

# Load-Aware Spectrum Distribution in Wireless LANs

Thomas Moscibroda, Ranveer Chandra, Yunnan Wu, Sudipta Sengupta, Paramvir Bahl, Yuan Yuan  
Microsoft Research, One Microsoft Way, Redmond, WA 98052  
{moscitho, ranveer, yunnanwu, sudipta, bahl}@microsoft.com

**Abstract**—Traditionally, the channelization structure in IEEE 802.11-based Wireless LANs has been fixed: Each access point (AP) is assigned one channel and all channels are equally wide. In contrast, it has recently been shown that even on commodity hardware, the channel-width can be adapted dynamically purely in software. Leveraging this capability, we study the use of *dynamic-width channels*, where every AP adaptively adjusts not only its center-frequency, but also its *channel-width* to match its traffic load. This gives rise to a novel optimization problem that differs from previously studied channel assignment problems. We propose efficient spectrum-distribution algorithms and evaluate their effectiveness through analysis and simulations using real-world traces. Our results indicate that by allocating more spectrum to highly-loaded APs, the overall spectrum-utilization can be substantially improved and the notorious load-balancing problem in WLANs can be solved naturally.

## I. INTRODUCTION

One of the core design principles of IEEE 802.11-based networks is the use of a simple, fixed channelization structure. The entire available spectrum is divided into smaller channels of equal channel-width (bandwidth), and each IEEE 802.11 network is specified to operate on a specific set of channels. For example, the 2.4GHz ISM band has 3 non-overlapping channels each being 20MHz wide, and at any given time each Access Point (AP) operates on a particular channel.

In this paper, we argue that by moving beyond this fixed channelization structure, the network capacity, overall spectrum utilization and fairness of WLANs can be greatly increased. In WLANs, clients are often unevenly distributed across the network and different clients have different traffic requirements. Therefore, certain APs can become hotspots and have to handle high traffic load, while others remain under-utilized. With a-priori channels of fixed width, it is difficult to naturally adapt to such spatial and temporal disparity of traffic distribution [7], [12], [13]; the overall spectrum utilization in the network, and hence its capacity, is reduced. Also, the fact that some APs are heavily loaded while others are not, creates a location-induced fairness problem. For instance, a highly-loaded AP near a conference room may have to serve multiple clients on a single channel, thus hurting these clients' performance, while at the same time, other APs in the network may serve only very few clients or no clients at all. While there exist various approaches to mitigating these problems (including power allocation [6], client-AP assignment [8], [21], and channel-assignment [20], [18], or combinations thereof), none of them addresses the fundamental root cause of the problem: heavily-loaded network areas and the APs therein require *more spectrum* than others in order to serve their load equally well.

In this paper, we take a fresh look at channelization in WLANs. In particular, we study the problem of how spectrum should be allocated to APs if the *width of the communication channels* can be adaptively changed: wider channels (e.g. up to 40MHz) for heavily-loaded APs and narrower channels (say, down to 5MHz) for lightly-loaded APs. This way, adapting channel-widths of APs and their respective clients allows to naturally implement the maxim “*provide more spectrum where spectrum is needed*”, thus enabling a conceptually very simple and efficient solution to the load-balancing problem.

Our studying the channel width as an adaptable knob, rather than a fixed pre-set parameter is motivated by recent work that has shown how even on commodity hardware such as the Atheros chipset [1], the channel width can be changed dynamically and purely in software with very little overhead [10]. Furthermore, there also exist recent advances in hardware technology that enable wireless devices to dynamically change both their operating frequency and channel-width. For example, WiMAX allows nodes to use 11 possible channel widths [3]. The 2007 version of the IEEE 802.11 standard [4] proposes the use of 5, 10 and 20MHz channel-widths for operation in different parts of the spectrum.

In view of these new possibilities, the key algorithmic challenge that needs to be investigated is, what channel (i.e., center frequency and channel-width) should be allocated to which AP, i.e., how much spectrum should each AP get to serve its clients. In this paper, we address this challenge by first proposing a new, simple model that captures adaptive channel-width. Based on this model, we define the *adaptive-width channel-assignment problem*, a problem that turns out to be very interesting practically, but also theoretically. This problem's underlying combinatorial structure is substantially different from previously studied problems. Specifically, whereas the problem of channel assignment in the conventional fixed-channelization framework can be modeled as graph coloring or a variant thereof, e.g. [20], [18], variable channel-widths fundamentally introduces new algorithmic challenges such as *self-fragmentation*. Self-fragmentation can arise because, due to practical hardware constraints, each AP must be assigned a *contiguous band* of spectrum [10]. Hence, if channels are not allocated carefully, the total available spectrum at an AP  $A$  may be fragmented by its neighboring APs, in which case  $A$  can allocate only a small non-overlapping channel even though the total amount of free spectrum may be large.

In WLANs with adaptive channel width, the ultimate algorithmic design goal is therefore to allocate as much spec-

Scenario	$AP_1$	$AP_2$	$AP_3$	$AP_4$	U	SLF
Case 1: (fixed)	1/6	1	1/3	1	4	0.58
Case 1: (adaptive)	2/6	1/2	1/3	1/2	4	<b>0.97</b>
Case 2: (fixed)	1/6	X	1/3	1/2	3	0.82
Case 2: (adaptive)	2/6	X	1/3	1/2	<b>4</b>	<b>0.97</b>

TABLE I  
SPECTRUM-PER-CLIENT (NORMALIZED BY 20MHZ)

trum to the different APs as needed, while 1) avoiding self-fragmentation and 2) still allowing as much spatial reuse as possible. In this paper, we present a compact, but computationally inefficient integer linear program (ILP) that finds the optimal solution. We show that the problem of optimally assigning channels of variable width to APs is NP-complete and we present constant-factor approximation algorithms. Based on these theoretical foundations, we then devise three simple and efficient heuristic approaches. We show that these algorithms achieve close to optimal performance while drastically outperforming existing fixed channelization approaches. In particular, our analytic and evaluation results show that dynamically allocating channels of different widths to APs has the potential of greatly increasing the network’s overall spectrum utilization and fairness in a natural and conceptually simple way. Our evaluations are based on real-world trace-data as well as empirical measurements.

To summarize, we make three primary contributions:

- We explore the use of channel-width as a new, powerful knob in the design of WLANs. Particularly, we quantify the vast potential increase in both spectrum utilization and spectrum fairness that can result from using this knob.
- We define a simple formal model that captures variable-width channels in infrastructure-based networks. Based on this model, we formulate the key spectrum distribution problem as a combinatorial optimization problem and discuss its algorithmically interesting properties.
- We devise a variety of algorithms that efficiently allocate channels of variable width to different APs. We investigate these algorithms by deriving complexity results and analytical worst-case guarantees, and by providing extensive simulation-based comparisons with state-of-the-art fixed channel-width solutions.

