



Enhanced Spatial Reuse in Multi-Cell WLANs

Thomas Bonald, Ali Ibrahim, James Roberts

► To cite this version:

Thomas Bonald, Ali Ibrahim, James Roberts. Enhanced Spatial Reuse in Multi-Cell WLANs. Info-com, 2009, Rio de Janeiro, Brazil. 10.1109/INFCOM.2009.5062200 . hal-01244784

HAL Id: hal-01244784

<https://hal.science/hal-01244784>

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.

Enhanced Spatial Reuse in 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—When IEEE 802.11 access points (APs) share the same channel in a multi-cell WLAN, their downlink transmissions can interfere. Typically, an AP whose scheduled transmission to some user is blocked by another cell will apply the CSMA/CA back-off algorithm and continue to make reattempts to the same user. Through analytical models and simulations, we demonstrate that significant capacity gains can be attained by choosing an alternative destination for the reattempt. Results demonstrate that a simple random choice of alternative destination brings almost the same gain as a more sophisticated algorithm that seeks to maximize spatial reuse.

I. INTRODUCTION

The downlink traffic capacity of multiple interfering IEEE 802.11 Access Points (APs) can be improved by introducing opportunism in the way APs schedule their transmissions. The envisaged multi-cell WLAN consists of a number of APs transmitting on the same frequency channel to a population of users distributed over the coverage area. We assume a stochastic traffic model where finite size flows arrive according to a Poisson process. Traffic capacity is measured by the maximum load at which the network remains stable in the sense that the number of flows in progress does not explode.

Traffic capacity depends on the way the APs interfere. Under the IEEE 802.11 Distributed Control Function (DCF), channel access is governed by the CSMA/CA algorithm. Three features of this algorithm are relevant to our study: (i) the backoff mechanism implemented by each station effectively provides random access opportunities to the APs downloading flows to their users; (ii) the two small RTS/CTS frames exchanged between transmitter and receiver in an access attempt inhibit interfering stations from attempting a simultaneous transmission; (iii) when a packet retransmission is blocked, the AP will make repeated attempts to send the same packet on successive expirations of the backoff counter. Our proposal is to modify the last feature by allowing an AP to transmit to an alternative destination after observing an initial collision.

The problem is more easily explained with reference to a simple example. In Figure 1 we depict a network consisting of two APs and three user stations. User station u_1 is associated with AP_1 and u_2 and u_3 with AP_2 . Transmissions to u_1 and u_2 by AP_1 and AP_2 , respectively, cannot occur at the same time due to the channel reservation mechanism (one user station is inhibited on hearing the CTS sent by the other) while transmissions to u_1 and u_3 can. Now assume both APs are in backoff mode and contending for channel access. AP_1 has a packet destined to u_1 and AP_2 has a packet destined to u_2 . Depending on the backoff timers, one of the APs will be

the first to try to reserve the channel. Assume AP_1 is first. An RTS/CTS sequence is exchanged between AP_1 and u_1 preventing any transmission involving u_2 . When the backoff timer of AP_2 expires, an RTS is sent to u_2 , but u_2 is blocked by the AP_1 transmission and cannot send a CTS response. AP_2 concludes that a collision has occurred at u_2 and backs off for a random time before attempting a new transmission to the same destination. As long as AP_1 is transmitting to u_1 none of AP_2 's retransmissions will succeed. On the other hand, if AP_2 changes its destination and tries to transmit a packet to u_3 it will succeed.




Fig. 1. Two AP network.

We propose two opportunistic scheduling algorithms that would allow this choice of an alternative destination. The first simply changes for a randomly chosen alternative destination when the first is blocked. The second is more sophisticated and seeks to optimize the choice of destination to minimize future blocking in other AP cells. We show that a significant increase in traffic capacity can be realized depending on the network topology.

After an outline of related work we introduce the considered network and traffic models. We then present the analytical model, deriving capacity limits for a simple configuration by the fluid limit approach. Simulations are used to validate the results and extend application to a more complex network.

