

Traffic-Aware Channel Assignment in Wireless LANs

Eric Rozner, Yogita Mehta, Aditya Akella, Lili Qiu

ABSTRACT

The importance of channel assignment in wireless LANs has been long recognized and has attracted significant research attention. Despite much recent progress, several important challenges still remain: First, what is the right “performance metric” to optimize for? Most existing approaches to channel assignment try to minimize the number of mutually-interfering APs. However it is not clear if this metric accurately reflects client performance. Second, to what extent is the quality of channel assignment improved by incorporating the observed traffic demands at APs and clients? Recent work in IP traffic engineering has shed light on the tremendous effectiveness of using traffic demands in network engineering decisions. However, to date, no approach to channel assignment has taken traffic demands into account.

In this paper, we conduct extensive simulations over publicly-available wireless traffic traces, to study the relative efficacy of different optimization metrics in tracking client performance. We show that being “traffic-aware” could substantially improve the overall quality of a channel assignment, irrespective of the metric employed. We also examine the effect of incomplete traffic information on the quality of a channel assignment. Finally, we develop and evaluate a preliminary set of practical traffic-aware assignment algorithms that predict future demands based on historical information and use the predicted demands for assigning channels.

1. INTRODUCTION

Over the past few years, wireless 802.11 technology has been widely adopted to provide connectivity in enterprise and university campus settings. Despite the several advantages that wireless technology offers—including low cost, ease of installation and support for mobile users—using it efficiently in these settings is surprisingly difficult. Indeed, due to the broadcast nature of wireless communication, wireless LANs are often plagued by interference issues that can cause sub-optimal user performance.

The traditional solution to this issue is to exploit the frequency diversity offered by wireless technology and configure neighboring networks to operate on distinct “non-overlapping” channels—an approach commonly referred to as “channel assignment”. 802.11b and g support three non-overlapping channels each, while the 802.11a provides 12 such channels.

The conventional, and widely employed, approach to channel assignment in WLANs is for network administrators to conduct a detailed “site-survey” and come up with a set of candidate AP locations to maximize coverage. Then, the admins hand-configure the frequencies of APs to minimize the number of mutually-interfering base stations [17]. Recently, more informed approaches to channel allocation have been proposed both in the research community (e.g.,

Lee *et al.* [21], Mishra *et al.* [24, 23]) as well as the industry (e.g., Propagate Network’s Autocell [2], and Alcatel’s Airview [1]). These seminal approaches have shown that careful management of the wireless spectrum could yield substantial improvements in WLAN user performance.

Despite this progress, the WLAN channel assignment problem is far from solved and several important challenges still remain. In this paper, we focus on two key issues:

1. *What “performance metric” should channel assignment optimize for?* Existing approaches to channel assignment aim to minimize the interference among neighboring APs [17, 23] or at client locations [24].¹ Ultimately, however, the quality of a channel assignment is determined by the network performance experienced by wireless users. It is not clear if any of the currently employed metrics accurately reflects the eventual performance of end-clients.
2. *Does the quality of a channel assignment improve when dynamic traffic demands in the WLAN are taken into account?* In current approaches, channel assignment is either performed only once, or repeated occasionally when new APs are added or the locations of clients change significantly. In particular, current approaches do not adapt to prevailing traffic conditions. A quick analysis of publicly available wireless traces (e.g. [12]) shows that traffic volumes in a wireless LAN could vary significantly over time and across access points. Recent studies [16] have also shown that traffic volumes on WLANs have grown substantially over the years. The application traffic mix is becoming increasingly diverse, making traffic bursty and unpredictable. An increasing number of embedded wireless devices are sharing the narrow frequency spectrum in WLANs. At the same time, the density of AP deployments is growing to ensure seamless connectivity for all devices. In the future, these trends are likely to become even more pronounced. Unfortunately, the limited number of channels (especially orthogonal channels) provided by popular 802.11 technologies (b, g) is insufficient to support these deployments. Therefore, we believe that it is important to develop new spectrum management schemes that better accommodate these trends.

To address the first challenge, in this paper, we perform an in-depth study of several optimization metrics—such as the number of interfering APs, noise at AP and/or client locations and expected client throughput. We use real traffic traces to evaluate their relative ability in accurately tracking the performance of wireless clients.

More importantly, to optimize spectrum usage in light of the changing trends in WLAN usage, we develop and systematically study the notion of *traffic-aware channel assignment algorithms* for WLANs. Our work is motivated by the recent string of studies in IP traffic engineering (see, for example, [4, 34, 32, 33]) that have established the tremendous potential of leveraging traffic demands in making network-

¹The interference at an AP location is defined as the sum of the strengths of the signals of all other in-range APs.

wide engineering decisions.

Our study employs trace-driven simulations conducted over large, publicly-available wireless traffic traces collected at Dartmouth College in 2004 [12]. Our analysis of this data shows that the traditional optimization metrics often result in sub-optimal channel allocation from the perspective of overall network throughput. We evaluate a collection of other metrics that better capture spectrum usage and client-perceived performance—namely, noise measured at AP locations, noise at client locations and aggregate noise in the WLAN. We show that using these metrics can improve the median network performance by up to 60%. Our results highlight the importance of using the right metric for optimization.

To evaluate the potential benefits of traffic-aware assignment, we simulate a setting where an “oracle” can accurately compute the right channel assignment to accommodate traffic variations in the immediate future. Our evaluation of the oracle shows that being traffic-aware could significantly improve the quality of a channel assignment (in terms of network throughput), irrespective of the metric being optimized for. We formulate a new metric, aggregate client throughput, and show that a traffic-aware channel allocation using this metric can improve the median network performance by a factor of 2 compared with conventional channel assignment approaches.

We also evaluate settings where only partial traffic information is available. Specifically, when only AP-side traffic information is available (and finer-grained client-side information is not), we observe that the network performance due to the traffic-aware channel assignments is $\sim 10\%$ inferior than the full information case. In some situations, the network performance is no better than a traffic-agnostic channel assignment. We also investigate if spectrum usage can be improved further by relying on mechanisms such as dynamic transmit power control [25]. Our results show that traffic-aware channel assignment is crucial even if other approaches for improving spectrum efficiency are already in use.

In practice, the “oracle” assumption is unrealistic: traffic-aware algorithms can only rely on historical traffic measurements to derive channel assignments. Yet, they must remain robust to wild fluctuations in traffic. Therefore, we propose initial approaches for traffic demand prediction, and we modify our channel assignment algorithms to use predicted demands. We evaluate the quality of the resulting channel assignments and show that the performance they offer is reasonably close to the oracle-based algorithms (within 5–10%).

