{\mu},$$

$$\phi'(r) = \frac{\varrho \bar{\phi}}{\bar{X}} f'(r) X(-f(r)).$$

We deduce (12), with

$$\theta = \frac{\varrho^2 \bar{\phi}}{\bar{X} \left(1 - \frac{\varrho \bar{\phi}}{2} \right)}.$$

\square