II. RELATED WORK

The considered problem, as described in the introduction, is related to the so-called Head of Line (HoL) blocking problem of FIFO base station schedulers studied in [1], [2]. A packet can be blocked when the destination user is experiencing poor radio conditions leading to repeated errors. Throughput is improved by choosing to transmit to an alternative user whose radio conditions are more favorable. Our proposed non-FIFO algorithms for choosing the users to which an AP should transmit can also be considered as opportunistic scheduling, as considered in [3], [4], [5], [6], for example. Like [6], we evaluate their performance in terms of the limiting demand at which the network remains stable under a stochastic traffic model. We are not aware of any other work that

has considered HoL blocking or opportunistic scheduling for multi-cell WLANs.

The above cited work concerns isolated base stations or APs. A number of other papers have considered how throughput can be increased by coordinating transmission in a cellular network [7]–[9]. It is necessary to realize the optimal trade off between a reduction in the number of simultaneous transmissions and the gain in rate realized by reducing interference. It is demonstrated in [8], [9] that considerable capacity gains are possible compared to a baseline scenario where all stations transmit simultaneously. Our work considers not a cellular network but a set of interfering WLAN access points and proposes distributed scheduling algorithms.

III. NETWORK CAPACITY UNDER RANDOM TRAFFIC

We present a simplified slotted time version of the multi-cell WLAN network and describe the considered traffic model.

A. A simplified network model

Consider a set of Access Points \mathcal{A} operating on the same channel and a set of user classes \mathcal{U} . With each AP i is associated a subset of user classes \mathcal{U}_i . Each class models a set of users having the same interference and traffic characteristics. These users would typically be located within a small geographical area. At any given time, a set of users are active in that they receive traffic from their AP: the AP sends them packets at a rate determined by the WLAN collision resolution algorithm. APs with at least one active user are said to be backlogged.

Collision resolution in the real multi-cell system is extremely complex to model. To progress in analysing the problem at hand we therefore make some simplifying assumptions. We model interference between users by a function χ from $\mathcal{U} \times \mathcal{U}$ to $\{0, 1\}$. For $j, k \in \mathcal{U}$, $\chi(j, k) = 1$ iff transmissions to class- j and class- k users cannot be scheduled simultaneously. Note that $\chi(j, k) = 1$ for all $j, k \in \mathcal{U}_i$.

Time is slotted with each slot representing the transmission time of a single packet including all overheads: RTS/CTS exchange, ACK transmission, interference spaces SIFS and DIFS and the backoff overhead. All APs are synchronized and collisions are assumed to be resolved at the start of each time slot.

B. Traffic model

We consider a flow-level traffic model. Users become active at random instants of time, require a finite random-sized file to be downloaded and cease to be active when the transfer is completed. An active user has a flow in progress. We assume class- j users have a Poisson flow arrival rate λ_j (flow/s) and an exponential flow size distribution of mean $1/\mu$ bits. This assumption is not crucial for studying network stability. Class- j traffic intensity is denoted $\rho_j = \lambda_j/\mu$.

C. Channel reservation

Resource allocation occurs in two phases, a channel reservation phase and a transmission phase. During the reservation

phase, backoff and packet scheduling mechanisms combine to define the set of users to which the APs will transmit during the ensuing transmission phase.

We assume channel reservation operates as follows. At the beginning of a time slot, all backlogged APs contend for channel access. One randomly chosen AP is assumed to win a first backoff contention. The scheduling mechanism in the AP selects the active user to which it will transmit.

All users in neighbouring cells that interfere with the scheduled transmission (as determined by an RTS/CTS exchange) are inhibited from channel access during the present slot. The same process is repeated to choose another AP and its scheduled transmission, among those that are not inhibited by previously scheduled transmissions, and so on until all APs are either scheduled or inhibited. All scheduled packets are transmitted in the ensuing transmission phase after which a new slot begins. By assumption, the channel reservation phase consumes a negligible amount of time, even when it includes multiple backoff phases.