The rest of the paper is organized as follows. In Section 2, we survey related work. In Section 3, we define several optimization metrics for traffic-agnostic channel allocation. We define oracle-backed approaches for traffic-aware channel assignment in Section 4. Section 5 provides details of our simulation set up and datasets. Our evaluation results are presented in Section 6. In Section 7, we present initial approaches for practical traffic-aware assignment. In Section 8, we discuss drawbacks of our approach and possible extensions. We conclude in Section 9.

2. RELATED WORK

We first review past work on the channel assignment problem. Then, we discuss related work in IP traffic engineering.

2.1 Channel Assignment

Channel assignment for improving the efficiency of spectrum usage is a well-studied problem. In particular, the problem has received much attention in the context of cellular networks [20]. In general, approaches for cellular channel allocation are unsuitable for our purposes (i.e., traffic-aware channel assignment in WLANs): Cells in a cellular network are arranged in a very regular fashion and have uniform, large coverage areas, unlike the regions covered by indoor access points (APs). As a result, channel assignment in cellular networks is a static, one-time task. In contrast, depending on the number and location of clients, load on APs and the presence or movement of obstacles, channel assignment across WLAN APs may need to change over time.

Popular approaches to channel assignment in 802.11 networks can be broadly classified into three categories based on the deployment scenario: (1) channel assignment and load balancing in enterprise/campus deployments; (2) channel assignment for mitigating interference in “chaotic” wireless deployments [5]; and, (3) channel allocation and management for capacity improvement in multi-hop mesh networks. While our focus in this paper is on #1, we also review related work in the other two areas for the sake of completeness.

Campus/Enterprise WLANs. The typical approach to assigning channels across APs in WLANs is the following [17]: First, network administrators conduct a detailed “RF site survey” of the campus. The goal is to determine the location and the number of access points required to adequately cover potential users. Next, the admin manually assigns available non-overlapping channels to individual access points (3 such channels exist in 802.11b and 12 in 802.11a). Once assigned, the channels remain unchanged for long periods of time.

More recently, several commercial products have been developed to automate channel assignment across WLANs. While some perform static channel assignment (see, for example DeviceScape [13]), other recent products perform automated, dynamic channel selection based on the current operating conditions (see, for example, AutoCell from Propagate Networks [2] and Alcatel OmniAccess AirView Software [1]). Of the latter set of tools, most also offer interference mitigation via transmit power control, along with load balancing across APs.

While these products are beginning to become popular, very little is known about their design, the system-wide metrics they optimize (i.e., do these tools minimize noise, or maximize per-flow throughput?), and their potential benefits. In our work, we study a variety of system-wide metrics, such as, the number of pairs of interfering APs, total system noise and aggregate throughput. We present an in-depth analysis of their effectiveness in tracking the performance of end-users.

Among research studies, there are two sets of approaches to channel assignment in WLANs: (1) AP-centric approaches, such as Lee *et al.* [21], identify “expected high-demand points” in a given WLAN deployment, and perform static channel assignment at APs to maximize signal strength at the demand points (effectively, extending the “site survey” approach above). (2) Client-centric approaches, such as Mishra *et al.* [24], argue that clients have a better view of interference (since interference directly impacts their performance), and therefore channel assignment must take client-side views of interference into account. In contrast with our

work, neither approach argues if optimizing for a particular metric (such as signal strength, or client-perceived interference) results in relatively better system-wide performance. Also, unlike past studies, our work shows that taking client and AP traffic patterns into account could result in better-quality channel allocation.

Chaotic networks. The term “chaotic networks” [5] characterizes modern home deployments of 802.11 access points which are often unplanned (since home users spontaneously deploy APs) and unmanaged (since users do not configure their APs carefully). Most home users manually configure the channel on their access points, and, more often than not leave the default configurations on their APs as is [5]. In general, efficient spectrum management in chaotic settings cannot be achieved by static, optimal channel assignment [5]. Other techniques, such as transmit power control must be performed in conjunction with channel assignment. Spectrum management in chaotic settings is further complicated by the lack of cooperation among APs.

We believe that our techniques for spectrum management based on client traffic patterns can be extended to chaotic settings as well (with APs making local decisions about which channel to operate on, based on locally observed traffic characteristics). We leave these extensions for future work.

Multihop mesh networks. Traditional approaches to wireless mesh networking [30, 26] advocate operating all wireless APs or antennas on the same channel. This is crucial to achieving successful transmission across multiple wireless hops. However, this significantly impacts the capacity of mesh network since transfers across a pair of wireless nodes can prevent a parallel transfer across other “in-range” pairs.

Two classes of solutions have been proposed to improve network capacity in such situations: The first class of solutions, proposed by Raniwala *et al.* [27, 28] advocates equipping mesh network nodes with multiple network interface cards (NICs) operating on different channels. The second, proposed by Bahl *et al.* [8], advocates a new link-layer mechanism called SSCH, wherein neighboring mesh nodes perform synchronized (or loosely synchronized) channel hops to better exploit frequency diversity. In both cases the ultimate goal of channel assignment is very different than in WLAN-settings: in WLANs, neighboring APs should be assigned to distinct channels where possible to mitigate interference; in mesh settings, neighboring nodes must be assigned similar channels, or overlapping hopping sequences, to ensure that data can be successfully transmitted.

2.2 Traffic Engineering in ISP Networks

Traffic demands have been shown to have tremendous utility for network provisioning and route optimization in ISP networks [6, 7, 31]. A wide range of traffic engineering approaches have been developed to incorporate traffic demands. At a high level, these approaches maintain a history of observed traffic demand matrices, and optimize routing for the representative traffic demands extracted from the observed traffic during a certain history window. They differ in how representative demands are derived.

For example, Agrawal *et al.* [4] use a traffic matrix in a one-hour window during daily peaks as the representative demand. Zhang *et al.* [32, 33] consider multiple representative traffic matrices and find an optimal set of routes to minimize expected or worst-case cost for these repre-

sentative matrices. TeXCP [19] and MATE [14] conduct online traffic engineering and react to instantaneous traffic demand.

Inspired by these results from the wireline world, we ask whether being traffic-aware has similar benefits for managing wireless network spectrum.

3. OPTIMIZATION METRICS

In this section, we outline five different metrics for optimization in channel assignment. Some of the metrics are generalizations of those employed in contemporary channel assignment schemes. In addition, we outline a few other metrics which, we believe, are better at capturing spectrum usage and, in turn, client performance.

In all cases below, we assume that the channel assignment does not depend on the volume of traffic sent or received by clients and access points. Where necessary, we assume that the locations of clients and APs are known. In other words, the approaches outlined next are static or infrequent—the assignment is either performed only once, or repeatedly occasionally (e.g. when a new AP is added or when client locations change significantly).

Also, while our focus is on allocating non-overlapping channels (e.g., channels 1, 6, 11 in 802.11b), we note that the metrics could be easily extended to allocating partially-overlapping channels as well (e.g., all 11 channels in 802.11b) [25].

