A Traffic-Oriented Approach for Channel Assignment in WLAN

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

***Abstract*—The IEEE 802.11 architecture offers multiple independent frequency channels. Efficient exploitation of the available frequency spectrum is one of the most promising approaches to increase network capacity. In this paper we propose an heuristic evaluation of network capacity that takes account of the dynamic nature of traffic. We use this method to derive static frequency allocations that maximize capacity. Numerical evaluations show our method can provide significant gains compared to the widely adopted received signal strength approach.**

I. Introduction and related work

The traffic capacity of a multi-cell IEEE 802.11 WLAN depends on the way frequencies are allocated to neighbouring Access Points (APs). In this paper we propose a method to compute a static allocation that maximizes capacity under an assumed dynamic traffic model. Users arrive randomly in time and space and associate with a certain AP. Each user requires a random amount of bits to be transferred by the network and leaves the network upon service termination. Traffic capacity is defined as the maximum load (arrival rate $\times$ flow size) under which the number of active users remains finite. We consider the problem from the point of view of a network operator having a complete knowledge of the position of APs and their transmission range and the spatial traffic distribution.

Previous approaches to the channel assignment problem used received power from neighbouring APs as a measure of channel quality [1], [2]. However, received power cannot capture all types of interference in a WLAN. Recent studies [3], [4], [5] have pointed out this problem. The authors propose a dynamic channel assignment that accounts for the numbers of mutually interfering users. The proposed solution is a variant of the classical graph coloring problem called weighted graph coloring. The idea is to assign a weight, typically a number between 0 and 1, to each pair of neighbouring APs that captures essentially the mutual number of interfering clients. This correctly accounts for interference. However it may not be suitable for a highly dynamic environment where users continuously enter and leave the network since it requires the exchange of a large amount of control information to maintain an up-to-date view of the network topology.
Our approach is to determine a static allocation that maximizes capacity for an expected traffic density profile. Other studies have tried to incorporate the notion of traffic demand [6]. However the capacity of each AP is assumed to be fixed and this does not account for interference between APs that share the same channel. A distributed channel allocation technique is proposed in [7] where each AP broadcasts the number of associated users to neighbouring APs. The AP then computes the number of neighbouring clients on each frequency channel and selects the channel that ensures the highest per-user throughput. The algorithm requires periodic coordination between APs which is hard to achieve since the APs may not be able to hear each other. A distributed scheme that requires only local measurement of channel quality is proposed in [8]. The authors show that the algorithm converges to a "good" channel assignment provided one exists, i.e., there is a sufficient number of frequencies. A generalisation to channel dependent interference is proposed in [9].

Our approach differs in that we seek a static allocation that maximizes capacity for a given spatial traffic distribution. Allocations would be changed only to reflect changing demand profiles and not in response to highly dynamic changes in the instantaneous user population. This proposed method only requires knowledge about the position of the APs and their transmission range and takes proper account of interference. It is derived from our previous work on an analytical model of a multi-cell WLAN. A significant original contribution is an heuristic based on a fluid model that allows an estimation of the traffic capacity of a general network topology.
To clarify the issues, we begin in Section II by illustrating the interference phenomena arising in a small scale example network. In Section III we present our model. An approximate method to compute network capacity is given in Section IV. Numerical results are presented in Section VI. We conclude the paper in Section VII.

II. A simple example

To understand how the IEEE 802.11 allocates transmission opportunities in a multi-cell WLAN we consider a simple network consisting of two APs and four user stations. The distance between the two APs is d. The transmission range is denoted by R. Let T be the throughput of the wireless link after accounting for all overheads: SIFS, DIFS, RTS/CTS, ACK and the random backoff time. T is the bit rate that would be achieved by an AP in the absence of any type of interference. Let L denote the mean packet length in bits. A transmission without interference would consume on average L/T seconds to complete.

User stations A and B (resp. C and D) are associated with AP 1 (resp. AP 2). Let p_A and p_B (resp. p_C and p_D) denote the fractions of AP 1 (resp. AP 2) packets that are destined to A and B (resp. C and D) with $p_A + p_B = 1$ (resp. $p_C + p_D = 1$). We consider four cases depending on the distance d:

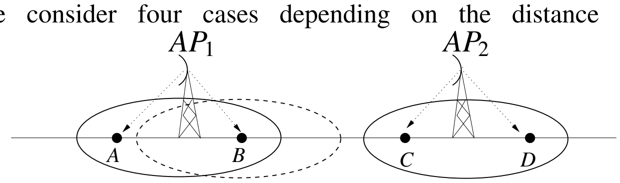

Fig. 1. Case 1, $d > 3R$.

Case 1, $d > 3R$: The distance between B and C is greater than the transmission range, therefore a CTS (if RTS/CTS is used) or an ACK sent by B will not interfere with DATA transmission sent from AP 2 to C. This means that simultaneous transmissions to B and C are feasible. The two APs are sufficiently distant from each other to enable complete spatial reuse. The rates ϕ_A, ϕ_B, ϕ_C and ϕ_D obtained by user stations A, B, C and D depend on the fraction of packets scheduled by the APs as follows:

$$\phi_A = p_A T, \quad \phi_B = p_B T, \quad \phi_C = p_C T, \quad \phi_D = p_D T.$$

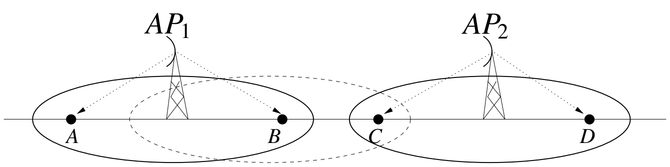

Fig. 2. Case 2, $2R < d < 3R$.

Case 2, $2R < d < 3R$: The distance between C and B is less than R. Therefore either C will be blocked by the CTS sent from B (if RTS/CTS is used) or an ACK from B will cause DATA collisions at C. In both cases simultaneous transmissions to C and B are not feasible while simultaneous transmissions to other destinations are. The rate obtained by each station is now more difficult to analyse.