D. Capacity region

Let x_j be the number of active class- j users. We refer to the vector $x = (x_j)_{j \in \mathcal{U}}$ as the network state. We suppose the scheduling algorithm can be characterized by state-dependent service rates ϕ_j . These rates, defined at the time scale of the flow-level process, are assimilated to the proportion of slots that would be scheduled for the user class assuming the mix of active users remains fixed. The evolution of the network state $x(t)$ is then Markovian with arrival rates λ_j and service rates $\phi_j\mu$. We define the capacity region of a given scheduling algorithm as the set of all traffic intensity vectors $(\rho_j)_{j \in \mathcal{U}}$ such that the Markov process $x(t)$ is ergodic.

IV. SCHEDULING ALGORITHMS

We consider three scheduling algorithms providing progressively larger capacity regions.

A. RETransmit: RET

This scheduler corresponds to the way IEEE 802.11 currently works. After each successful transmission, the AP selects a new user at random (in fact, the user of the packet that is currently head of line) and attempts to transmit a packet to it in the next slot. If this user turns out to be inhibited by previously scheduled transmissions, the slot will remain unused by the considered AP. This choice is maintained for successive slots until transmission is successful.

B. Change Destination: CD

CD ensures that the set of scheduled transmissions in every slot is maximal, i.e., any other unscheduled transmission would interfere with at least one member of the set. Instead of waiting for a selected user to become free as in RET, the AP simply selects another free destination.

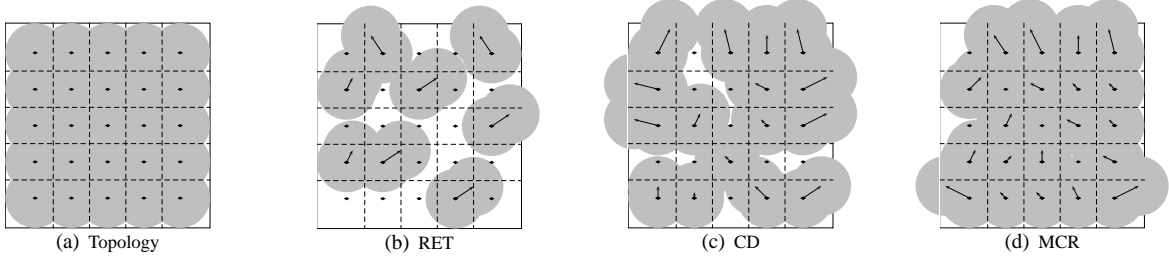


Fig. 2. Spatial reuse capacity of RET, CD and MCR.

C. Minimal Cost Reduction: MCR

MCR seeks to choose the least interfering destinations. Suppose at a certain stage, we have already scheduled some APs and know the set of free users \mathcal{F} . Define, for $j \in \mathcal{U}$,

$$\delta_j = \sum_{i \in \mathcal{A}} \sum_{k \in \mathcal{U}_i \cap \mathcal{F}} \frac{x_k \chi(j, k)}{\sum_{l \in \mathcal{U}_i \cap \mathcal{F}} x_l} / \sum_{i \in \mathcal{A}} \sum_{k \in \mathcal{U}_i} \frac{x_k \chi(j, k)}{\sum_{l \in \mathcal{U}_i} x_l}. \quad (1)$$

The numerator of δ_j measures the fraction of free users that would become blocked by a transmission to class j from its associated AP. The denominator is the same fraction assuming all users are free so that δ_j is a measure of the relative cost of choosing user class j . MCR aims at each stage to minimize this cost. The intention in using the relative variation of transmission cost is to schedule users when they consume a minimal amount of resources compared to their nominal resource consumption.

D. Contention regions

Figure 2 illustrates the spatial reuse capacity of each of the above algorithms. Figures 2(b), 2(c) and 2(d) show a snap-shot of a realization of the respective algorithms. Each scheduled transmission inhibits connections in a certain space called the contention region. This is the union of two disks, one centred at the AP, the other at the receiver. No transmission can take place between an AP and a receiver if either lies in the contention region of a previously scheduled transmission. In this example RET has the lowest spatial reuse capacity with only 8 APs transmitting while CD and MCR schedule 16 and 20, respectively.