1. Channel Separation. Here, the goal of channel assignment is to ensure that APs that are within interference range of each other operate on distinct channels whenever possible. This simple objective reflects the traditional goal of channel assignment in wireless LANs given the placement of APs [17].

Let C_i denote the channel assigned to AP i , $d(i, j)$ denote the distance between i and j , I denote the interference range, and A denote the set of all APs. Also, if $d(i, j) < I$, define $Separation(i, j) = \min(|C_i - C_j|, 5)$, otherwise $Separation(i, j) = 5$. The channel separation objective can then be expressed as:

$$Maximize : \sum_{i, j \in A} Separation(i, j)$$

Based on the pathloss model from [29], if P is the strength in dBm of a transmitted signal, then the strength of the signal at a recipient located dm away is $P - (40 + 3.5 \cdot 10 \cdot \log(d.))$.² In this model, the interference range, I , is defined by the equation $P - (40 + 3.5 \cdot 10 \cdot \log(I)) = T$, where T is the carrier sense threshold in dBm.

2. Noise at Access Points. In this case, the goal is to find a channel assignment that yields the lowest aggregate interference or noise (we use the two terms interchangeably) at AP locations. This metric captures the spectrum usage in a WLAN more directly than channel separation.

To compute noise levels, we assume that APs and clients are always active. The noise at an AP i , denoted by N_i , is the sum of the strengths of signals from all in-range APs and clients (excluding clients associated with i) operating on the same channel as i . In other words, for $i \in A$:

$$N_i = \sum_{\substack{j \in AUW, C_i = C_j, \\ AP(j) \neq i, d(i, j) < I}} P - (40 + 3.5 \cdot 10 \cdot \log(d(i, j)))$$

²We use constants that correspond to measurements reported in [18].

Here, W is the set of all clients. If j is a client, then $AP(j)$ denotes its access point. If j is an AP, then $AP(j) = j$. With this definition of noise, the objective of channel assignment can be expressed as:

$$\text{Minimize} : \sum_{i \in A} N_i$$

3. Noise at Clients. In contrast with the above, the goal here is to capture spectrum usage from the perspective of clients. The goal of this metric is to minimize the noise observed at client locations. We assume that APs and clients always have data to transmit. The noise at a client i is the sum of the strengths of signals from all other in-range APs and clients (including those associated with $AP(i)$) operating on the same channel. In other words, for $i \in W$:

$$N_i = \sum_{\substack{j \in AUW, C_i = C_j, \\ j \neq AP(i), d(i,j) < I}} P - (40 + 3.5 \cdot 10 \cdot \log(d(i,j)))$$

The objective then is— $\text{Minimize} : \sum_{i \in W} N_i$.

4. Aggregate Noise. This objective goes a step further and combines the channel assignment goals of the previous two objective functions. The goal is to minimize the aggregate noise level in a WLAN:

$$\text{Minimize} : \sum_{i \in AUW} N_i$$

5. Client Throughput. While the above metrics try to characterize spectrum usage in different ways, they do not directly account for the network performance of clients in the wireless LAN. In contrast, the final metric we study incorporates client-perceived performance into the optimization objective.

To define this metric, we make a few simplifying assumptions about traffic requirements of wireless nodes: (1) All APs wish to transmit the same aggregate volume of traffic to their clients; (2) Clients of a given AP wish to transmit the same volume of traffic to the AP; (3) The outgoing and incoming traffic volumes at an AP are identical.

In this setting, letting T_i denote the throughput of client i 's transmissions to its AP, the objective simply is:

$$\text{Maximize} : \sum_{i \in W} T_i$$

The per-client throughput can be computed using recently developed flow throughput and starvation models, such as that of Garetto *et al.* [15]. Instead, as we describe in Sections 4.2 and 5, we rely on simulations in NS-2 [3].

4. TRAFFIC AWARENESS

The above approaches to channel assignment are *traffic-agnostic*—that is, they do not take the actual traffic volumes of individual clients and APs into account. As we show in Section 5, traffic volumes vary substantially across APs. Traffic at a single AP could also vary significantly over short and long time-scales. Next, we outline *traffic-aware* channel assignment mechanisms that take dynamic variations in traffic into account.

4.1 Traffic Demands

Before outlining the algorithms, we define the term “traffic demand”. We also outline approaches for collecting traffic demand information in practice:

AP traffic demand: At a given time, the traffic demand of an AP is the aggregate volume of traffic per unit time that the AP wishes to send to its clients.

Client traffic demand: At a given time, the traffic demand of a client is the volume of traffic per unit time that the client wishes to send to its AP.

In practice, AP and client traffic demands can be gathered using SNMP [11]. Most commercial access points export an SNMP management interface that can be polled at 5 minute intervals to obtain: (1) total bytes sent by the AP (*IfOutOct*); (2) bytes received at the AP (*IfInOct*); and, (3) the number of active clients currently associated with the AP (*NumClients*).

The first statistic provides a direct estimate of AP traffic demand at 5-minute time granularity:

$$AP_Demand[t - 5, t] = \frac{IfOutOct(t) - IfOutOct(t - 5)}{\Delta(t)}$$

Similarly, the last two can be used to approximate client traffic demands:

$$Client_Demand[t - 5, t] = \frac{IfInOct(t) - IfInOct(t - 5)}{\Delta(t) \cdot NumClients(t)}$$

This assumes uniform demands across all clients of an AP. We note that finer grained per-client information can be obtained by collecting and correlating `syslog` and `tcpdump` statistics (this approach was employed in [22]).

4.2 Traffic-Aware Approaches

We now describe traffic-aware channel assignment algorithms. We assume that our algorithms have access to an *oracle*—i.e., they possess perfect information about instantaneous traffic demands in a WLAN, and how the demands vary in the immediate future. Therefore, our oracle-backed algorithms can dynamically alter channel assignments over time to accommodate future spikes or dips in traffic.

In practice, traffic-aware algorithms can only use historical traffic information in assigning channels; yet, the channel assignments must be robust to wild, unexpected fluctuations in demands. In Section 7, we relinquish the notion of an oracle and present a preliminary evaluation of practical traffic-aware approaches.