## II. MOTIVATION

In existing WLANs, each AP is assigned a fixed width 20 MHz channel, and if possible, neighboring APs are placed on orthogonal frequencies. When the traffic is uniformly distributed across the network, such a scheme increases capacity and reduces interference. However, in dynamic conditions, using fixed-width channels can be problematic and suboptimal. When the number of APs is fewer than the number of available channels, the spectrum is not fully utilized since each AP uses only one channel. On the other hand, if the number of APs is large, two or more neighboring APs are inevitably assigned the same channel, which can create a varying degree of interference [21].

To illustrate how adaptive channel width can help overcome these challenges, we consider a simple example. Figure 1




Fig. 1. A network with four mutually interfering APs. If channel-widths are fixed, each AP is allocated a 20MHz channel. In the adaptive scheme,  $AP_1$  is allocated 40MHz,  $AP_2$  gets 20MHz,  $AP_3$  and  $AP_4$  get 10MHz each.

shows a scenario with four APs all within interference range of one another. In Case 1 (left),  $AP_1$  has 6 clients,  $AP_3$  has 3 clients, while the remaining two APs have one client each. In Case 2 (right), client A moves away from  $AP_2$  and associates to  $AP_4$ . Suppose that all clients generate the same traffic load. We compare the performance of using fixed-width channels with adaptive-width channels. In the fixed-width case, the spectrum is divided into 4 channels of 20 MHz each.<sup>1</sup> In the adaptive-width case, channels may be 10, 20, or 40 MHz. Table I lists the *spectrum-per-client* at each AP. Also included is the total spectrum utilization (U), and a *spectrum-per-load fairness index* (SLF). In the case in which every client generates the same traffic load, this index corresponds to Jain’s fairness index and is computed as  $(\sum c_i)^2 / (n \sum c_i^2)$  summed up over all clients  $i$ ;  $c_i$  is the share of client  $i$ ’s spectrum, and  $n$  is the total number of clients.

In Case 1, fixed-width channelization leads to an unfair spectrum distribution among different APs. A client associated  $AP_2$  or  $AP_4$  can make use of the entire 20MHz spectrum, whereas the same 20MHz spectrum has to be shared among 6 clients at  $AP_1$ . In contrast, with an allocation of 40 MHz to  $AP_1$ , 20 MHz to  $AP_2$  and 10 MHz to the remaining APs, the spectrum-per-load distribution improves significantly because APs with many clients ( $AP_1$ ) receive a wider part of the spectrum to serve its clients. Adaptive channelization can also help to improve system capacity. In Case 2, for instance, if client A moves from  $AP_2$  to  $AP_4$ , an adaptive approach can reallocate the 10 MHz spectrum formerly used by  $AP_2$  to  $AP_4$ , thus giving  $AP_4$  a total of 20 MHz.

**Existing approaches to load-balancing:** There exist several alternative means of alleviating the load-imbalance problem illustrated in Figure 1. One idea, for instance, is to balance the load by assigning some clients to more distant APs or by adjusting transmission powers [6], [21], [8]. In comparison to these solutions, leveraging the ability to adaptively change channel-widths can provide a conceptually simple and more natural solution. Instead of trying to artificially balance the load across APs, this scheme retains the natural locality-induced client-AP association and simply assigns the spectrum according to the specific needs. This does not force clients to associate to far-away APs, and therefore does not reduce their data rate. Each AP can then use any MAC-layer protocol (e.g., CSMA, TDMA, etc...) to fairly and efficiently distribute the spectrum among its clients.

<sup>1</sup>Since there are in fact only 3 completely non-overlapping channels, our estimate for the fixed-width case is optimistic.

### III. DESIGN APPROACH

#### A. System Architecture – Overview

We consider a network architecture based on the techniques developed in [10], in which the channel width of different APs and clients can be changed adaptively, based on their respective traffic load. That is, each AP is allocated a certain contiguous part of the spectrum (a channel of a certain width), which it can then use to serve its clients. Notice that this general architecture involves two, potentially orthogonal, problems: 1) How to distribute the spectrum to the different APs and 2) how each AP uses its allocated channel to serve its associated clients. In this paper, we consider only the first of these problems and assume that within an AP—i.e., for the communication between an AP and its associated clients—any existing contention resolution protocol may be used (CSMA, TDMA,...). In any case, measurements show that APs that are allocated a wider part of the spectrum can deliver more throughput to their clients [10].

Our algorithms are targeted for enterprise networks in which all APs are connected via a backbone network. Each access point is capable of gathering some measure that represents its current *traffic load*. The simplest possible such *load* measure would be the number of clients currently associated with this AP. However, more sophisticated and accurate measures could take into account the traffic demands of each client. Each AP periodically reports its load to a centralized server that is attached to the network’s backbone network and maintains a view of the traffic distribution across the network in a local database.<sup>2</sup> Periodically, the centralized server—based on information stored in its database—runs one of our spectrum distribution algorithms, and assigns a channel-width and center-frequency to each APs. The APs then inform their clients of the new communication channel, upon which they switch to the new channel. Notice that this can be done at little overhead and without breaking any connections [10].

#### B. Model Abstractions

In order to study the potential of load-aware channel-width allocation to APs, we develop a simple model that captures the essence of the spectrum distribution problem: what channel (channel-width and center-frequency) to assign to each AP. It also allows us to analyze and understand the respective merits of different allocation algorithms. The model makes the following key abstractions:

- When setting up and managing a WLAN network, several degrees of freedom may be tuned to optimize the network’s throughput and/or fairness, including transmission powers [6], client-AP association schemes [8], modulation schemes, density of deployment, and even the locations of the APs. In the sequel, we assume these variables to be

<sup>2</sup>Alternatively, using more decentralized, distributed solutions are also possible and an interesting direction for future research. Since the main focus of this work is to identify and quantify the potential gain when abandoning fixed-width channels in WLANs, we focus on the conceptually simpler centralized solution.

fixed, which keeps our results clean from complex inter-dependencies.

- As mentioned above, we are not concerned with the exact protocol that APs use to communicate with the associated clients. In practice, experimental measurements support the intuition that the total throughput achieved by all clients associated to an AP grows linearly in the channel-width used by that AP and its clients [10], [26]. This is also in line with Shannon’s capacity formula.
- We use the simple and standard conflict-graph-based model of inter-AP interference. While there has recently been important work on studying interference models that more accurately capture physical reality (e.g. [22], [23]), our choice is justified for two reasons. First, it is conservative and ignores additional optimizations that could further enhance our system. Second, as we discuss next, our algorithms are designed to allocate non-overlapping spectrum bands to neighboring APs, deliberately limiting interference among APs and their clients.