Remarque: Note that this type of interference is not modelled by the classical models of Bianchi [10] and Kumar et al. [11]. These models consider only networks of single collision domains where simultaneous transmissions are not feasible. The Bianchi and Kumar models compute only the collision probability that results from equal values of backoff counters. In a multi-cell WLAN this type of collision would be very rare because only a small number of APs are within the hearing range of each other. Thus we cannot use these models to compute the rate obtained by each station for the network of Figure 2. In this type of network the most important resource allocation mechanism is how simultaneous transmissions are scheduled.

In recent work [12] we proposed an analytic model to characterise the performance of a multi-cell WLAN. The basic idea of this model is as follows. Suppose that APs 1 and 2 are contending for channel access and that AP 1 wants to transmit a packet to B. If the destination of AP2's packet is D, the transmission to B needs L/T seconds. On the other hand, if the destination is C, then on average the transmission time will be $2L/T$ because B and C interfere with each other. Thus, on average, AP 1 needs

$$p_D(L/T) + p_C 2(L/T) = (1 + p_C)(L/T), \text{ seconds}$$

to complete. The mean transmission time is increased by a factor $(1 + p_C)$. This factor is due to interference from AP 2. Since AP 1 transmits to A and B with probabilities p_A and p_B then the mean time needed to transmit a packet from AP 1 is

$$p_A(L/T) + p_B(1 + p_C)(L/T) = (1 + p_B p_C)(L/T).$$

The overall mean transmission time of AP 1 is increased by a factor $(1+p_B p_C)$. This factor models the effect of interference between the two cells. The throughput obtained by each user station A and B may then be written

$$\phi_A = \frac{p_A}{(1 + p_B p_C)} T, \quad \phi_B = \frac{p_B}{(1 + p_B p_C)} T.$$

Similarly the throughputs of stations C and D are

$$\phi_C = \frac{p_C}{(1 + p_B p_C)} T, \quad \phi_D = \frac{p_D}{(1 + p_B p_C)} T.$$

Note that T appears only as a multiplicative factor for the rate allocated to user stations. Thus without loss of generality it can be assumed to be equal to 1. ϕ_A and ϕ_B correspond in this case to a normalised throughput with the normalising factor being the throughput of the physical link. The above analysis yields simple tractable formulas for the throughput allocated to each user station in a multi-cell WLAN. Such formulas cannot be obtained by either Bianchi or Kumar models because they do not consider the possibility of more than a single transmission in the network.

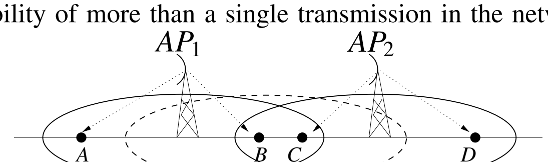

Fig. 3. Case 3, $R < d < 2R$.

Case 3, $R < d < 2R$: AP 2 is now within the range of B. If the RTS/CTS scheme is used, AP 2 is blocked when AP 1 is transmitting to B. Thus only transmissions to A and D are allowed simultaneously. If RTS/CTS is not used, simultaneous transmissions will not succeed unless they are destined to A and D respectively. To simplify, we assume RTS/CTS is used. We proceed as in the previous case. We compute the mean time required to transmit a packet to each destination. For A a transmission will take L/T if AP 2 is transmitting to D and $2L/T$ if AP 2 is transmitting to C. For B, a transmission will take on average $2L/T$ s because B interfere with both C and D. The mean time required to transmit a packet by AP 1 is therefore

$$p_A(p_D(L/T){+}p_C 2(L/T)){+}p_B 2(L/T) = (1{+}p_B{+}p_A p_C)(L/T).$$

The throughputs of stations A and B in this case are

$$\phi_A = \frac{p_A}{1+p_B+p_Ap_C}T, \quad \phi_B = \frac{p_B}{1+p_B+p_Ap_C}T.$$

The throughput of AP 1 is therefore reduced by a factor $(1+p_B+p_Ap_C)$ due to interference from AP 2.

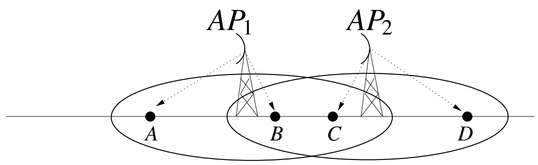

Fig. 4. Case 4, $d < R$.

Case 4, $d < R$: The APs are now within the range of each other. All destinations interfere with each other. The mean time required to transmit a packet is $2L/T$, the throughput of user stations are therefore given by

$$\phi_A = \frac{p_A}{2}T, \quad \phi_B = \frac{p_B}{2}T, \quad \phi_C = \frac{p_C}{2}T, \quad \phi_D = \frac{p_D}{2}T.$$

In cases 1 and 4, the received signal power correctly captures the interference between the two APs. However, it fails in cases 2 and 3. Thus a signal-based channel allocation scheme would allocate the same frequency channel to APs 1 and 2 even though they interfere.

The signal-based approach allows for a simple distributed implementation of channel assignment at the expense of performance loss due to interferences of type 2 and 3. For centralised deployments, these performance losses can be recovered by choosing an appropriate metric that accounts for all types of interference. A major contribution of this paper is a metric that effectively models all these types of interference.

Another problem with the signal-based approach is that it is unaware of the traffic demands of APs. The channel allocation depends only on how well the APs are able to detect each other. We consider on the contrary that a good channel assignment should integrate the notion of traffic demand and should optimise the performance obtained by end users. Another contribution of this paper is to define an objective function that accounts for the performance observed by end users and a method to compute this function.

Our approach is complementary to channel allocation methods proposed for use in cellular networks [13], [14], [15], [16]. These studies propose to use heuristic optimisation techniques, such as the genetic algorithm or simulated annealing, to solve underlying combinatorial optimisation aspects of the channel allocation problem. All these solutions can be used to optimise our objective function. The contribution of this paper is not to propose another optimisation heuristic but rather to define an interference metric and an objective function that accounts for the correct type of interference and the performance seen by end users.