Channel Separation. To incorporate traffic-awareness, we modify the channel separation metric described in Section 3 so that interfering APs with high individual demands are first assigned to non-overlapping channels. In other words, we “scale up” the channel separation metric by the demands of individual APs. Symbolically, letting D_i denote the current demand at AP i , the goal of the channel assignment can be expressed as:

$$\text{Maximize} : \sum_{i \in A} D_i \sum_{j \in A, j \neq i} Separation(i, j)$$

Noise-based metrics. Similar to the above, we modify the definition of noise at an access point i (client i) to incorporate the traffic demand at i ($AP(i)$). Furthermore, the

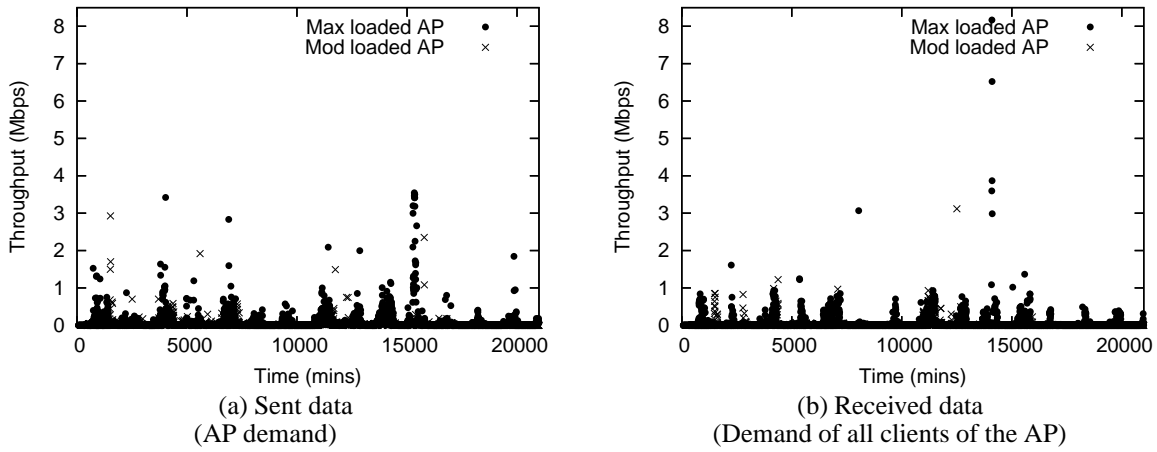


Figure 1: Time series of traffic for a heavily loaded and a moderately loaded AP from LibBldg in the Dartmouth Dataset.

noise calculation will only consider currently active clients and APs operating on a similar channel. In other words, for an access point i :

$$N_i = \sum_{\substack{j \in AUW, C_i = C_j, \\ IsActive(j), \\ AP(j) \neq i, d(i,j) < I}} D_i \cdot (P - 40 - 3.5 \cdot 10 \cdot \log(d(i,j)))$$

For a client i :

$$N_i = \sum_{\substack{j \in AUW, C_i = C_j, \\ IsActive(j), \\ j \neq AP(i), d(i,j) < I}} D_{AP(i)} \cdot (P - 40 - 3.5 \cdot 10 \cdot \log(d(i,j)))$$

Throughput. To make the throughput-based metric traffic-aware, we rely on simulations in NS-2: given current traffic demands at access points and clients, we instrument NS-2 to find a channel assignment that maximizes the total client throughput in simulation.

Rather than rely on a model to estimate client throughput, we employ direct measurements from simulations for the following reasons: (1) It’s hard to model the performance of wireless networks accurately; (2) Packet-level simulators are known to yield reasonably accurate performance results. Moreover, our focus is on the relative performance of different channel assignments, not their absolute values; (3) Finally, we note that our algorithms and approach are broadly applicable to other modeling schemes or packet-level simulators.

5. SIMULATION APPROACH

In this section, we first describe the datasets we employ in our study. Then, we outline the simulation set-up used in the rest of the paper.

5.1 Datasets

We conduct trace-driven simulations over publicly available data sets collected at Dartmouth College [12, 16].

While the Dartmouth College traces covered several campus buildings, the analysis in this paper focuses on three

specific buildings: “AcadBldg10”, “ResBldg94” and “LibBldg2”. These buildings contain 21, 12 and 20 access points respectively. Other buildings of similar type (e.g. other AcadBldg’s) had fewer access points.

The Dartmouth traces include SNMP statistics and number of active clients sampled every 5 minutes at all APs. As described in Section 4.1, we use this data to derive the aggregate traffic at an AP—in Mbps, both sent and received—for each 5 minute interval. We analyzed two weeks’ worth of SNMP data, collected between Feb 1st and Feb 15th, 2004.

To illustrate the time-varying nature of the traffic, in Figures 1(a) and (b), we show a time-series of the traffic sent and received, respectively, at a heavily-loaded and a moderately-loaded access point in LibBldg. We note the general trend of long, quiescent (or “low traffic”) periods followed by relatively shorter bursts of activity. The magnitude of the traffic bursts also vary significantly over time and across the two APs.

5.2 Simulation Set-up

We use the publicly available version 2.29 of NS with support for multiple non-overlapping channels.

We simulate the wireless LANs in three buildings: AcadBldg10, LibBldg2 and ResBldg94. The coordinates for the APs in the Dartmouth dataset are publicly available; we employ these coordinates directly in our simulation.

We distribute clients of an AP randomly within a distance of 30m centered around the AP. We use unidirectional constant-rate UDP sources located at APs and clients to model AP and client traffic demands, respectively. We assume that clients are not mobile.

Our simulations progress in rounds, where a single round covers a given SNMP measurement interval. Within a round, we use Simulated Annealing for searching the space of plausible channel allocations and obtaining the best assignment assignment for a certain metric. Of all the search algorithms we tried (e.g. we also tested random walk and greedy search), we found that simulated annealing offered faster convergence and greater accuracy. We use 30 iterations for simulated annealing. We stress that it is likely that the globally optimal channel allocation may not be found within these iterations.

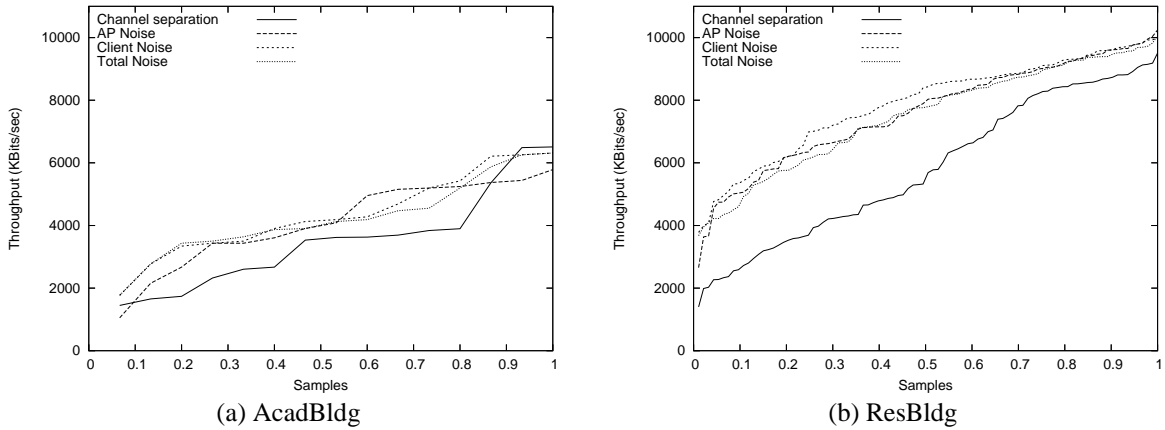


Figure 2: Effectiveness of different metrics for traffic-agnostic assignment.