#### C. Non-Overlapping Channel Assignment

Empirical measurements show that whenever possible, neighboring APs are best assigned non-interfering parts of the spectrum. In traditional WLANs with fixed channelization, it is often unavoidable to assign overlapping channels to neighboring APs if the number of APs in a vicinity exceeds the number of independent channels. Several works have investigated the cost of such overlaps and proposed corresponding algorithms [20]. In contrast, having adaptive channel-widths provides much more flexibility in avoiding overlapping channels. For instance, if the minimum channel-width option is 5MHz, as many as 16 APs in close physical proximity can be assigned mutually orthogonal channels in a 80MHz total spectrum. In an ideal setting with unlimited channel width options, overlapping channels could always be avoided altogether.

This raises the interesting question whether neighboring APs should always be allocated non-overlapping channels? Specifically, one conceivable alternative to assigning non-overlapping channels is to allocate the entire available spectrum to all APs, and purely rely on the 802.11 contention mechanism to access the spectrum. Theoretically, such a scheme should have several pitfalls. Transmissions using wider channels can be more susceptible to interference [26], [10]. Furthermore, building a wide-band receiver to cover the entire spectrum (say 160 MHz in the 5.8 GHz band) is expensive, and will consume more battery power. Finally, the throughput of interfering APs drops more than half if they use the same portion of the spectrum due to contention overhead and losses caused by interference.

Recent empirical measurements in [10] suggest that assigning non-overlapping channels is indeed beneficial. Specifically, it was shown that in spite of some cross-channel leakage, the average combined throughput of the two flows when sharing one channel of width  $2X$  MHz width is less than when they are split on adjacent non-overlapping  $X$  MHz channels.

This gain of partitioning the spectrum is particularly striking if different flows transmit at different data rates because of the *rate anomaly problem* [16]. The other reasons for the gain stems from reduced contention overhead and from the fact that narrower channels have a smaller per-packet relative overhead [10].

**Simulations in QualNet:** We corroborate these findings with a simulation-based evaluation of the 6 AP scenario of Section VI (Figure 3) in QualNet [2]. Six clients are associated to each AP, and each AP is sending separate CBR flows to every client. We measured the total throughput of the system in two scenarios: (i) APs using the 4 orthogonal channels used in our deployment at 12 Mbps each (We used the channel allocation presented in Chandra et. al. [11]); and (ii) all APs on one channel using 48 Mbps data rate. The second scenario is an approximation of all APs using the entire spectrum (4 channels), i.e. they can pack 4 times the number of bits per second. To factor out the impact of encoding (and hence data rates) on packet losses, we place the clients very close to the AP they are associated with. The key results are that for these scenarios, the total throughput of the system was 30.4 Mbps when APs used 4 channels. In contrast, the throughput drops to 17.85 Mbps when all nodes shared the same channel, albeit at 4 times the data rate.

Both our empirical measurements and our simulations shows that having two APs on two separate channels gives better performance than having both on a single channel of double the channel width. And for this reason, we have designed our algorithms in Section V to always try and assign non-overlapping channels. We note, however, that a partial overlap in the AP's channels *might* give better system throughput than completely non-overlapping channels as has been shown in a recent work [22]. We are currently extending our model to capture this phenomenon.

#### IV. PROBLEM FORMULATION & COMPLEXITY

Based on the discussion in the previous section, we formalize the spectrum allocation problem as follows. Consider a WLAN consisting of  $n$  access points  $AP_1, \dots, AP_n$ . Given the fixed locations and transmission powers, let  $G = (V, E)$  be the network's *conflict-graph* [17], [25]: There is an edge between two APs if they have significantly overlapping coverage regions and should therefore avoid transmitting on the same frequency.<sup>3</sup> The conflict graph is static and is updated only rarely. For an  $AP_i$ , we denote by  $N(i)$  the set of all neighboring APs that are potentially in conflict with  $AP_i$ ,  $N(i) = \{AP_j \mid (i, j) \in E\}$ .

Each AP  $AP_i$  has a *load*  $L_i$  associated to it. The load corresponds to the amount of traffic the AP needs to service. In practice,  $L_i$  can either be determined using simple heuristics

<sup>3</sup>In a practical system, the interference relationship between neighboring APs can be determined in an ad hoc fashion (e.g., by APs using beacon messages to probe their proximity to other APs, or by client feedback) [25]; or it may be statically provided as part of the network planning. Also, more sophisticated conflict-graphs can be used if desired. For instance, it can take into account interference between clients and client-locations as in [21].




Fig. 2. Network in which an optimal solution in terms of spectrum utilization is unfair. T and F denote the allocations in a spectrum-optimal and fair solution, respectively.

(e.g.,  $L_i$  is the number of clients associated to the AP) or by more sophisticated means that include aspects such as the clients' data rates.

**Algorithms:** The set of loads  $L_1, \dots, L_n$ , along with the interference graph, forms the input to a *spectrum assignment algorithm*. This algorithm assigns a channel  $\mathcal{I}_i = [S_i, S_i + B_i]$  to every  $AP_i$ , where  $S_i$  is the assigned channel's lower-end frequency, and  $B_i$  is the channel-width.

**Measures:** For a given spectrum assignment, we consider two measures for every AP: the *total-spectrum*  $T_i$  and the *per-unit-load-spectrum*  $PL_i$ , or short, PUL-spectrum. The total-spectrum is simply the total spectrum available to  $AP_i$ , i.e.,  $T_i = B_i$  if the channel assigned to  $AP_i$  is non-overlapping. The PUL-spectrum is defined as  $PL_i = B_i/L_i$  and captures how much spectrum is allocated to  $AP_i$  per unit of load that this AP has to serve.

Given these definitions, we define the non-overlapping version of our key spectrum distribution problem.

**Adaptive-Width Channel-Assignment Problem:** Given AP loads  $L_1, \dots, L_n$  and an conflict-graph  $G = (V, E)$ , find a an assignment of a contiguous channel  $\mathcal{I}_i = [S_i, S_i + B_i]$  to each AP. An assignment is called *non-interfering* if, for any pair of neighboring APs,  $AP_i$  and  $AP_j$  with  $(i, j) \in E$ , the assigned channels  $\mathcal{I}_i$  and  $\mathcal{I}_j$  are non-overlapping.

There are two important objectives that an adaptive-width channel-allocation should achieve: 1) *high spectrum utilization* and 2) *per-load fairness*.

**Spectrum Utilization  $T_{Sys}$ :** The goal is to reuse as much spectrum as possible in the system, i.e., to maximize the overall spectrum utilization  $T_{Sys} = \sum_{i \in V} T_i$ .