V. ANALYTICAL DETERMINATION OF CAPACITY

In this section we derive the capacity region of the network shown in Fig. 1 using the fluid limit approach [10], [11]. Note that CD and MCR are equivalent for this network.

We first recall the following results from [12] for the RET algorithm. It is shown there that service rates can be approximated by:

$$\phi_j = \frac{\xi_j}{\sum_{n \in \mathcal{A}} \sum_{k \in \mathcal{U}_i} \sum_{l \in \mathcal{U}_n} \xi_k \xi_l \chi(k, l)}, \quad j \in \mathcal{U}_i, \quad (2)$$

where the ξ_j are given by

$$\xi_j = \frac{x_j}{\sum_{k \in \mathcal{U}_i} x_k}, \quad j \in \mathcal{U}_i. \quad (3)$$

The denominator of (2) can be interpreted as the mean number of slots required by AP i to transmit a packet and the

numerator ξ_j indicates that the capacity of AP i is shared equally between all users of this cell.

RET scheduling: First note that if $x_1 > 0$ and $x_2 + x_3 > 0$ service rates are as follows:

$$\phi_1 = \frac{1}{1 + \xi_2}, \quad \phi_2 = \frac{\xi_2}{1 + \xi_2}, \quad \phi_3 = \frac{\xi_3}{1 + \xi_2}.$$

If one cell is empty the other behaves like a Processor Sharing (PS) server.

Proposition 1: The capacity region of RET is defined as follows:

$$\begin{cases} \rho_2 + \rho_3 + \frac{\rho_1 \rho_2}{\rho_2 + \rho_3} < 1 & \text{if } \rho_1 < \rho_2 + \rho_3, \\ \rho_1 + \rho_2 < 1 & \text{otherwise.} \end{cases}$$

Proof of Proposition 1: The fluid volume corresponding to class- j users is denoted by \bar{x}_j . The corresponding proportion of fluid at AP i is denoted by $\bar{\xi}_j$, with $j \in \mathcal{U}_i$. We first consider the following two cases:

Case $\bar{x}_1 = 0, \bar{x}_2 > 0, \bar{x}_3 > 0$: Class-1 users observe an $M/M/1$ PS queue with capacity $1/(1 + \bar{\xi}_2)$. When $x_1 = 0$, the service rate of class-2 is $\bar{\xi}_2$ while when $x_1 > 0$ it is $\bar{\xi}_2/(1 + \bar{\xi}_2)$. Since the probability that $x_1 = 0$ is $1 - \rho_1(1 + \bar{\xi}_2)$, we deduce the average fluid service rate:

$$\bar{\phi}_2 = \bar{\xi}_2(1 - \rho_1 \bar{\xi}_2). \quad (4)$$

Similarly, we get:

$$\bar{\phi}_3 = \bar{\xi}_3(1 - \rho_1 \bar{\xi}_2). \quad (5)$$

Case $\bar{x}_1 > 0, \bar{x}_2 = 0, \bar{x}_3 = 0$: Class 2 and 3 users see a “non-conservative” Discriminatory Processor Sharing (DPS) server with rates $\phi_2 = x_2/(2x_2 + x_3)$ and $\phi_3 = x_3/(2x_2 + x_3)$. The probability that the DPS queue is empty is $1 - 2\rho_2 - \rho_3$. Class-1 users are served at a rate that depends on the state of cell 2. Its mean value is given by:

$$\pi(0, 0) \times 1 + \sum_{x_2 + x_3 > 0} \pi(x_2, x_3) \frac{1}{1 + \xi_2},$$

where π is the stationary distribution of the DPS queue. The sum $\sum_{x_2 + x_3 > 0} \pi(x_2, x_3)/(1 + \xi_2)$ is also the mean service rate of cell 2. It must therefore be equal to the cell 2 arrival rate $\rho_2 + \rho_3$. We deduce the fluid limit rate:

$$\bar{\phi}_1 = 1 - \rho_2. \quad (6)$$

Stability conditions: Observe that the stochastic model satisfies $\phi_2/x_2 = \phi_3/x_3$ for all $x_2, x_3 > 0$. This equality also holds for the fluid limit, $\bar{\phi}_2/\bar{x}_2 = \bar{\phi}_3/\bar{x}_3$. Thus as long as \bar{x}_2 and \bar{x}_3 are positive we have:

$$\frac{d}{dt} \log \left(\frac{\bar{x}_3}{\bar{x}_2} \right) = \frac{\rho_3}{\bar{x}_3} - \frac{\rho_2}{\bar{x}_2}.$$