We note that traffic-aware channel assignment is most effective when the network utilization is high and a significant number of access points are active. To study the importance of traffic-awareness, in our simulations, we focus on intervals with $\geq 50\%$ simultaneously active APs. We assume that an AP is active if the total volume of traffic it sends and receives exceeds 10Kbps. Also, we scale up the traffic demands in these intervals to increase the net utilization.

6. EVALUATION RESULTS

In this section we evaluate both traffic-agnostic and traffic-aware channel assignment algorithms. For a given time interval, we quantify the “effectiveness” of a channel assignment by computing the net throughput achieved by all network flows under the assignment. Our results cover 1 of the 14 days of traffic traces, namely Feb 1, 2004 (Sunday).³ We note that the traffic volumes were generally low on this date, especially for AcadBldg and LibBldg. Also, compared to ResBldg, we observed far fewer intervals where $\geq 50\%$ APs were simultaneously active.

6.1 Comparison of Optimization Metrics

First, we focus on traffic-agnostic allocation algorithms. Our aim is to understand the relative effectiveness of different optimization metrics in capturing spectrum usage and client performance.

Figures 2(a) and (b) provide a comparison of four of the metrics—channel separation, AP noise, client noise and aggregate noise—for WLANs in AcadBldg and ResBldg. Each figure plots a CDF of the net throughput performance of flows in the WLAN, when a certain channel assignment approach is employed. The x-axis is the fraction of intervals with $\geq 50\%$ simultaneously active APs.

First, we notice that the simplest and conventional metric, namely channel separation, provides the worst overall performance. The sub-optimality of this metric is particularly pronounced for ResBldg (Figure 2(b)). We note that the median throughput performance offered by this metric is about 5Mbps. Also, the net throughput is below 6Mbps roughly 55% of the time. In contrast, for the noise-based metrics, the median performance is in excess of 8Mbps,

³Unfortunately, our simulations for the rest of the data did not finish on time.

and the throughput is below 6Mbps for only 20–25% of samples.

Although the differences between the metrics are less pronounced for AcadBldg⁴ Figure 2(a)), we note a 2X difference in the 20-th percentile throughput, and a 10% difference in the median throughput, between the noise-based metrics and channel separation.

The noise-based metrics are able to better capture spectrum usage, and therefore offer substantially better performance than channel separation. We do note that difference across the three noise metrics is negligible, indicating that they are equally effective at capturing spectrum usage and client performance in practical situations.

These results show that *in the traffic-agnostic case, the choice of the optimization metric is closely coupled with the quality of the channel assignment in the WLAN.*

6.2 Traffic-Awareness Helps

We now evaluate the oracle-based, traffic-aware channel assignment algorithms. We assume that both AP and client demands are known. Our aim is to understand to what extent traffic-awareness can improve performance in practice.

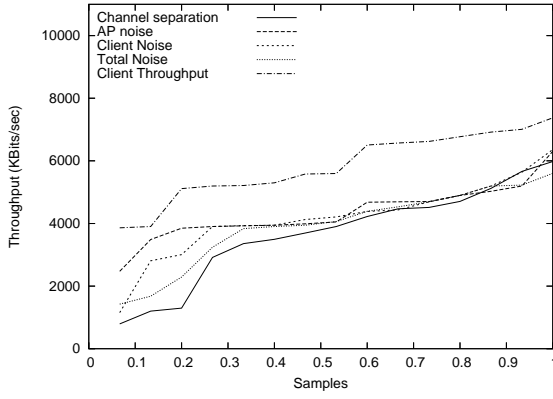
In Figures 3(a) and (b), we compare the effectiveness of different traffic-aware approaches for AcadBldg and ResBldg. Again, we plot a CDF of the net throughput of flows.

We contrast these figures against their traffic-agnostic counterparts in Figures 2(a) and (b), respectively. We note that incorporating traffic awareness definitely improves the overall performance of channel assignments, irrespective of the metric considered (i.e., all curves “shift upward”). However, the extents of improvements vary across the metrics.

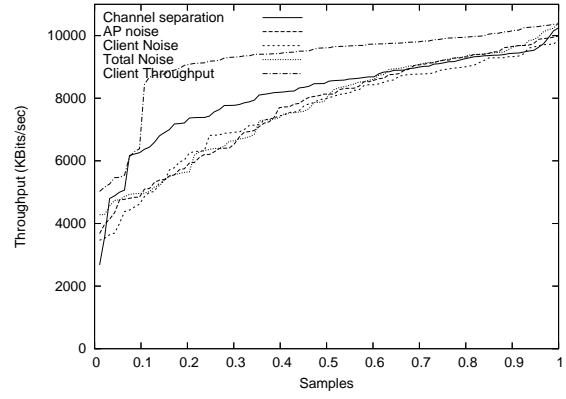
Among the four metrics described in the previous section, we note that channel separation gains substantially from incorporating traffic awareness: Focusing on ResBldg, for example, the network performance is under 8Mbps roughly 70% of the time for the traffic-agnostic setting (Figure 2(b)). For the traffic-aware channel separation (Figure 3(b)), this fraction falls to 35%.

In contrast, for the noise-based metrics, we observe very little improvement from incorporating traffic-awareness. Also, the improvements are even less pronounced for AcadBldg (Figure 3(a)).

⁴Again, this was mainly due to our simulations covering a single day (Sunday) of traces. We also explore other reasons in a later section.

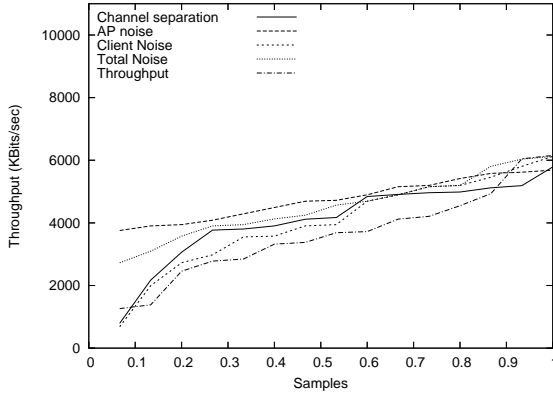


(a) AcadBldg

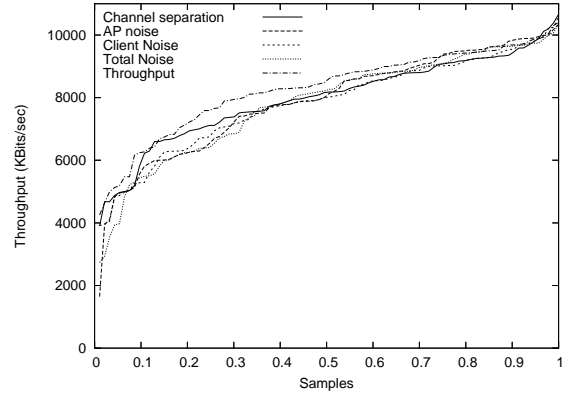


(b) ResBldg

Figure 3: Effectiveness of traffic aware assignment.