**Per-Load Fairness:** For fairness, various definitions can be considered and the optimization criterion can be defined appropriately. We incorporate fairness by requiring a minimum PUL-spectrum fairness across all nodes. Specifically, we consider a *local fairness condition* in which every AP must receive at least its fair share of spectrum in its neighborhood.<sup>4</sup> In particular, we define  $\phi(i) = L_i / (L_i + \sum_{j \in N(i)} L_j)$  as the *fair spectrum-share* that  $AP_i$  should receive. APs with high load  $L_i$  get a proportionally larger share of the spectrum. In the *Fair Adaptive-Width Channel-Assignment Problem*, we seek to maximize the spectrum utilization under the condition that every AP receives at least a spectrum of  $B_i \geq \alpha \phi(i) \cdot B_{tot}$ ,

<sup>4</sup>Other fairness definitions can be incorporated into our problem formulation. Our fairness definition has the advantage of being entirely *local*. That is, if there are dense and sparse regions in the network, the naturally higher fairness restrictions in the dense network parts do not affect the sparse regions.

where the parameter  $\alpha$  characterizes the maximally tolerable unfairness in the system. Note that this fairness definition implies a local PUL-spectrum fairness constraint.

**Fairness vs. Capacity Trade-off:** The parameter  $\alpha$  characterizes the trade-off between achieving high spectrum utilization and fairness, which are typically contradicting aims. Consider the star graph with uniform demands shown in Figure 2. An allocation maximizing spectrum utilization assigns each leaf AP the entire spectrum, while giving no channel to the center AP. While achieving spectrum utilization (the entire spectrum is reused 4 times), such a solution starves clients associated to the AP in the center. A completely fair solution, on the other hand, consists of assigning each AP a channel-width spanning half of the totally available spectrum. The above problem definition addresses this fairness vs. throughput trade-off by fixing a lower bound on the degree of fairness that must be maintained between different APs.

**Complexity:** We can show that the fair dynamic-width channel-assignment problem is NP-hard for  $\alpha > 2/3$ . Due to lack of space, the proof is deferred to the full version.

*Theorem 4.1:* The dynamic-width channel-assignment problem is NP-hard for any fairness parameter  $\alpha > 2/3$ . This holds even in restricted geometric graph models such as the unit disk graph.

## V. ALGORITHMS

As motivated in Section III-C, our algorithms are designed to assign non-overlapping channels to neighboring APs whenever possible, by making the channels narrower if needed. The *adaptive-width channel-assignment problem* is different from *coloring problems* or *multicoloring problems* that have been extensively studied in the networking and theory communities. The reason is that, unlike in (multi)coloring problems, the interval assigned to each AP must consist of a *contiguous* chunk of spectrum. This contiguity constraint can lead to *fragmentation* of the spectrum; APs may be unable to reserve a large contiguous block of the spectrum even though the totality of unused spectrum in their neighborhood would be sufficiently high. The problem is therefore also of theoretical interest. We devise three different algorithmic approaches.

### A. Optimal ILP Formulation

The optimal solution of a problem instance can be characterized by the following integer linear program (ILP). Variables  $b_i$  and  $s_i$  denote the channel width and lower-end frequency allocated to  $AP_i$ . The ILP determines the optimal spectrum utilization in a network with arbitrary channel-width options:

$$\begin{aligned} \max \quad & \sum_{AP_i \in V} b_i \\ s_i + b_i - s_j - f_{ij} \cdot B_{tot} & < 0, \quad \forall (i, j) \in E \\ s_j + b_j - s_i - f_{ji} \cdot B_{tot} & < 0, \quad \forall (i, j) \in E \\ f_{ij} + f_{ji} & \leq 1, \quad \forall (i, j) \in E \\ s_i + b_i & \leq F_{top}, \quad \forall AP_i \in V \\ s_i & \geq F_{bottom}, \quad \forall AP_i \in V \\ \chi \cdot b_i & \leq L_i, \quad \forall AP_i \in V \\ b_i & \geq \alpha \phi(i) \cdot B_{tot}, \quad \forall AP_i \in V \\ f_{ij}, f_{ji} & \in \{0, 1\}, \quad \forall (i, j) \in E. \end{aligned}$$

For each pair of APs  $AP_i$  and  $AP_j$  with  $(i, j) \in E$ , the ILP has two binary indicator variables  $f_{ij}$  and  $f_{ji}$ . The first two constraints force these variables to behave as follows. The variable  $f_{ij}$  is 1 if and only if the top-frequency  $s_i + b_i$  of  $AP_i$  is “above” the bottom-frequency  $s_j$  of  $AP_j$ . Conversely,  $f_{ji} = 1$  if and only if  $s_j + b_j > s_i$ . Considering two intervals  $[s_i, s_i + b_i]$  and  $[s_j, s_j + b_j]$ , it is easy to observe that these intervals overlap if and only if  $s_i + b_i > s_j$  and  $s_j + b_j > s_i$ . The third constraint therefore guarantees that no two neighboring intervals in the graph overlap, i.e., the resulting channel assignment is non-overlapping. The next two constraints ensure that the assigned interval is located within the available spectrum  $[F_{bottom}, F_{top}]$ . The sixth one expresses that above a certain load-dependent point, raising the channel width does not increase throughput. Finally, the last inequality captures our local *per-load fairness* condition as discussed in Section III. Note that the ILP assumes channel-widths to be arbitrarily tunable. However, discrete widths can easily be incorporated by restricting the variables  $b_i$  to a corresponding set of integers. In Section VI, we examine the impact of this discrete set of channel width options.

While the ILP describes the theoretical optimum of any problem instance, it is computationally practicable only in small networks. Therefore, we now investigate computationally efficient approximate solutions.

### B. LP-Based Approximation

Our problem has the intricacy that *fragmentation* needs to be avoided, which makes it very difficult to formulate the problem as a linear program (LP) relaxation.

**A Packing Algorithm that avoids Fragmentation:** We start by considering the interval packing problem in isolation. Assume for the moment that the channel-width intervals allocated to the different APs was already determined. How should these intervals be packed? Intuitively, adhering to the following rules of thumb may help:

- R1. Pack large items first.
- R2. Try to fill up from one end.

However, in addition to being a packing problem, the problem also has the flavor of a complex (interval) *coloring* problem. The well-known greedy coloring algorithm visits nodes one-by-one, assigning each node the lowest color still available in its neighborhood. Clearly, this procedure colors any graph using at most  $\Delta(G) + 1$  colors, where  $\Delta(G)$  is the maximum node-degree. Similarly, if we were not constrained to assigning a *contiguous interval* to each AP, we could assure that all required spectrum can be packed in a total spectrum of

$$\delta(\mathbf{b}) \triangleq \max_{u \in V} \left( b_u + \sum_{v \in N(u)} b_v \right), \quad (1)$$

which is essentially the continuous counterpart of the  $\Delta(G) + 1$  coloring upper-bound. That is, without the contiguity constraint, the greedy coloring algorithm assures that the total spectrum requirement is  $\delta(\mathbf{b})$ . Based on these observations, we now present an approximation algorithm that combines both the packing and coloring aspects of the problem.