This shows that either $\bar{x}_2 + \bar{x}_3 \rightarrow 0$ or $\bar{\xi}_2 \rightarrow \rho_2/(\rho_2 + \rho_3)$ and $\bar{\xi}_3 \rightarrow \rho_3/(\rho_2 + \rho_3)$.

While all fluid components are positive, we have

$$\bar{\phi}_1 = \frac{\rho_2 + \rho_3}{2\rho_2 + \rho_3}, \quad \bar{\phi}_2 = \frac{\rho_2}{2\rho_2 + \rho_3} \text{ and } \bar{\phi}_3 = \frac{\rho_3}{2\rho_2 + \rho_3}.$$

If $\rho_1 > (\rho_2 + \rho_3)/(2\rho_2 + \rho_3)$ and $2\rho_2 + \rho_3 > 1$ the fluid limit dynamical system is unstable. If $\rho_1 < (\rho_2 + \rho_3)/(2\rho_2 + \rho_3)$ and $2\rho_2 + \rho_3 < 1$, the system is stable. If $\rho_1 > (\rho_2 + \rho_3)/(2\rho_2 + \rho_3)$ and $2\rho_2 + \rho_3 < 1$, then $(\bar{x}_2, \bar{x}_3) \rightarrow (0, 0)$ and $\bar{x}_1 > 0$. In view of (6) the fluid limit is stable iff $\rho_1 < 1 - \rho_2$. If $\rho_1 < (\rho_2 + \rho_3)/(2\rho_2 + \rho_3)$ and $2\rho_2 + \rho_3 > 1$, then $\bar{x}_1 \rightarrow 0$ and $\bar{x}_2 > 0$ and $\bar{x}_3 > 0$. By (4) and (5) the fluid model is therefore stable iff $\rho_2 + \rho_3 < 1 - \rho_1\rho_2/(\rho_2 + \rho_3)$.

It can be verified that the above set of stability conditions can be summarized by the inequalities in the proposition. ■

CD scheduling: First assume $x_1, x_2, x_3 > 0$. AP_1 transmits to u_1 if it wins the channel contention (with probability 1/2) or AP_2 wins and AP_2 transmits to u_3 (probability $\xi_3/2$). The service rate of u_1 is thus:

$$\phi_1 = \frac{1}{2}(1 + \xi_3). \quad (7)$$

AP_2 will transmit to u_2 if and only if it wins (probability 1/2) and AP_2 chooses u_2 as a destination (probability ξ_2). The service rate is:

$$\phi_2 = \frac{\xi_2}{2}. \quad (8)$$

Finally, AP_2 transmits to u_3 if AP_2 wins (probability 1/2) and chooses u_3 (probability ξ_3), or if AP_1 wins (probability 1/2). The service rate is:

$$\phi_3 = \frac{1}{2}(1 + \xi_3). \quad (9)$$

Proposition 2: The capacity region of CD is given by

$$\rho_2 + \max(\rho_3, \rho_1) < 1.$$

Proof of Proposition 2:

From equations (7), (8) and (9) we deduce the following:

$$\begin{cases} \phi_2 + \phi_3 = 1 & \forall x_1 \geq 0, x_2 \geq 0, x_3 > 0 \\ \phi_1 + \phi_2 = 1 & \forall x_1 > 0, x_2 + x_3 \geq 0. \end{cases}$$