(a) AcadBldg



(b) ResBldg

Figure 4: Effectiveness of traffic aware assignment without client-side information.

Figures 3(a) and (b) also plot the network performance from traffic-aware client-throughput based metric. We note that the throughput metric offers substantial improvement in overall performance compared to other metrics. For example, the overall performance for ResBldg is under 9Mbps for only 20% of the intervals, compared with 80% of the intervals for the other metrics. The median throughput performance from the throughput metric is 15% higher compared to other traffic-aware metrics.

Also, compared to the conventional method of channel assignment—namely, static assignment using the channel separation metric—traffic-aware assignment based on the throughput metric nearly doubles the median network performance (contrast the throughput curve in Figure 3(b) with the channel separation curve in Figure 2(b)). For AcadBldg (Figure 3(a)), we note that the traffic-aware throughput metric improves the median performance by 35% compared to the median throughput from other traffic-aware schemes.

The fact that ResBldg often experiences better improvement than AcadBldg (and LibBldg, although not shown) requires further analysis. We highlight two main differences between the AP topologies in order to help explain the variance in improvement gains:

1. The density of APs deployed in ResBldg is about 3.65

times higher than that of AcadBldg.

2. The amount of traffic generated in ResBldg is higher than AcadBldg; And, the number of intervals with $\geq 50\%$ active APs is four times higher in ResBldg.

As access point density and traffic load increase, the importance of selecting a good channel assignment also increases. Topologies with low AP density or minimal traffic demands are more likely to perform well over a broad-range of channel assignments. When demand and density increase, efficient spectrum management becomes vital to obtaining good performance.

The observations in this section bring two conclusions to the fore: (1) *Traffic-aware channel assignment offers substantially better network performance than traffic-agnostic assignment.* The importance of traffic-awareness is more pronounced in dense deployment settings; and (2) *The client-throughput based metric is the most effective at tracking the quality of channel assignment as perceived by clients.*

6.3 Partial Demand Information

We now shift focus to situations where only partial demand information is available. Specifically, we assume that only AP-side traffic demands are known. When assigning channels, we assume that client-side demands are all zero.

We evaluate the performance of the resulting assignment against the actual client-side demands.

In Figures 4(a) and (b), we compare the effectiveness of different traffic-aware approaches that do not use client-side information, for AcadBldg and ResBldg, respectively.

Comparing Figures 4(a) and (b) with the full-information results in Figures 3(a) and (b), we observe that not having complete information has a telling impact on the quality of a channel assignment. For example, focusing on ResBldg, the median throughput performance under the throughput metric drops from 9.5Mbps in the full-knowledge case to about 8.5Mbps in the partial-knowledge case. The performance of the channel separation metric is similarly inferior in the partial-knowledge scenario.

For AcadBldg in Figure 4(a), we note that the performance under throughput-based metric is severely impacted (compare Figure 4(a) and Figure 3(a)), and it is no better than the other optimization metrics. In fact, the performance is comparable to channel separation in the traffic-agnostic case (Figure 2(a)).

Metric	W/ client - W/o client (Kbps)	
	AcadBldg	ResBldg
Channel Separation	698	1676
AP Noise	614	1440
Client Noise	2320	1447
Total Noise	68	1387
Client Throughput	4139	1994

Table 1: 90-th percentile of the difference in throughput performance with and without client-side demands.

In Table 1, we show the 90-th percentile of difference in the throughputs between an assignment that uses client side information and one that does not. Of the five metrics, we note that the performance of the client-throughput based metric is affected the most—by as much as 4Mbps—due to the lack of extra information.

Together, these results point to an important fact: *the more the amount of information available about traffic conditions, the better the quality of the channel assignment.*

6.4 Transmit Power Control

Our simulations so far assumed a static transmission power setting on all access points. In this section, we explore the effectiveness of dynamic transmit power control. At the beginning of each simulated 5 minute interval, we re-assign power levels on APs so as to just be sufficient to cover the clients associated with them during the interval. Moreover, we restrict the power settings to belong to the set {0dBm, 7dBm, 13dBm, and 15dBm}. This reflects power control on modern commercial APs, which disallow arbitrary power settings.

We contrast three approaches: (1) Employing dynamic power control with a random assignment of underlying channels (we assume that client location is known beforehand); (2) Employing traffic-aware channel assignment (using the client-throughput metric); and (2) Employing dynamic power

control in conjunction with traffic-aware channel assignment.

Algorithm	Average Throughput (Kbps)		
	AcadBldg	LibBldg	ResBldg
Tx Power	4548	8727	6382
Channel Assignment	6941	9432	7622
Channel Assignment + Tx Power	6831	9450	7741

Table 2: Throughput performance of transmit power control with and without channel assignment. We also show performance from using traffic-aware channel assignment alone.

In Table 2, we show the average throughput of the three approaches. The average is computed over 200 intervals belonging to the first day of trace collection (as mentioned earlier, we focus on intervals with at least 50% active APs). We note that using transmit power control in isolation is not sufficient to extract the best possible performance from the WLAN. In fact, reducing transmit power could hurt the overall performance, since it degrades signal strength at clients and exacerbates the impact of interference.

We also note that traffic-aware channel assignment alone is able to significantly improve performance. Using the two approaches in combination provides little added benefit.

These results show that *traffic-aware channel assignment is necessary even if other approaches for improving spectrum efficiency are already in use.*

7. TRAFFIC-AWARENESS IN PRACTICE

The algorithms described in Sections 3 and 6 rely on an oracle that provides accurate information about upcoming traffic demands. In practice, a traffic-aware channel assignment algorithm can only employ historical demand information to obtain channel assignments. This gives rise to two practical issues: (1) How to use historical data to identify trends in demands and to predict future demands reasonably accurately? Here, channel assignment is based on predicted demands; (2) How to ensure that the resulting assignment is robust to mis-predictions and to wild fluctuations in demands?

7.1 Algorithms

We present a family of practical traffic-aware algorithms for channel assignment. These algorithms offer varying degrees of trade-offs between the two issues discussed above. We present a preliminary evaluation of some of these algorithms in Section 7.2

Exponentially-Weighted Average of Demand (EWMA).

This approach predicts AP demands at time t by using a simple weighted moving average of demands observed in previous intervals. More recent demands are given greater weight than older ones:

$$APDem_Pred(t) = w \cdot APDem_Actual(t-1) + (1-w) \cdot APDem_Pred(t-1)$$

We set the weight $w = 0.9$. We estimate the aggregate amount of data received at the AP ($APRecd_Pred$)

and the number of active clients ($NumClients_Pred$) in a similar manner. Then, the predicted demand of a client is $ClientDem_Pred(t) = APRecd_Pred(t)/NumClients_Pred(t)$.