Assume that the widths of all intervals followed a power series, i.e., each interval has length  $2^k$  for some integer  $k$ . Applying rule of thumb #1, we sort the items in decreasing order of their sizes and try to pack them one by one into the real axis  $[0, +\infty]$ . Applying rule of thumb #2, when packing each item, we always try to fill up from one end, closer to the origin. When packing in this way, it can be proven by induction that whenever an interval of size  $2^k$  is packed, all available chunks of the spectrum are of size at least  $2^k$  (in fact, they are an integer multiple of  $2^k$ ). Hence, in this case, we do not suffer from fragmentation: when packing intervals in this order, every interval can be placed in such a free spectrum chunk. And as pointed out before, the total spectrum required to pack all intervals is at most  $\delta(\mathbf{b})$ . Therefore, this method achieves for the joint packing and coloring problem the same performance that one can achieve for coloring.

If the channel-width intervals to be packed do not follow a power series, we can round them up accordingly. Suppose the given interval lengths are  $b_0 \geq b_1 \dots \geq b_N$ . Then we round each  $b_i$  to the next higher power of 2. Consequently, all intervals can be packed within a spectrum of

$$\max_{u \in V} \left( \tilde{b}_u + \sum_{v \in N(u)} \tilde{b}_v \right) \leq 2\delta(\mathbf{b}). \quad (2)$$

Finally, we can linearly map the assigned frequencies in  $[0, 2\delta(\mathbf{b})]$  to the entire available spectrum interval  $[F_{bottom}, F_{top}]$ . Doing so, we have packed demands  $\mathbf{b}$  in a maximum interval of  $2\delta(\mathbf{b})$ , which is at most by a factor of 2 (due to the rounding) worse than applying the greedy coloring algorithm to a relaxed problem where each node can make use of non-contiguous bands.

**Optimizing the Interval Lengths:** The packing algorithm presented in the previous subsection is effective in assuring the performance for the worst AP (with maximum load in its neighborhood). While this is good from the fairness perspective, it may harm throughput in scenarios in which some parts of the graph are dense, and others are sparse. We now show how the overall spectrum utilization can be enhanced without sacrificing fairness. We use the packing algorithm as a building block that packs any load vector  $\mathbf{b}$  into a spectrum of width  $[0, 2\delta(\mathbf{b})]$ . The idea is to employ a sequence of linear programs to search for a load vector with good worst-case performance  $\delta(\mathbf{b})$  and good spectrum utilization  $T_{Sys}$ . The packing algorithm can then be run over the resulting load vector  $\mathbf{b}$  to pack it into  $[0, 2\delta(\mathbf{b})]$ .

Consider the following linear program:

$$T_{Sys}(\alpha) \triangleq \max_{\mathbf{b}} \sum_u b_u, \quad \text{subject to:} \quad (3)$$

$$b_u \geq \alpha \phi_u \cdot B_{tot}, \quad \forall u \quad (4)$$

$$b_u + \sum_{v \in N(u)} b_v \leq B_{tot}, \quad \forall u. \quad (5)$$

Constraint (5) ensures that the computed vector  $\mathbf{b}$  results in a feasible solution with a greedy coloring algorithm. Constraint (4) maintains fairness by guaranteeing node  $u$  its

fair PUL-spectrum share. By varying the constant scaling parameter  $\alpha$  from 0 to some maximum value  $\alpha^*$ , different tradeoffs between fairness and spectrum utilization can be achieved. Using the maximum value  $\alpha^*$  maximizes the worst node's performance; this value can be determined using the following LP:

$$\alpha^* = \max_{b, \alpha} \alpha, \quad \text{subject to:} \quad (4)(5) \quad (6)$$

**Approximation Guarantees:** For any specific  $\alpha$ , the performance achieved by a slightly adapted version of the LP-based algorithm can be shown to be within a constant factor of the optimal algorithm. The proof is omitted due to lack of space.

*Theorem 5.1:* When modeling the wireless network as a disk graph, the LP-based algorithm achieves a spectrum utilization that is within a constant factor of the optimal solution. This holds for any per-load fairness  $\alpha$  and the constant depends on the underlying network model.

Furthermore, it is worth noticing that our proofs of the packing scheme imply a second important result.

*Corollary 5.2:* Every non-contiguous, but otherwise feasible spectrum allocation to APs can be turned into a feasible contiguous spectrum allocation at the cost of at most a factor of 2 in spectrum utilization.

**Practical Deployment:** Our LP-based algorithm leaves open various parameters for tuning the involved fairness vs. spectrum utilization trade-off. A simple way of employing it in practice is the following: First, determine the optimal fairness parameter  $\alpha = \alpha^*$  using LP 6. Then, using this  $\alpha$ , use the first LP to compute  $B_{total}(\alpha)$ . This amounts to a conservative approach that maximizes the total spectrum utilization (by "flattening" the loads at the nodes) while assuring the maximum level of fairness at the worst node. The LPs can either be solved directly using an LP solver, or we can apply efficient approximation algorithms [15].

### C. GreedyRaising: Simple Greedy Heuristics

The LP-based approximation algorithm in Section V-B provides provable worst-case performance guarantees. We now apply the lessons learned in the previous section and devise simpler heuristic solutions that do not require solving a linear program and, as we show in Section VI, still perform well.

The heuristics are based on the greedy-packing subroutine shown in Algorithm 1. This greedy packing routine takes as input an ordering of APs (for example, from heaviest to lightest load) and a channel-width for each AP. It then proceeds in order of the given ordering and, when considering  $AP_i$ , greedily attempts to pack a non-overlapping channel of width  $B_i$  into the spectrum. If there is a choice, the interval is packed at the lowest possible frequency.

Depending on the given ordering and channel-widths, the greedy-packing scheme may not succeed. If the desired channel-widths are too wide, it is theoretically impossible to correctly pack. However, even if it is theoretically possible to achieve a valid assignment of intervals to APs, the greedy allocation may make suboptimal decisions and get stuck in the process. In this case, the subroutine returns false, thereby indicating that narrower channel-widths should be tried.

---

**Algorithm 1** *GreedyPack*( $B_1, \dots, B_N, \mathcal{O}$ ) Routine

---

Input: Channel widths  $B_1, \dots, B_N$  and an ordering  $\mathcal{O}$  of APsOutput: If possible, a non-overlapping packing of channel widths into the available spectrum.  
Return false if no packing is found.

- 1: In the order of  $\mathcal{O}$ : **for each**  $AP_i \in V$  **do**
  - 2: pack an interval of channel-width  $B_i$  in the lowest possible non-overlapping frequency.
  - 3: **end for**
  - 4: **if** the interval of all APs was successfully packed within the total spectrum  $[F_{bottom}, F_{top}]$  **then**
  - 5: **return true**
  - 6: **else return false**
  - 7: **end if**
- 

---

**Algorithm 2** *GreedyRaising* Algorithm

---

Input: An ordering  $\mathcal{O}$  of APs

Output: A non-overlapping packing of channel widths intervals in the available spectrum.