These relations hold also for the fluid limits

$$\begin{cases} \bar{\phi}_2 + \bar{\phi}_3 = 1 & \forall \bar{x}_1 \geq 0, \bar{x}_2 \geq 0, \bar{x}_3 > 0. \\ \bar{\phi}_1 + \bar{\phi}_2 = 1 & \forall \bar{x}_1 > 0, \bar{x}_2 + \bar{x}_3 \geq 0. \end{cases}$$

Thus if $\rho_2 + \rho_3 < 1$ and $\rho_1 + \rho_2 < 1$ neither \bar{x}_1 nor \bar{x}_3 can be positive after a finite time. To show the stability of the fluid limit it remains to compute $\bar{\phi}_2$ when $\bar{x}_1 = 0, \bar{x}_2 > 0$ and

$\bar{x}_3 = 0$. This means that we must consider the following four cases

$$(\phi_1, \phi_2, \phi_3) = \begin{cases} (1/2, 1/2, 1/2), & x_1 > 0, x_3 > 0, \\ (1/2, 1/2, 0), & x_1 > 0, x_3 = 0, \\ (0, 1, 0), & x_1 = 0, x_3 > 0, \\ (0, 1, 0), & x_1 = 0, x_3 = 0. \end{cases} \quad (10)$$

First note that if $\rho_1 > 1/2$, \bar{x}_1 cannot be zero before \bar{x}_2 because $\bar{\phi}_1 = 1/2$. This means \bar{x}_2 attains zero before \bar{x}_1 and the fluid limit is stable. Now, if $\rho_1 < 1/2$, class-1 users observe an $M/M/1$ server of rate 1/2. Thus x_1 is zero with probability $1 - 2\rho_1$. The rate allocated to class-2 users depends only on whether x_1 is positive or zero and is given by (10). We deduce $\bar{\phi}_2 = 2\rho_1(1/2) + (1 - 2\rho_1)1 = 1 - \rho_1$.

Now, since $\rho_2 + \rho_1 < 1$, we have $\rho_2 < \bar{\phi}_2$. On the other hand, it is obvious that $\rho_2 + \rho_3 < 1$ and $\rho_1 + \rho_2 < 1$ are necessary conditions for stability. We conclude that the capacity region is indeed as stated in the proposition. ■

Figure 4(a) plots the RET and CD capacity regions when $\rho_1 = 0.5$. The region for RET is strictly included in that for CD showing that RET is not Pareto efficient since both ρ_2 and ρ_3 can be increased. The figure also shows that CD realizes the maximal capacity region achievable by any scheduling algorithm since the stated stability conditions are necessary for the stability of any scheduling algorithm.

VI. SIMULATION RESULTS

We first report results of ns-2 simulations of the simple network of Fig. 1. These simulations take proper account of the CSMA/CA protocol and make no assumption of slotted operation. The normalized results depicted in Fig. 4(b) show that the gain from opportunistic scheduling is more significant than that predicted by the analytical model. The main difference is that the capacity of RET suffers from a phenomenon of short term unfairness: the CSMA/CA algorithms typically allow one AP to capture the channel for repeated transmissions (its backoff timer is smaller) thus blocking the other for long periods and amplifying the resulting loss in capacity. CD effectively counters this phenomenon.



Fig. 3. Infinite linear network in 1D.

To evaluate the interest of MCR, we have simulated a more complex network but preserve the simplified slotted contention model. The network consists of a linear array of APs, as depicted in Fig. 3. The distance between two APs is d and user traffic is uniformly distributed on the line with traffic density ϱ (in bits/s/m). Transmission range is 1.

For this network it is possible to characterize the maximal capacity C_{\max} (expressed in bits/s) defined as $\varrho_{\max} \times$ cell length where ϱ_{\max} is the maximal traffic density that can be achieved by any scheduling algorithm:

$$C_{\max} = \begin{cases} \frac{d}{1 + \frac{d}{n}}, & \frac{1}{n} < d \leq \frac{1}{n-1}, \\ \frac{1}{n}, & \frac{1}{n-1} \leq d < \frac{1}{n-2}, \end{cases}$$




Fig. 4. Capacity region of the network of Fig. 1.

where n is a positive integer (proof is omitted).