In the remainder of this paper, we simply generalise the above ideas and formalise them within a mathematical framework that allows the analysis of arbitrary network topologies.

III. The model

A. Network and traffic

The network consists of a set $\mathcal{A}$ of APs. With each AP is associated a set of user classes. A user class represents a set of users having the same interference characteristics. Typically, the class is composed of users located within a small geographical area. We denote by $\mathcal{U}_i$ the set of user classes associated with AP i and by $\mathcal{U}$ the set of all user classes. We characterise interference between user classes $j \in \mathcal{U}_i$ and $l \in \mathcal{U}_k$ by a function χ from $\mathcal{U} \times \mathcal{U}$ to $\{0,1\}$ such that $\chi(j,l) = 1$ if and only if transmissions to j and l cannot occur simultaneously.

Class-j traffic is assumed to arrive according to a Poisson process with traffic intensity λ_j $(flows/s)$ and flow sizes are assumed to be exponentially distributed with mean flow size $1/\mu$ $(bits/flow)$. We define the traffic intensity of class-j as $\rho_j = \lambda_j/\mu$ (bits/s).

We consider M independent channels. Each AP is equipped with a single radio interface. A channel assignment θ is a function from $\mathcal{A}$ to $\{1, ..., M\}$ that assigns a frequency channel to each AP. We denote by θ_i the frequency assigned to AP i.

B. Rate functions

We use x_j to denote the number of class-j users and we refer to the vector $(x_j)_{j\in\mathcal{U}}$ as the network state. The AP transmits an equal number of packets to each user. Therefore, the fraction ξ_j of AP i packets that are destined to class j users can be written as

$$\xi_j = \frac{x_j}{\sum_{l\in\mathcal{U}_i} x_l}, \quad j \in \mathcal{U}_i.$$

Note that $\sum_{j\in\mathcal{U}_i} \xi_j = 1$. The vectors $(\xi_j)_{j\in\mathcal{U}_i}$, $i \in \mathcal{A}$ are the equivalent of the probabilities p_A, p_B, p_C and p_D of Section II and they determine how the APs interact with each others.

From the examples of Section II we observe also that for each couple of interfering stations, the average packet transmission time of a particular AP is increased by a factor equal to the product of their packet distribution probabilities. In case 2, only B and C interfere, the transmission time of AP 1 is increased by the factor p_Bp_C. In case 3, A interferes with C and B interferes with C and D. The packet transmission time is increased by $p_Ap_C + p_Bp_C + p_Bp_D = p_B + p_Ap_C$ because $p_C + p_D = 1$.

Thus the mean transmission time of AP i would be increased by the factor

$$\sum_{l\in\mathcal{U}_i}\sum_{n\in\mathcal{U}_k} \xi_l\xi_n\chi(l,n)$$

to account for interference from AP k. To account for interference from all other APs we simply add these factors. Also, since

$$\sum_{l\in\mathcal{U}_i}\sum_{n\in\mathcal{U}_i} \xi_l\xi_n\chi(l,n) = 1,$$

the rate allocated to each user class j can be written as

$$\phi_j = \frac{\xi_j}{\sum_{k\in\mathcal{A}}\sum_{l\in\mathcal{U}_i}\sum_{n\in\mathcal{U}_k} \xi_l\xi_n\chi(l,n)}. \quad (1)$$

The denominator in (1) can be interpreted as the mean time required by AP i to transmit a packet. Its reciprocal is the capacity (in *bits/s*) of AP i in state $(x_j)_{j\in\mathcal{U}}$. The numerator ξ_j indicates that the capacity of AP i is shared equally between all its active users. The service rates are readily generalized to a multi-channel network with allocation θ: the factor $\chi(l,n)$ in (1) must be replaced by $\chi(l,n)\mathbb{I}_{\{\theta_i=\theta_k\}}$ where $\mathbb{I}$ is the indicator function.

Thus the number of users in each class behaves like the population of a multiclass queueing system with state the dependent service rates $\phi_j,\ j\in\mathcal{U}$.

C. Objective function

By definition, the capacity region of the network, with respect to the channel allocation θ, is the set of traffic intensities $(\rho_j)_{j\in\mathcal{U}}$ such that the network is stable, i.e., the number of users in all classes remains finite. To measure capacity, we assume intensities are proportional to a vector $(\eta_j)_{j\in\mathcal{U}}\in\mathbb{R}^{|\mathcal{U}|}$ and define

$$\gamma^*=\sup\{\gamma\in\mathbb{R}^+,\gamma(\eta_j)_{j\in\mathcal{U}}\text{ is inside the capacity region}\},$$

Thus γ^* measures the relative efficiency of an assignment and our objective can be formulated as follows

$$\begin{cases}\text{maximise} & \gamma*\\ \text{subject to} & \theta\in\Theta,\end{cases}\tag{2}$$

where Θ denotes the set of all possible channel allocations. By choosing γ^* as an objective function for the channel allocation problem our model takes into account appropriately both the traffic demand of each cell and interference between cells. The remaining difficulty is how to compute γ^* for arbitrary network configurations.

In [12], it has been shown that a *sufficient* condition for network stability is

$$\sum_{j\in\mathcal{U}_i}\rho_j<\frac{1}{\sum_{k\in\mathcal{A}}\sum_{l\in\mathcal{U}_i}\sum_{n\in\mathcal{U}_k}\alpha_l\alpha_n\chi(l,n)},\tag{3}$$

where

$$\alpha_j=\frac{\rho_j}{\sum_{l\in\mathcal{U}_i}\rho_l},\qquad j\in\mathcal{U}_i.$$

The left hand side of (3) is the traffic offered in cell i while the right hand side can be interpreted as the cell capacity. The network is therefore stable if the offered traffic is less than the capacity. The capacity of the cell depends mainly on the traffic distribution and interference characteristics of nearby cells. It is also shown in [12] that the above condition is *necessary* for symmetric networks. However this formula is not valid for asymmetric networks. One of the major contributions of this paper is to provide a simple polynomial time algorithm that can compute a good approximation for γ^* and is valid for an arbitrary network.