We use the predicted AP and client demands at time t to obtain a channel assignment.

Optimal for the Previous Interval (PREV). Here, the channel assignment for time t is simply the optimal channel assignment for the traffic demands in time $t - 1$ (or the most recently sampled time interval, if there are no samples available for $t - 1$). In other words, PREV is simply EWMA with $w = 1$. We note that compared to EWMA, PREV is more sensitive to short term traffic fluctuations.

Optimal Over a Time Window (PREV_N). The PREV approach could prove ineffective even under common traffic patterns, e.g., a periodic bursty traffic setting. Our next approach, PREV_N, tries to address this drawback by simultaneously optimizing for all traffic demands observed over a sizeable history window. In other words, PREV_N will derive a channel assignment for time t that maximizes the *total* client throughput for traffic demands from the past N intervals.

Symbolically, the goal of the channel assignment is:

$$\text{Maximize : } \sum_{i=1..N} \text{Throughput}(\text{Demands}(t - N))$$

Where $\text{Throughput}(\text{Demands}(t))$ denotes the total client throughput under the traffic demands from interval t .

Peak Demand in a Time Window (PEAK_N). This approach is a variant of PREV_N: Instead of optimizing for *all* sets of demands in a time window, PEAK_N obtains a channel assignment that is optimal for the *worst-case* demand-set within the history window. This allows the channel assignment to be more responsive to sudden increases in aggregate network utilization.

The above channel assignments can be made robust to mis-predictions and to fluctuations in traffic by having each approach optimize the channel assignment for a random sample of M demand-sets *in addition* to the predicted demands. Note that these demands-sets are generated randomly, and are not sampled from previously observed demands.

7.2 Evaluation of Practical Algorithms

We now show the network performance from employing the predicted demand-based algorithms described above. Specifically, we compute the network throughput averaged over a few hours worth of traces from Feb 1, 2004. We focus on the basic set of algorithms that do *not* employ additional random demands.

The table outlines the performance of EWMA, PREV and PEAK_N for $N = 2, 4, 8$, when using the client-throughput metric for optimization. The performance of the oracle-backed algorithm is also shown. While we do not evaluate PREV_N, we expect the results to be roughly identical to PEAK_N.

We note that the EWMA algorithm performs relatively poorly compared to the others. At the other extreme, PEAK_8 offers the best average performance, and is only 6% away from the performance of the oracle.

While these results are preliminary, they show that *traffic-aware channel assignment based on predicted traffic demands is effective and promising*. We leave a more thorough evaluation of the algorithms for future work.

Algorithm	Average Throughput (Kbps)	
	LibBldg	ResBldg
Oracle	10426	12126
EWMA	9309 (10.7%)	10959 (9.6%)
PREV	9459 (9.3%)	11207 (7.6%)
PEAK_2	9617 (7.8%)	10489 (13.5%)
PEAK_4	9812 (5.9%)	10717 (11.6%)
PEAK_8	9752 (6.5%)	11373 (6.2%)

Table 3: Comparison of practical traffic-aware algorithms optimizing for client throughput. The relative difference with the oracle-based algorithm is shown in parentheses.

8. EPILOGUE

This paper extends the current body of work on channel assignment in two new directions: First, we point to the importance of selecting a metric that accurately reflects spectrum usage and client performance. Second, and more importantly, we highlight the value of employing dynamic traffic information in assigning channels. Our results show that complete information is essential to obtain the best quality channel assignments.

Below, we address several issues pertaining to our traffic-aware approaches, including their effectiveness in a practical deployment, infrastructure support, and impact on client behavior.

Traffic-Aware Assignment vs Traffic Volumes. We note that traffic-aware channel assignment is interesting and effective in situations where the WLAN utilization is high, and a sizeable number of active nodes (APs or clients) are deployed close to each other. As recent studies have shown [16], today's campus WLANs are seeing substantially higher wireless traffic volumes. A barrage of new applications, including peer-to-peer file sharing, streaming video/audio and even voice-over-IP, are now frequently run over the wireless network. This has significantly altered the application mix, the burstiness and the predictability of traffic. Also, ever-increasing numbers of embedded wireless nodes are now sharing WLANs with the more conventional user laptops. As these trends continue into the future, we expect traffic-aware channel assignment to become increasingly crucial.

Infrastructure Support. To effectively incorporate traffic aware channel assignments, WLANs must deploy additional infrastructure to collect demand information as well as to disseminate channel assignment decisions to APs in a timely manner. Common management tools, such as SNMP, coupled with recent infrastructure proposals for WLAN monitoring and management [9] could be employed for this purpose.

In addition to tracking traffic demands, the monitoring infrastructure may need to collect a variety of other statistics. This includes client location and mobility patterns [10], path loss to clients, packet loss and noise levels at clients and APs as a function of location etc. We envision that the information will be centrally processed to perform a "what-if" analysis of channel assignment, similar to the

simulation-based approach studied in this paper.

We leave the actual implementation and deployment of a network monitoring and channel assignment infrastructure for future work.

Client-side Behavior. An important issue left unaddressed by our work is how clients respond to changes in channel assignment. Whenever an AP changes its channel, its clients will have to re-associate on the new channel. However, the 802.11 standard does not precisely define the re-association policy for clients.

One approach is for the wireless client to probe for APs using probe request packets. The APs can respond using a probe response packet (this is similar to AP's beacon packet). Alternatively, wireless stations can simply listen passively for beacons, which are transmitted every 100ms. The client associates with the AP and channel offering the highest RSSI. We are yet to explore the impact of re-association on ongoing client transfers. We do expect that reducing the beacon interval size on APs (e.g., to 50ms) is a simple way to contain the impact on client performance, if any.

802.11a. Our analysis has focused on 802.11b and g networks which support fewer operating frequencies than technologies like 802.11a. It is conceivable that traffic aware channel assignment is less critical in 802.11a networks. However, as WLAN deployment densities grow, and as multiple independently-administered WLANs operate in close proximity of each other, we believe that static allocation of non-overlapping channels—no matter how many—is unlikely to offer good performance.

9. SUMMARY

The centrality of channel assignment to improving the efficiency of spectrum usage in WLANs has been long recognized and well-studied. Several proposals have been made over the years, but they have all focused predominantly on static channel assignments that optimize simple system-level metrics (e.g., minimizing the number of interfering APs).

In this paper, we extend the body of work on channel assignment in two important and promising directions: First, we show that, in practice, the choice of the optimization metric is critical to obtaining channel assignments that ensure good client-perceived performance. We outline and evaluate several metrics each capturing spectrum usage and network performance to varying extents. We show that carefully crafted metrics could substantially improve the performance of WLAN clients.