- 1: Set parameters  $\theta_{upper} := BIGVALUE$  and  $\theta_{lower} := 0$  and let *successful* := *FALSE*;
  - 2: **while**  $\theta_{upper} - \theta_{lower} > 0.01$  **do**
  - 3:  $\theta := (\theta_{upper} + \theta_{lower})/2$ ;
  - 4: Let  $\phi'_i := \theta \cdot L_i / (L_i + \sum_{j \in N(i)} L_j)$  for each  $AP_i \in V$ .
  - 5: Let  $B_i$  be the largest channel width option s.t.  $B_i \leq \phi'_i \cdot B_{tot}$
  - 6: *successful* := *GreedyPack*( $B_1, \dots, B_N, \mathcal{O}$ ).
  - 7: **if** *successful*
  - 8: **then**  $\theta_{upper} := (\theta_{upper} + \theta_{lower})/2$ ;
  - 9: **else**  $\theta_{lower} := (\theta_{upper} + \theta_{lower})/2$ ;
  - 10: **end if**
  - 11: **end while**
  - 12: In the order of  $\mathcal{O}$ : **for each**  $AP_i \in V$  **do**
  - 13: Let  $\hat{B}_i$  be the next higher channel width option of  $B_i$ .
  - 14: *successful* := *GreedyPack*( $B_1, \dots, \hat{B}_i, \dots, B_N, \mathcal{O}$ ).
  - 15: **if** *successful* = *TRUE* **then**  $B_i := \hat{B}_i$ .
  - 16: **end for**
- 

The basic idea of our GreedyRaising heuristics is the following. Starting from a feasible initial assignment, the heuristics “probes” APs one-by-one and checks whether greedy-packing remains successful if the AP’s channel-width is raised. More specifically, GreedyRaising considers all APs in a given sequence  $\mathcal{O}$ . When considering an AP, its channel-width is increased to the next higher channel-width option, and the greedy-packing subroutine is called in order to see whether it still succeeds. If it does, the higher channel-width is adopted; if not, its channel-width is reset to its original value.

The only thing that remains to be defined is the ordering  $\mathcal{O}$  in which APs are considered in both the greedy packing subroutine and the main algorithm. We distinguish two different orderings and evaluate their relative merits:

**Most-Congested-First:** In this ordering, APs are sorted in decreasing order of their load.

**Smallest-Last:** Consider an ordering  $\mathcal{O}$  and let  $\tau_i$  be the number of APs that are neighbors of  $AP_i$  that appear *before*  $AP_i$  in  $\mathcal{O}$ . The smallest-last ordering is an ordering that minimizes the maximum  $\tau_i$  over all APs in the network [19]. This ordering has been studied in the context of coloring problems and is based on the following observation. When




Fig. 3. Floor plan and AP locations on the floor of an office building. Solid lines represent two interfering APs, and dashed lines indicate that APs interfere at one of the clients.




Fig. 4. Number of active clients at different times of the day.

considering  $AP_i$  in the greedy-packing routine,  $\tau_i$  reflects the number of potentially interfering intervals that have already been packed in  $AP_i$ ’s neighborhood. Intuitively, the fewer such intervals, the easier it is to pack  $AP_i$ ’s allocated channel. Considering the APs in smallest-last order minimizes the maximum obstruction that any AP faces when its channel width interval is packed. The smallest-last ordering can be computed efficiently in a single pass [19].

Our evaluations show that both GreedyRaising heuristics perform similarly well and outperform schemes based on fixed channels. In many network topologies, the smallest-last heuristic has the tendency to perform somewhat better.

## VI. EVALUATION

**Simulation Settings:** We evaluate our algorithms using two data sets: a small-scale enterprise WLAN and a large enterprise/campus WLAN deployment.

*Small WLAN:* We use the wireless usage data collected in [11]. This dataset contains monitoring information of 6 APs on the floor of an office building. The floor plan of the office building and the AP locations is illustrated in Figure 3. Lines between APs show the conflict (interference) relationship. Two APs are treated as interfering if there is one client associated with one AP that can hear beacons from the other AP.

Figure 4 shows the maximum number of clients that are simultaneously associated to each AP during every hour from 8 AM to 8 PM on Tuesday of a work week. The plots for the other days show a similar trend; they are omitted in the interest of space. Clearly, there is a spatial and time disparity in network usage across different APs. At any given time, APs at some locations serve a significantly larger number of clients than the others. For example, AP 4, which is located close to several conference rooms had up to 16 clients from 2 PM to 3 PMs. Furthermore, client populations at the APs






Figure 5: Performance comparisons for the small network: fixed channel-width vs. ILP with 4 channel width options vs. ILP with 11 channel width options vs. ILP without discrete channel width constraints vs. ideal approach.

vary significantly over time. The set of heavily-loaded APs also changes at various times of the day across different days.

The conflict graph of Figure 3, together with the number of clients associated to each AP for each hour, constitute a test sample. We use the data for Monday to Thursday, each covering 13 hours; hence there are altogether  $13 \times 4 = 52$  tests. We will present the results for 6 tests covering cases with low, medium, and high variances in the client-distribution.

*Large WLAN:* This test scenario is generated based on measurement data reported in [7], which measured the number of clients associated to each AP for a network of 177 APs that spans three buildings. The result is reported as a cumulative distribution function of the number of clients associated to an AP. In this trace, 50% of the APs serve less than 5 users, while 10% of the APs serve over 15 users; the average number of clients served by each AP is 8. We generated a 50 AP test scenario by drawing according to the distribution reported in [7]. Then we randomly placed the APs in a  $1000\text{m} \times 1000\text{m}$  square area. We generated two conflict graphs, for interference ranges of 50m and 75m, respectively. They constitute two tests, which we label as “sparse” and “dense” in the plots.

Using the test scenarios mentioned above, we first compare the performance of the proposed adaptive-width scheme with (i) a state-of-the-art solution that uses fixed-width channels, and (ii) the ideal upper-bound assuming best possible adaptation. For brevity, the former will be referred to as “fixed-width approach” and the latter will be referred to as “ideal approach”. The ideal approach can be viewed as either a perfect time-domain adaptation approach or a perfect frequency-domain adaptation approach where an AP can be allocated multiple non-contiguous bands of arbitrary width. Specifically, here we are evaluating the performance of the ILP scheme described in Section V-A. This serves as the optimal spectrum-domain adaptation scheme without complexity constraints. In the next subsection we will evaluate how far the low-complexity approximation algorithms can approach the ILP’s performance.

The overall available spectrum is assumed to be 80 MHz. We consider three variants of the ILP scheme. The first one assumes a discrete set of allowed channel-widths,  $\{5, 10, 20, 40\}$  MHz. The second one assumes a different set of channel-widths,  $\{3, 5, 6, 7, 10, 12, 14, 20, 24, 28, 40\}$  MHz (this is the set of channel-widths defined in WiMAX). The third variant assumes any channel-width is allowed.