D. Interference metric

For $i,k\in\mathcal{A}$ define

$$I_{ik}=\sum_{l\in\mathcal{U}_i}\sum_{n\in\mathcal{U}_k}\alpha_l\alpha_n\chi(l,n).\tag{4}$$

The term I_{ik} can be seen as a measure of interference between APs i and k. These factors abstract the complex interaction between user stations in a multi-cell WLAN.

If A,B,C and D of Section II represent user classes (not simple user stations) with traffic intensities ρ_A,ρ_B,ρ_C and ρ_D then the interference metrics I_{ik} represented by the matrix

$$\begin{pmatrix} I_{11} & I_{12}\\ I_{21} & I_{22}\end{pmatrix},$$

would corresponds to:
Case 1:

$$\begin{pmatrix}1 & 0\\ 0 & 1\end{pmatrix}$$

Case 2:

$$\begin{pmatrix}1 & \frac{\rho_B}{\rho_A+\rho_B}\frac{\rho_C}{\rho_C+\rho_D}\\ \frac{\rho_B}{\rho_A+\rho_B}\frac{\rho_C}{\rho_C+\rho_D} & 1\end{pmatrix}$$

Case 3:

$$\begin{pmatrix}1 & \frac{\rho_B}{\rho_A+\rho_B}+\frac{\rho_A}{\rho_A+\rho_B}\frac{\rho_C}{\rho_C+\rho_D}\\ \frac{\rho_C}{\rho_C+\rho_C}+\frac{\rho_B}{\rho_A+\rho_B}\frac{\rho_D}{\rho_C+\rho_D} & 1\end{pmatrix}$$

Case 4:

$$\begin{pmatrix}1 & 1\\ 1 & 1\end{pmatrix}$$

In the next section we propose a method to compute γ^* using only the traffic intensity vector $(\rho_j)_{j\in U}$ and the interference metrics I_{ik}.

IV. Approximating traffic capacity

Stability condition (3) was derived using the fluid limit approach of Dai [17]. Unfortunately, to apply this approach to derive the precise capacity region of general multi-class networks proves problematic. The difficulty resides in computing fluid service rates when a subset of user classes have empty fluid volumes. This requires the joint distribution of a coupled queueing system that is generally extremely hard to characterise.

In this section we therefore propose an heuristic to evaluate γ^*. This yields exact values for symmetric networks and for networks with only two cells but it is only an approximation for other cases. We replace the model of [17] by an approximate fluid model. Let $\widehat{x}_j$ denote the fluid volume of class j in this model. The approximate fluid model is constructed as follows.

We divide time into a succession of intervals of length dt. At the beginning of each interval, a fluid quantity $\rho_j dt$ is added to the fluid volume of class j for all classes. During the remaining part of the interval, fluid volumes are drained at the service rates $\phi_j(\widehat{x})$ as given by (1) with $\widehat{x}=(\widehat{x}_j)_{j\in\mathcal{U}}$. The approximate fluid model is, by definition, the limiting model when dt tends to 0. This model is identical to that of Dai [17]

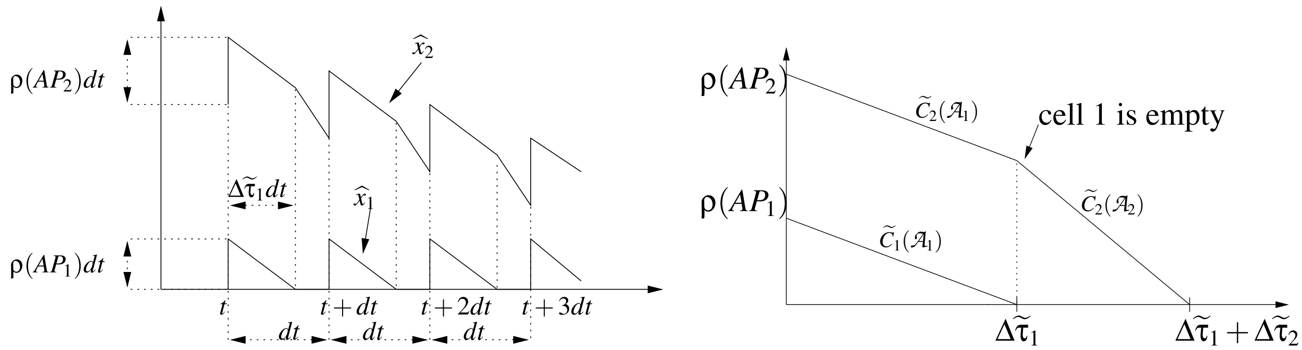

(a) With arrivals.

(b) Without arrivals.

Fig. 5. Approximate fluid model.

if all user classes are non empty (in the sense that residual fluid prior to the injection of the $\rho_j dt$ is non-zero). However, it is only an approximation when a subset of classes are empty. We replace the complex stochastic behaviour of empty classes by a simpler deterministic one.

Figure 5(a) illustrates the operation of the approximate fluid model for a two-cell network. Cell 1 is empty while cell 2 has a positive fluid volume. At time t, infinitesimal fluid volumes $\rho(AP_1)dt$ and $\rho(AP_2)dt$ are added to cells 1 and 2, respectively. Between t and $t + \Delta\widetilde{\tau}_1 dt$, cells 1 and 2 are drained at rates denoted $\widetilde{C}_1(\mathcal{A})$ and $\widetilde{C}_2(\mathcal{A})$. These rates reflect the mutual interference of cells 1 and 2. At time $t + \Delta\widetilde{\tau}_1 dt$, cell 1 empties and the service rate of cell 2 increases to $\widetilde{C}_2(\{2\})$ until the end of the interval. Cell 2 service capacity in this approximate fluid model is $\widehat{C}_2(\{2\})$, the average over the interval of rates $\widetilde{C}_2(\mathcal{A})$ and $\widetilde{C}_2(\{2\})$. The depicted instantiation is stable since the derivative of $\widehat{x}_2$ is negative.