We rely on packet level simulation to determine the capacity of RET, CD and MCR. We use a simple propagation model. Stations (APs and users) can communicate if and only if they are within transmission range. The capture model is such that a packet will be received successfully if and only if there is no other transmitting station located within the transmission range of the receiver. Time is slotted and scheduling proceeds in two phases as described previously. Users associate with the nearest APs. Flow sizes are distributed geometrically with a mean size of 100 packets. To evaluate capacity, i.e., the maximal traffic load supported by the network, we measure the flow departure rate under heavy traffic.




Fig. 5. Linear network in 1D.

In Figure 5 we plot simulation results for RET, CD and MCR together with the maximal capacity. The simulations have been performed on a network of 15 cells and the figure shows the capacity of the central cell. We observe that when $d > 2$, both CD and MCR achieve optimal capacity. Both CD and MCR are extremely close suggesting simple randomized opportunism is sufficient. Capacity gain with respect to RET attains 40% at $d = 1.6$.

VII. CONCLUSION

We have demonstrated the potential loss of traffic capacity in multi-cell WLANs caused by a form of head of line blocking. This occurs when a scheduled transmission in one AP cell interferes with a previously scheduled transmission in another AP. In IEEE 802.11, the blocked transmission will be reattempted several times despite the likelihood that interference is persistent and the new attempts will not succeed.

We have demonstrated that significant capacity gains are possible by opportunistic scheduling. Two scheduling algorithms of increasing complexity have been proposed. For the representative toy networks considered in our evaluation, the best part of the gains are achieved by the simplest algorithm called CD. This scheduler just picks a feasible alternative destination when the first choice is blocked. It appears to constitute a reasonable compromise between traffic efficiency and implementation complexity.

Simulations show that capacity gains in a real network is likely to be greater than in the simplified model considered for analysis. APs can capture the channel for repeated packet transmissions thus extending the period during which contending APs suffer head of line blocking.

REFERENCES

- [1] P. Bhagwat, P. Bhattacharya, A. Krishna, and S. K. Tripathi, "Enhancing throughput over wireless LANs using channel state dependent packet scheduling," in *INFOCOM '96*.
- [2] L. B. Jiang and S. C. Liew, "An adaptive round robin scheduler for head-of-line-blocking problem in wireless LANs," in *WCNW 2005*.
- [3] J. M. Holtzman, "Asymptotic analysis of proportional fair algorithm," in *PIMRC*, 2001.
- [4] X. Liu, E. K. P. Chong, and N. B. Shroff, "Opportunistic transmission scheduling with resource-sharing constraints in wireless networks," *IEEE Journal on Selected Areas in Communications*, 2001.
- [5] H. J. Kushner and P. A. Whiting, "Convergence of proportional-fair sharing algorithms under general conditions," *IEEE Transactions on Wireless Communications*, 2004.
- [6] S. Borst, "User-level performance of channel-aware scheduling algorithms in wireless data networks," *IEEE/ACM Transactions on Networking*, 2005.
- [7] A. Bedekar, S. Borst, K. Ramanan, P. Whiting, and E. Yeh, "Downlink scheduling in CDMA data networks," in *GLOBECOM*, 1999.
- [8] S. Das, H. Viswanathan, and G. Rittenhouse, "Dynamic load balancing through coordinated scheduling in packet data systems," in *INFOCOM 2003*.
- [9] T. Bonald, S. Borst, and A. Proutiere, "Inter-cell coordination in wireless data networks," in *European Transactions on Telecommunications*, 2006.
- [10] J. Dai, "On positive harris recurrence of multiclass queueing networks: a unified approach via fluid limit models," *The Annals of Applied Probability*, vol. 5, no. 1, pp. 49–77, 1995.
- [11] S. P. Meyn, "Transience of multiclass queueing networks and their fluid limit models," *The Annals of Applied Probability*, vol. 5, pp. 946–957, 1995.
- [12] T. Bonald, A. Ibrahim, and J. Roberts, "Traffic capacity of multi-cell w lans," in *SIGMETRICS*, 2008.