Second, and more importantly, we explore the affect of dynamically adapting the channel assignment to prevailing traffic conditions. Using simulations based on large, publicly-available traces, we show that, in practice, traffic-aware channel assignment that use the right optimization metrics could potentially offer up to 2X better performance than conventional approaches. As AP deployments become denser and traffic volumes grow, traffic-aware channel assignment becomes increasingly important. We also show that it is important to collect complete information, covering the traffic requirements of both clients and access points; Otherwise, in some situations, the network performance may be no better than traffic-agnostic assignments.

Finally, we present initial results on practical traffic-aware assignment algorithms that use historical traffic information to derive good channel assignments. Our preliminary

results show that these approaches are effective in practice. We therefore believe that developing informed channel assignment algorithms that employ predicted traffic demands is a promising research direction to pursue.

10. REFERENCES

- [1] Alcatel AirView Software. <http://www.alcatel.com>.
- [2] Autocell. <http://www.propagatenetworks.com/product/>.
- [3] The network simulator - ns-2. <http://www.isi.edu/nsnam/ns/>.
- [4] S. Agarwal, A. Nucci, and S. Bhattacharyya. Measuring the shared fate of IGP engineering and interdomain traffic. In *Proceedings of the 13th International Conference on Network Protocols (ICNP) '05*, 2005.
- [5] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self Management in Chaotic Wireless Deployments. In *ACM MobiCom*, Cologne, Germany, Aug. 2005.
- [6] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja, and X. Xiao. *Overview and Principles of Internet Traffic Engineering*, RFC 3272, May 2002.
- [7] D. O. Awduche. MPLS and traffic engineering in IP networks. *IEEE Communication Magazine*, pages 42–47, December 1999.
- [8] P. Bahl, R. Chandra, and J. Dunagan. SSCH: Slotted Seeded Channel Hopping For Capacity Improvement in IEEE 802.11 Ad Hoc Wireless Networks. In *ACM MobiCom*, Philadelphia, PA, Aug. 2004.
- [9] P. Bahl, J. Padhye, L. Ravindranath, M. Singh, A. Wolman, and B. Zill. DAIR: A Framework for Managing Enterprise Wireless Networks Using Desktop Infrastructure. In *Workshop on Hot Topics in Networks (HotNets-IV)*, Nov 2005.
- [10] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM (2)*, pages 775–784, 2000.
- [11] J. Case, M. Fedor, M. Schoffstall, and J. Davin. A simple network management protocol (SNMP). In *Internet Engineering Task Force, RFC 1098*, May 1990. <http://www.snmp.com/>.
- [12] Dartmouth campus-wide wireless traces. <http://www.cs.dartmouth.edu/campus/>.
- [13] Devicescape. Channel Planning Device. <http://www.devicescape.com/docs/smap/AdminGuide/ChannelPlanning.php#wp1%000031493>.
- [14] A. Elwalid, C. Jin, S. Low, and I. Widjaja. MATE: MPLS adaptive traffic engineering. In *Proceedings of IEEE INFOCOM '01*, Anchorage, AK, April 2001.
- [15] M. Garetto, T. Salonidis, and E. Knightly. Modeling Per-flow Throughput And Capturing Starvation in CSMA Multi-hop Wireless Networks. In *Proc. of IEEE Infocom*, March 2006.
- [16] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proc. of ACM MOBICOM '2004*, Sept. 2004.
- [17] Jim Grier. Assigning 802.11b Access Point Channels. <http://www.wi-fiplanet.com/tutorials/article.php/972261>.
- [18] G. Judd and P. Steenkiste. Using Emulation to Understand and Improve Wireless Networks and Applications. In *Proceedings of NSDI 2005*, Boston,

- MA, May 2005.
- [19] S. Kandula, D. Katabi, B. Davie, and A. Charny. Walking the tightrope: Responsive yet stable traffic engineering. In *Proceedings of ACM SIGCOMM '05*, Philadelphia, PA, August 2005.
 - [20] I. Katzela and M. Naghshineh. Channel assignment schemes for cellular mobile telecommunications: A comprehensive survey. *IEEE Personal Communications*, pages 10–31, 1996.
 - [21] Y. Lee, K. Kim, and Y. Choi. Optimization of AP Placement and Channel Assignment in Wireless LANs. In *Workshop on Wireless Local Networks (WLN)*, *IEEE LCN*, Nov. 2002.
 - [22] X. Meng, S. Wong, Y. Yuan, and S. Lu. Characterizing flows in large wireless data networks. In *Proc. of ACM MOBICOM '2004*, Sept. 2004.
 - [23] A. Mishra, S. Banerjee, and W. Arbaugh. Weighted Coloring Based Channel Assignment for WLANs. In *Mobile Computing and Communications Review*, July 2005.
 - [24] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. Arbaugh. A Client-driven Approach for Channel Management in Wireless LANs. In *INFOCOM06*, Barcelona, Spain, Apr. 2006.
 - [25] A. Mishra, E. Rozner, S. Banerjee, and W. Arbaugh. Exploiting partially overlapped channels in wireless networks: Turning a peril into an advantage. In *Proc. of Internet Measurement Conference (IMC)*, 2005.
 - [26] Self-organizing neighborhood wireless mesh networks. <http://www.research.microsoft.com/mesh>.
 - [27] A. Raniwala and T. Chiueh. Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network. In *Proc. of IEEE Infocom*, March 2005.
 - [28] A. Raniwala, K. Gopalan, and T. Chiueh. Centralized algorithms for multi-channel wireless mesh networks. In *Proc. of ACM Mobile Computing and Communications Review (MC2R)*, April 2004.
 - [29] T. Rappaport. *Wireless Communications: Principles and Practice*. Prentice-Hall, Englewood Cliffs, NJ.
 - [30] MIT Roofnet. <http://www.pdos.lcs.mit.edu/roofnet/>.
 - [31] X. Xiao, A. Hannan, B. Bailey, and L. Ni. Traffic engineering with MPLS in the Internet. *IEEE Network Magazine*, pages 28–33, March 2000.
 - [32] C. Zhang, Z. Ge, J. Kurose, Y. Liu, and D. Towsley. Optimal routing with multiple traffic matrices: Tradeoff between average case and worst case performance. In *Proceedings of the 13th International Conference on Network Protocols (ICNP) '05*, Boston, MA, November 2005.
 - [33] C. Zhang, Y. Liu, W. Gong, J. Kurose, R. Moll, and D. Towsley. On optimal routing with multiple traffic matrices. In *Proceedings of IEEE INFOCOM '05*, Miami, FL, April 2005.
 - [34] Y. Zhang and Z. Ge. Finding critical traffic matrices. In *Proceedings of DSN '05*, Yokohama, Japan, June 2005.