**Performance Metrics:** We use three performance metrics. The first metric is the *total spectrum utilization*  $T_{Sys}$ . The second metric is a *global spectrum-per-load fairness* (SPF)

$F_{global}$  that is similar to Jain’s fairness index. Specifically, the global fairness index is defined as follows: For every unit of load  $j$  at an  $AP_i$ , we define  $C_j = T_i/L_i$ . The fairness index is then, i.e.  $(\sum C_j)^2 / (n \sum C_j^2)$ , where the sums are over all units of load and  $n$  is the total number of clients. This global fairness index reflects the uniformity of spectrum obtained by every unit of traffic load in the network. The above fairness index has shortcomings. In a network with dense and sparse parts, a completely fair solution according to  $F_{global}$  may be far from a good solution, because it wastes potential spectrum in the network’s sparse parts. To reflect this consideration, we also consider the *local per-load fairness* introduced in Section IV. That is, this local fairness index, denoted by  $F_{local}$ , is  $F_{local} = \min_{AP_i} \frac{B_i}{L_i / (L_i + \sum_{j: j \in N(i)} L_j) \cdot B_{total}}$ . It is the ratio of an AP’s allocated spectrum  $B_i$  over the AP’s fair spectrum share, minimized over all APs.

**Small WLAN Results:** We begin with the small WLAN setup shown in Figures 3 and 4. Note that the minimum number of colors in a proper vertex coloring of this conflict graph is 4. Hence, the “fixed-width approach” assigns 4 orthogonal channels to the APs, each of width 20MHz. The four channels correspond to channels 1, 4, 8, and 11 in the IEEE 802.11 standard. Although in reality, these channels are not completely non-overlapping, in order to make a conservative comparison with our adjustable channel width allocation, we assume the fixed-width approach can assign four completely non-overlapping channels.

Figure 5 shows the performance comparisons for the small network, for the three performance metrics discussed above. In terms of spectrum utilization, the fixed channel-width approach achieves 144 MHz whereas the ideal bound is 192 MHz in all six scenarios. The ILP without discrete channel-width constraints achieves 192 MHz, showing that requiring the band to be contiguous does not decrease the spectrum utilization. The average spectrum utilization of ILP with 11 channel-width options is 182.4 MHz, which is better than the state-of-the-art fixed channel-width approach. In terms of global fairness index  $F_{global}$ , the average value for the fixed channel-width approach is 0.54; the ILP with 4 channel-width options achieves 0.63; the ILP with 11 channel-width options achieves 0.65; the ILP without constraints as well as the ideal approach achieve 0.77. Hence our algorithms provide slightly better global fairness than the fixed channel-width approach. However, similarly as in the case of spectrum utilization, our algorithms are significantly better in terms of the local fairness






Figure 6: Large network. Spectrum utilization (left), Global fairness index (middle), Local fairness index (right)

index  $F_{local}$ <sup>5</sup>. Again, the fixed width approach is the worst, whereas the ILP without discrete channel-width constraints has about the same performance as the ideal solution.

**Large WLAN Results:** To compare with existing WLAN deployments in which each AP is assigned one (fixed width) channel, we compute the spectrum utilization with existing deployment for the 50-node topology as follows. Since the conflict graph  $G$  in this case is not colorable with four colors, we cannot avoid conflicting channel assignments. We proceed by first obtaining a (conflicting) channel assignment on four channels and then compute the spectrum utilization assuming that if an AP conflicts with  $k$  of its neighboring APs on the same channel, then this AP can use effectively a fraction  $1/(k+1)$  of its assigned 20MHz spectrum. To obtain a good channel assignment, we use the weighted coloring channel assignment algorithm in [20].

Figure 6 gives the performance comparisons for the larger WLAN, for the two test cases. In both cases, the spectrum utilization of the fixed channel-width approach is significantly worse than the other approaches (up to 4X smaller). Similar to the small WLAN case, the ILP with discrete channel width options achieves a much better spectrum utilization, and the ILP without discrete constraints achieves a spectrum utilization close to that of the ideal solution. However, in this case, having 11 channel-width options does not significantly improve the spectrum utilization. In terms of the local fairness index, the fixed channel-width approach can be as much as 8X worse than the the adaptive channel-width approach.

**Low Complexity Approximation Algorithms:** The ILP scheme described in Section V-A, although optimal, has a high complexity, limiting its use to only small networks. In this subsection we compare the ILP scheme with the low complexity approximation algorithms in Section V-B and Section V-C. Figure 7 shows the throughput comparisons for the small network and the large network, where the ILP approach is compared with (i) the LP approach followed by greedy tuning, (ii) the greedy raising algorithm using the most-congested-first ordering, (iii) the greedy raising algorithm using the smallest-last ordering. In this setting, there are four channel-width options,  $\{5, 10, 20, 40\}$  MHz; the results for 11 options are similar and are omitted. The crucial observation is that the low complexity heuristics algorithms can achieve close to optimal spectrum utilization. The same is true for both the global and local fairness index measures.

<sup>5</sup>The local fairness index can exceed 1 if the spectrum assigned to every AP exceeds its fair share. This can happen if spectrum is reused very efficiently by the algorithm.




Figure 7: Comparison of spectrum utilization—Large network (top), Small network (bottom): ILP vs heuristics. “LP-4widths” refers to the LP algorithm using 4 channel-width options. “GR-MCF-4widths” refers to the greedy raising algorithm with most-congested-first ordering and 4 channel-width options.

**Summary of Observations:** In summary, the proposed adaptive channel-width architecture, which allocates each AP a contiguous channel with a tunable width, can significantly outperform the fixed channel-width approach in terms of spectrum utilization. In some cases, the increase in spectrum utilization can be as much as 4X. Secondly, adapting channel-widths based on the loads at the APs also helps in substantially improving fairness in the network, thus providing a good solution to the load-balancing problem. And finally, the performance of our heuristic algorithms is close to the optimal upper bound that assumes non-contiguous, arbitrary-width channels.

## VII. RELATED WORK

Several algorithms and techniques have been proposed to balance load across APs in a WLAN.

Bejerano et. al. [8] showed the importance of association control for fairness and load balancing in WLANs. Cell Breathing [6] adjusts power levels of APs to achieve association control. It decreases the transmit power of heavily loaded APs to force some clients away from it. Similarly, it increases the transmit power of lightly loaded APs to attract clients. Other approaches, e.g. [27], explicitly modify clients to associate with the best AP. Although association control is important, it might not be the best approach for load balancing if done in isolation. For example, cell breathing might associate a client with a farther away AP. Although, this can reduce the number of clients associated to an AP,

moved clients might get much lower data rates if the new AP is far, or they could reduce the throughput of other clients around it due to the rate anomaly problem [16]. In contrast, our approach of load-aware spectrum allocation allows clients to stay associated to the best APs. We note, however, that association control is complementary to our approach of using adaptive channel widths, and both these mechanisms can be combined for better WLAN performance.