The stability condition of this approximate fluid model is in fact much easier to compute than that of the exact fluid model as stated in the theorem below.

Theorem 1: Consider the approximate fluid model with no traffic arrivals, i.e.

$$\frac{d\widetilde{x}_j(t)}{dt} = -\phi_j(\widetilde{x}(t)), \tag{5}$$

and with initial fluid volumes equal to the traffic intensities, i.e. $\widetilde{x}(0) = (\rho_j)_{j\in\mathcal{U}}$. Let $\widetilde{\tau} = \inf\{t \in \mathbb{R}^+, \widetilde{x}(t) = 0\}$. Then the approximate fluid model is stable iff $\widetilde{\tau} < 1$.

It can easily be verified that if we scale the traffic intensity vector by a positive factor γ, then $\widetilde{\tau}$ is scaled by the same factor. It follows that if we start with the traffic intensity vector $\gamma^*(\rho_j)_{j\in\mathcal{U}}$ then the time $\widetilde{\tau} = 1$. We deduce that

$$\gamma^* = 1/\widetilde{\tau}. \tag{6}$$

. An outline of the proof of Theorem 1 is given below.

Algorithm 1 summarises the steps to compute $\widetilde{\tau}$. The algorithm is initialised in steps 1, 2 and 3. The time $\widetilde{\tau}$ is initialised to 0. The set of non-empty cells $\mathcal{A}_*$ is initialised with $\mathcal{A}$. The remaining workload w_i at cell i is assigned the total traffic offered to cell i. The algorithm then iterates until all cells become empty. During each iteration at least one of the cells become empty. The cell that empties is the one that has the minimum ratio of w_i/C_i. The capacities C_i are

Algorithm 1 Calculate $\widetilde{\tau}$

1: $\widetilde{\tau} \leftarrow 0$;
2: $\mathcal{A}_* \leftarrow \mathcal{A}$;
3: **for each** $i \in \mathcal{A}$, $w_i \leftarrow \sum_{j\in\mathcal{U}_i} \rho_j$;
4: **while** $\mathcal{A}_* \neq \emptyset$ **do**
5:     **for each** $i \in \mathcal{A}_*$, $C_i \leftarrow \frac{1}{\sum_{k\in\mathcal{A}_*} I_{ik}\mathbb{I}_{\{\theta_i=\theta_k\}}}$;
6:     $\Delta\widetilde{\tau} \leftarrow \min_{i\in\mathcal{A}_*}\left(\frac{w_i}{C_i}\right)$;
7:     $i_\emptyset \leftarrow argmin_{i\in\mathcal{A}_*}\left(\frac{w_i}{C_i}\right)$;
8:     **for each** $i \in \mathcal{A}_*$, $w_i \leftarrow w_i - C_i\Delta\widetilde{\tau}$;
9:     $\widetilde{\tau} \leftarrow \widetilde{\tau} + \Delta\widetilde{\tau}$;
10:     $\mathcal{A}_* \leftarrow \mathcal{A}_* \backslash \{i_\emptyset\}$;
11: **end while**

updated at the begining of each iteration in step 5 to take into account the fact that only those cells in $\mathcal{A}_*$ are active. The algorithm then determines in steps 6 and 7 which cell in $\mathcal{A}_*$ empties first as well as the needed time to become empty $\Delta\widetilde{\tau}$. The remaining workloads w_i are updated at the end of each iteration to reflect that all cells in $\mathcal{A}_*$ have been drained during the time interval $\Delta\widetilde{\tau}$.

For the network of Section II, Algorithm 1 would return the following values for γ^*

$$\gamma^* = \begin{cases} \frac{1}{(\rho_A+\rho_B)I_{12}+\rho_C+\rho_D}, & \text{if } \rho_A + \rho_B < \rho_C + \rho_D, \\ \frac{1}{(\rho_A+\rho_B+(\rho_C+\rho_D)I_{12}}, & \text{otherwise,} \end{cases}$$

where I_{12} is given by one of the Equations of Section III-D.

A. Outline proof of Theorem 1

Let

$$\rho(AP_i) = \sum_{j\in\mathcal{U}_i} \rho_j, \qquad i \in \mathcal{A},$$

denote the traffic offered to cell i. For each $\mathcal{B} \subset \mathcal{A}$ let

$$\widetilde{C}_i(\mathcal{B}) = \frac{1}{\sum_{k\in\mathcal{B}} I_{ik}\mathbb{I}_{\{\theta_i=\theta_k\}}}, \qquad i \in \mathcal{B},$$

be the capacity of cell i when the set of non-empty cells is $\mathcal{B}$ and there is no traffic arrival. Without loss of generality we assume that the cells in the fluid model with no arrivals empty in the order $1, 2, ..., N$. Otherwise the cells can be relabeled to obtain that order. Let $\mathcal{A}_i = \{i, ..., N\}$ and let $\Delta\widetilde{\tau}_i$ be such that $\sum_{k=1}^{i} \Delta\widetilde{\tau}_k$ is the instant when cell i become empty. Within each sub-interval of length $\Delta\widetilde{\tau}_k$, AP i transmits at rate $\widetilde{C}_i(\mathcal{A}_k)$. Since cell i empties at the end of the i^{th} interval and it began with an initial fluid volume equal to $\rho(AP_i)$ we have

$$\rho(AP_i) = \sum_{k=1}^{i} \Delta\widetilde{\tau}_k \widetilde{C}_i(\mathcal{A}_k), \qquad i \in \mathcal{A}. \tag{7}$$

The times $\sum_{k=1}^{i} \Delta\widetilde{\tau}_k$ when cells empty can be interpreted as measures of congestion. Congested cells with large traffic volumes and/or small capacities need much more time to evacuate their traffic than cells with less traffic or higher capacities.

Figure 5(b) illustrates the operation of the approximate fluid model with no arrivals for the two-cell network scenario. Cells 1 and 2 have initial fluid volumes at time $t = 0$ equal to their traffic intensities $\rho(AP_1)$ and $\rho(AP_2)$. Up to time $\Delta\widetilde{\tau}_1$ when cell 1 becomes empty, the capacities of cells 1 and 2 are $\widetilde{C}_1(\mathcal{A}_1)$ and $\widetilde{C}_2(\mathcal{A}_1)$. After $\Delta\widetilde{\tau}_1$, cell 2 has a larger capacity $\widetilde{C}_2(\mathcal{A}_2) > \widetilde{C}_2(\mathcal{A}_1)$. It becomes empty at time $\Delta\widetilde{\tau}_1 + \Delta\widetilde{\tau}_2$.

In the remaining part of this section we outline the proof of Theorem 1. To prove the stability (resp. instability) of the approximate fluid model we need to show that the sum of fluid volumes tends to 0 (resp. to $+\infty$) when time tends to $+\infty$ and that independently of the initial fluid volume. The next lemma shows that it is sufficient to start with an initial fluid volume proportional to the vector of traffic intensities.

Lemma 1: To prove stability with any initial conditions, it is sufficient to prove stability with the specific initial values:

$$\widehat{x}_j(0) = \rho_j \Delta, \quad j \in \mathcal{U} \quad \text{and} \quad \Delta \in \mathbb{R}^+.$$

We omit the proof due to space limitations. Thus, without loss of generality, it is sufficient to prove Theorem 1 with an initial condition as stated in the previous lemma. In the following two propositions we characterise the capacity of the approximate fluid model when a subset of cells are empty.

Proposition 1: Suppose $\widetilde{\tau} < 1$ and cells $1, ..., i-1$ are empty. The capacity $\widehat{C}_k(\mathcal{A}_i)$ of cell $k \in \mathcal{A}_i$ with traffic arrival is

$$\widehat{C}_k(\mathcal{A}_i) = \sum_{m=1}^{i-1} \Delta\widetilde{\tau}_m \widetilde{C}_k(\mathcal{A}_m) + \Big(1 - \sum_{m=1}^{i-1} \Delta\widetilde{\tau}_m\Big) \widetilde{C}_k(\mathcal{A}_i). \tag{8}$$

Proposition 2: Suppose $\widetilde{\tau} > 1$ and cells $1, ..., i-1$ are empty. Let i^* be the largest integer such that $i^* \leq i$ and $\sum_{m=1}^{i^*-1} \Delta\widetilde{\tau}_m < 1$. Then the capacity $\widehat{C}_k(\mathcal{A}_i)$ of cell $k \in \mathcal{A}_i$ with traffic arrivals is

$$\widehat{C}_k(\mathcal{A}_i) = \sum_{m=1}^{i^*-1} \Delta\widetilde{\tau}_m \widetilde{C}_k(\mathcal{A}_m) + \Big(1 - \sum_{m=1}^{i^*-1} \Delta\widetilde{\tau}_m\Big) \widetilde{C}_k(\mathcal{A}_{i^*}). \tag{9}$$

Proposition 1 means that if there is a subset of empty cells, for instance cells $1, ..., i-1$, then the capacity at which fluid is drained in any particular non-empty cell k is a mean of the capacities of cell k assuming only cells $1, ..., m$ are empty for $0 \leq m \leq i-1$. Proposition 2 is similar to Proposition 1. However the mean should be taken assuming only cells $1, ..., m$ are empty for $0 \leq m \leq i^* - 1$.

Outline proof of Theorem 1: Let $\widehat{w}_i(t) = \sum_{j \in \mathcal{U}_i} \widehat{x}_j(t)$ denote the workload of cell i and let $\widehat{x}_j(0) = \rho_j$, i.e., we set $\Delta = 1$ as it can be readily verified that if $\Delta \neq 1$, then all affected quantities will be scaled by Δ.

Case $\widetilde{\tau} < 1$: we prove by recurrence on i, using Equation (7) and Proposition 1, that if cells $1, ..., i-1$ are empty then the next cell to become empty is cell i.

Case $\widetilde{\tau} > 1$: let i_0 denote the largest integer such that $\sum_{m=1}^{i_0-1} \Delta\widetilde{\tau}_m < 1$. We prove that for cell $k \in \mathcal{A}_{i_0}$ we have

$$\frac{dw_k(t)}{dt} = (\sum_{m=1}^{i_0} \Delta\widetilde{\tau}_m - 1) \widetilde{C}_k(\mathcal{A}_{i_0}) > 0.$$

Thus $\lim_{t\to} w_k(t) = +\infty$ and the system is unstable since none of the cells $i_0, ..., N$ will become empty.

■

V. Networks in $\mathbb{R}^2$

In the remainder of this paper we focus on the particular case of networks in the $\mathbb{R}^2$ space. Traffic is assumed to be uniformly distributed and user classes are now points in the plane so that sums in (3) are replaced by integrals (see [12]). We use a simple fixed range propagation model where the transmission range of each AP is a disc of radius R. Users are assumed to associate with the nearest AP. Thus the set of points served by an AP, i.e., the cell, is a subset of its transmission range. We denote by A_i the area of cell i. We also use a simple capture model where a packet is correctly received if and only if there is no other transmitting station with its transmission range. Thus

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

where v_1 and v_2 denote the position of the APs of u_1 and u_2, respectively. Since traffic is uniformly distributed in $\mathbb{R}^2$, the interference factor I_{ik} between cells i and k is

$$I_{ik} = \frac{1}{A_i A_k} \int_{\mathcal{U}_i \times \mathcal{U}_k} \chi(u_i, u_k) du_i du_k. \tag{10}$$

It can be evaluated numerically using the following procedure. For each point u_i of cell i, draw two disks of radius R, the first centred at AP i, the second at u_i. The union of these two disks constitutes the exclusion region for any transmission towards u_i. Then compute the area of cell k that intersects with this exclusion region with the convention that the intersection is equal to all the area of cell k if AP k is located inside the exclusion region. Add all of these intersection areas together and then divide by the product $A_i A_k$.