Other approaches such as SMARTA [5], MDG [9], or DenseAP [24] have realized the limitations of performing association control in isolation, and propose an integrated approach for improving WLAN performance. They adapt the AP transmission power and channel, and also control client AP associations. Furthermore, these systems propose dense AP deployments to reduce the client-unfairness problem. Each AP only serves a few clients, and therefore the impact of client-unfairness is reduced. However, the problem still exists as long as each AP serves more than one client and the client distribution is not uniform. In these scenarios, SMARTA, MDG, and DenseAP can benefit from using channel width as an additional tuning parameter.

There exists prior work on allocating channel widths based on demand in other types of networks. In OFDMA [14], an AP can split its channel width among its clients to reduce the impact of interference. However, the allocated channels can be fragmented. Further OFDMA work deals with allocating channel widths to single communication links between clients, and not with balancing load among various APs [10]. In another work, the KNOWS system [28], [29] proposes the use of contiguous bands to maximize the usage of fragmented spectrum in mesh networks. To the best of our knowledge, ours is the first work to use channel width as a parameter to solve the load balancing problem in WLANs.

## VIII. CONCLUSION

The fixed-width channelization technique in IEEE 802.11 networks inherently has difficulties in coping with the spatially non-uniform and temporally dynamic user demand that is prevalent in many infrastructure networks deployed today. To remedy this situation, various load-balancing techniques have been investigated. In this paper, we argue that we can leverage recent work on adaptive channel-width systems to move beyond pre-determined channels of fixed width. This provides a very powerful and conceptually simple solution to the load-balancing and spectrum allocation problem in WLANs; a solution that has the potential of significantly increasing both fairness and spectrum utilization.

Before this potential can ultimately be tapped, a number of practical issues need to be solved. For example, the IEEE 802.11 scanning process will become more heavyweight, since every client may have to try different channel-width options on every center frequency to determine the presence of an AP. We are currently in the process of solving this and other problems on a real-testbed. There are also numerous directions for future algorithmic research, such as incorporating power control or client-AP association into the optimization framework.

## REFERENCES

- [1] Atheros Communications, <http://www.atheros.com/pt/whitepapers/atheros-superg-whitepaper.pdf>.
- [2] QualNet, <http://www.qualnet.com/>.
- [3] Wimax forum whitepapers: <http://www.wimaxforum.org/>.
- [4] IEEE 802.11-2007: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 2007.
- [5] N. Ahmed and S. Keshav. SMARTA: A Self-Managing Architecture for Thin Access Points. In *Proc of ACM CoNext*, 2006.
- [6] P. Bahl, M. T. Hajiaghayi, K. Jain, V. Mirrokni, L. Qiu, and A. Seber. Cell Breathing in Wireless LANs: Algorithms and Evaluation. *IEEE Transactions on Mobile Computing*, 2006.
- [7] M. Balazinska and P. Castro. Characterizing mobility and network usage in a corporate wireless local-area network. In *MobiSys*, 2003.
- [8] Y. Bejerano, S. Han, and L. Li. Fairness and load balancing in wireless lans using association control. In *Proc of ACM MOBICOM*, 2004.
- [9] I. Broustis, K. Papagiannaki, S. Krishnamurthy, M. Faloutsos, and V. Mhatre. MDG: Measurement-Driven Guidelines for 802.11 WLAN Design. In *Proc of MobiCom*, 2007.
- [10] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A Case for Adapting Channel Width in Wireless Networks. In *Proc of SIGCOMM*, 2008.
- [11] R. Chandra, J. Padhye, A. Wolman, and B. Zill. A Location-Based Management System for Enterprise Wireless LANs. In *Fourth Symposium on Networked Systems Design and Implementation (NSDI)*, April 2007.
- [12] David Kotz and Kobby Essien. Analysis of a Campus-wide Wireless Network. *Wireless Networks*, 11:115–133, 2005.
- [13] Diane Tang and Mary Baker. Analysis of a Local-Area Wireless Network. In *Proceedings of the Annual ACM International Conference on Mobile Computing (MobiCom)*, August 2000.
- [14] K. Fazel and S. Kaiser. *Multi-Carrier and Spread Spectrum Systems*. John Wiley and Sons, Inc., 2003.
- [15] N. Garg and J. Koenemann. Faster and Simpler Algorithms for Multicommodity Flow and other Fractional Packing Problems. In *Proc of FOCS*, 1998.
- [16] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance Anomaly of 802.11b. In *Proc of IEEE Infocom*, 2003.
- [17] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu. Impact of Interference on Multi-hop Wireless Network Performance. In *Proc of ACM MobiCom*, 2003.
- [18] B.-J. Ko and D. Rubenstein. Distributed Self-stabilizing Placement of Replicated Resources in Emerging Networks. *IEEE/ACM Transactions on Networking*, 13(3):476–487, 2005.
- [19] D. W. Matula and L. L. Beck. Smallest-last ordering and clustering in graph coloring algorithms. *Journal of the ACM*, 30(3):417–427, 1983.
- [20] A. Mishra, S. Banerjee, and W. Arbaugh. Weighted Coloring based Channel Assignment in WLANs. *Mobile Computing and Communications Review*, July 2005.
- [21] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. Arbaugh. Client-driven Channel Management for Wireless LANs. In *IEEE INFOCOM*, 2006.
- [22] A. Mishra, V. Shrivastava, S. Banerjee, and W. Arbaugh. Partially Overlapped Channels not Considered Harmful. *SIGMETRICS Perform. Eval. Rev.*, 34(1):63–74, 2006.
- [23] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol Design Beyond Graph-Based Models. In *ACM HotNets*, 2006.
- [24] R. Murty, J. Padhye, R. Chandra, A. Wolman, and B. Zill. Designing High-Performance Enterprise Wireless Networks. In *Proc of NSDI*, 2008.
- [25] J. Padhye, S. Agarwal, V. Padmanabhan, L. Qiu, A. Rao, and B. Zill. Estimation of Link Interference in Static Multi-hop Wireless Networks. In *In Proc. of Internet Measurement Conference (IMC)*, 2005.
- [26] T. S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall, 1996.
- [27] K. Sundaresan and K. Papagiannaki. The Need for Cross-layer Information in Access Point Selection Algorithms. In *Proc of Internet Measurement Conference (IMC)*, pages 257–262, 2006.
- [28] Y. Yuan, P. Bahl, R. Chandra, P. A. Chou, I. Farel, T. Moscibroda, S. Narlanka, and Y. Wu. KNOWS: Cognitive Networking Over White Spaces. In *Proc of IEEE DySpan*, 2007.
- [29] Y. Yuan, P. Bahl, R. Chandra, T. Moscibroda, and Y. Wu. Allocating Dynamic Time-Spectrum Blocks in Cognitive Radio Networks. In *Proc of ACM MOBIHOC*, 2007.