The channel assignment algorithm is centralised. The network operator provides the algorithm with the positions of the APs and a single additional parameter: the transmission range R. The algorithm computes the coefficients I_{ik} according to

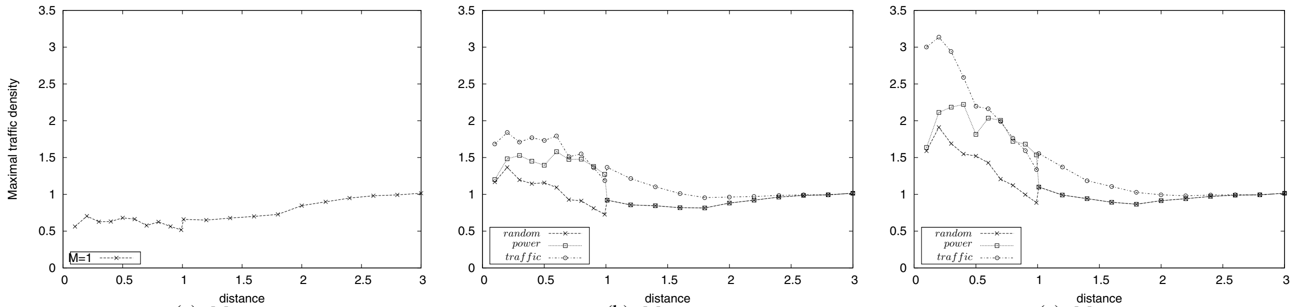

(a) $M = 1$ (b) $M = 3$ (c) $M = 6$

Fig. 6. Maximal traffic density: regular networks.

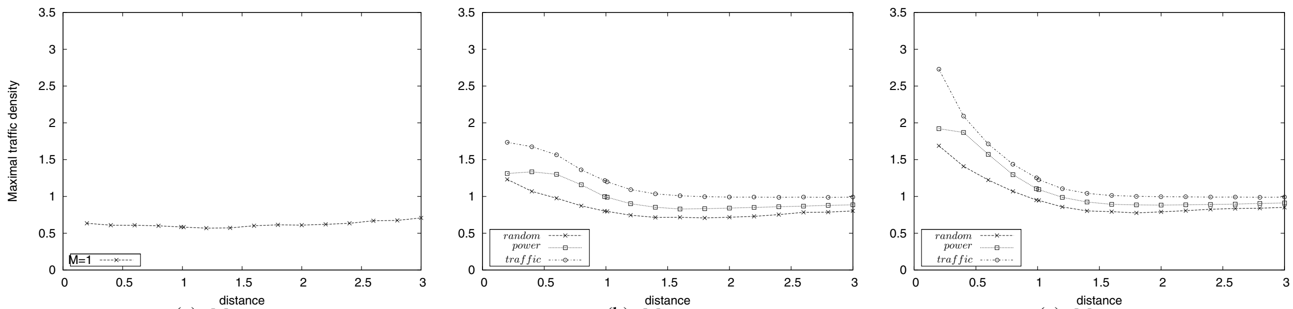

(a) $M = 1$ (b) $M = 3$ (c) $M = 6$

Fig. 7. Maximal traffic density: semi-regular networks.

Equation (10) and then solves the optimisation problem (2) using for example a local search approach. The objective function of the optimisation problem is evaluated using Algorithm 1.

We compare our results with the widely used received signal strength approach. There are many variants of this model. They differ mainly by their objective function and their implementation, e.g., distributed or centralised, adaptive or static. However, they are all similar in the sense that channel assignment is based only on a common underlying graph that we call the hearing graph. The vertices of this graph are the APs and an edge between two vertices means that the APs are able to hear each other. The interference metric can be evaluated locally and in a distributed manner by each AP. In this paper we use a simple variant of the approach that is consistent with the capture and propagation models mentioned above. Two APs are able to hear each other if they are separated by a distance less than R. Thus the adjacency matrix of the hearing graph can be written

$$H_{ik} = \mathbb{I}_{\{d(i,k)<R\}}.$$

The channel allocation mechanism consists in minimizing the number of adjacent APs that use the same frequency channel.

$$\begin{cases} \text{maximise} & \kappa(\theta) = \frac{1}{\sum_{i,k} H_{ik}\mathbb{I}_{\{\theta_i=\theta_k\}}} \\ \text{subject to} & \theta \in \Theta. \end{cases} \tag{11}$$

VI. Numerical results

To illustrate potential capacity gains, we present numerical results for the network model of Section V. We consider a traffic density of $1/\pi$ $(bits/s/m^2)$. By definition, the traffic density is the traffic intensity per unit of area. The factor $1/\pi$ $(bits/s/m^2)$ is a normalising constant. We have chosen this factor because it is the maximal traffic density that can be achieved in an isolated cell. We present our results in terms of γ^* where γ^*/π $(bits/s/m^2)$ is the maximal traffic density that can be achieved by the network. We compare the performance of three channel allocation mechanisms. The first is the random allocation scheme where an AP selects a random frequency channel. The second referred to as *power* corresponds to the signal-based approach described above. The third, referred to as *traffic*, corresponds to the system (2).

We use a local search based algorithm to solve the $NP-$hard optimisation problems (2) and (11). The local search approach consists in defining a neighbouring structure over the search space Θ. The algorithm starts with a random channel assignment. Then at each iteration it moves to a neighbouring assignment that maximises the objective function. The algorithm terminates when no neighbouring assignment can increase the objective function. The neighbouring structure we have chosen is the following: two channel assignments θ_1 and θ_2 are considered as neighbours if and only if θ_1 differs from θ_2 in at most one AP, i.e., θ_1 and θ_2 assign the same channels to all APs except at most one AP.

Regular networks:: In Figure 6 we consider a regular grid of $7 \times 7 = 49$ APs. We plot the maximal traffic density γ^* as the distance d between two APs varies from 0 to $3R$. Figure 6(b) (resp. 6(c)) corresponds to $M = 3$ (resp. $M = 6$). The random and signal-based allocations achieve the same traffic

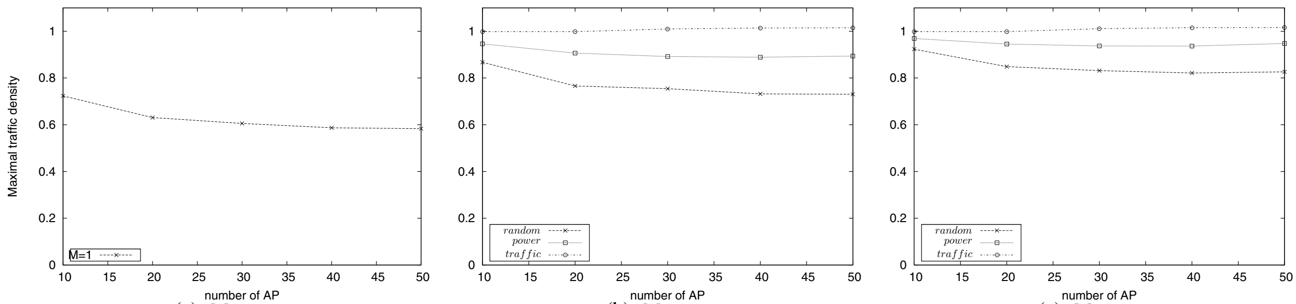

(a) $M = 1$ (b) $M = 3$ (c) $M = 6$

Fig. 8. Maximal traffic density: random networks.

density as long as $d > R$ while the traffic density allocation provides a gain that attains 40%. For d slightly less than R, the performance of *power* and *traffic* is the same since most interference comes from first tier cells and is well captured by the signal-based model. Note that traffic densities greater then 1 are achieved due to reduced cell sizes and the use of multiple frequency channels.

Semi-regular networks: In Figure 7 we plot the maximal traffic densities of a 7×7 semi-regular network versus the side length of the cells of the underlying grid. In a semi-regular network the APs are placed randomly inside the cells of a regular grid. For each plot the results are averaged over 20 different network topologies and then each network topology is averaged over 20 different allocations. We observe that the maximal traffic density increases smoothly as the cell size decreases. The traffic density allocation yields up to 20% more capacity than the signal-based allocation and up to 50% more than the random allocation.

Random networks:: Maximal traffic densities as a function of the number of APs are plotted in Figures 8(b) and 8(c) for $M = 3$ and 6. The APs are distributed uniformly over a square area with side length equal to $10R$. The gains are 10% and 20% with respect to the signal based and random allocations.

VII. Conclusion

In this paper we have proposed a channel assignment method for WLANs. Our method incorporates traffic demand and uses an appropriate interference model that captures the correct type of interference in multi-cell WLANs. It relies on a novel heuristic for estimating the traffic capacity of a multi-cell network by means of a fluid approximation. Numerical results suggest there is scope for a significant gain in capacity over classical signal based assignments.

In future work we intend to extend the presented approach to account for dynamic traffic conditions. We also aim to derive a distributed approach appropriate for WLANs without a common administrative authority able to apply our centralized algorithm.

References

[1] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot, "Measurement-based self organization of interfering 802.11 wireless access networks," in *INFOCOM*, 2007.

[2] K. K. Leung and B. J. Kim, "Frequency assignment for IEEE 802.11 wireless networks," in *VTC*, 2003.

[3] A. Mishra, S. Banerjee, and W. Arbaugh, "Weighted coloring based channel assignment for wlans," *SIGMOBILE Mob. Comput. Commun. Rev.*, 2005.

[4] A. Mishra, V. Shrivastava, D. Agrawal, S. Banerjee, and S. Ganguly, "Distributed channel management in uncoordinated wireless environments," in *MobiCom*, 2006.

[5] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. Arbaugh, "A client-driven approach for channel management in wireless LANs," in *INFOCOM*, 2006.

[6] Y. Lee, K. Kim, and Y. Choi, "Optimization of ap placement and channel assignment in wireless lans," in *LCN*, 2002.

[7] H. Luo and N. K. Shankaranarayanan, "A distributed dynamic channel allocation technique for throughput improvement in a dense WLAN environment," in *ICASSP*, 2004.

[8] D. J. Leith and P. Clifford, "A self-managed distributed channel selection algorithm for WLANs," in *4th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, 2006.

[9] ——, "Channel dependent interference and decentralized colouring," in *NETCOOP*, 2007.

[10] G. Bianchi, "Performance analysis of the ieee 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, 2000.

[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.

[12] T. Bonald, A. Ibrahim, and J. Roberts, "Traffic capacity of multi-cell wlans," in *SIGMETRICS*, 2008.

[13] D. Kunz, "Channel assignment for cellular radio using neural networks," *IEEE Transactions on Vehicular Technology*, 1991.

[14] M. Duque-Anton, D. Kunz, and B. Ruber, "Channel assignment for cellular radio using simulated annealing," *IEEE Transactions on Vehicular Technology*, 1993.

[15] W. K. Lai and G. G. Coghill, "Channel assignment through evolutionary optimization," *IEEE Transactions on Vehicular Technology*, 1996.

[16] K. A. Smith, "A genetic algorithm for the channel assignment problem," *GLOBECOM*, 1998.

[17] J. Dai, "On positive harris recurrence of multiclass queueing networks: a unified approach via fluid limit models," *The Annals of Applied Probability*, 1995.