

MobiCom Poster Abstract: Traffic-Aware Channel Assignment in Wireless LANs

Eric Rozner*
erozner@cs.utexas.edu

Yogita Mehta*
yamehta@cs.utexas.edu

Aditya Akella†
akella@cs.wisc.edu

Lili Qiu*
lili@cs.utexas.edu

*The University of Texas, Austin, TX, USA

†The University of Wisconsin, Madison, WI, USA

I. Introduction

Campus and enterprise wireless networks are increasingly characterized by ubiquitous coverage and high traffic demands. An efficient channel assignment of access points (APs) in these networks has the potential to greatly improve wireless LAN performance. Therefore, previous studies have investigated methods for determining effective channel assignments. Existing approaches, however, do not adapt to prevailing traffic conditions in the network. An analysis of wireless traces (e.g. [2]) shows that traffic volume in a wireless LAN can vary significantly across APs. Motivated by this observation, we investigate whether the quality of a channel assignment can be improved by incorporating observed traffic demands at APs and clients in wireless networks. Our results show that being traffic-aware can substantially improve the effectiveness of a channel assignment: in some cases it nearly doubles performance. We further apply traffic demand prediction algorithms to make traffic-aware assignment practical.

II. Optimization Metrics

In this section, we outline two metrics used for optimizing channel assignment that are generalizations of metrics employed in contemporary channel assignment schemes. In the next section, we augment each of these metrics to become traffic-aware. Throughout this paper, we focus on 802.11b/g networks, but these metrics can be easily extended to 802.11a.

1. Traffic-agnostic, client-agnostic channel separation. 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 metric is to *Maximize* : $\sum_{i,j \in A} Separation(i, j)$.

2. Traffic-agnostic, client-aware channel separation. The above metric only considers interference among APs. In real networks, minimizing interference introduced by client transmissions

is also important. Indeed, our analysis of real traffic traces shows that clients transmit a significant volume of traffic. Therefore, we extend the above client-agnostic channel separation to *Maximize* : $\sum_{i,j \in A \cup B, BSS(i) \neq BSS(j)} Separation(i, j)$, where B denotes the set of clients in the network. We assume that the client locations are known a-priori. In effect, this metric factors in the channel separation between any two interfering APs, any two interfering clients that are associated with different APs, and any interfering AP-client pair.

III. Traffic Awareness

To incorporate traffic demands, we modify the traffic-agnostic channel separation metrics so that interfering nodes with higher demands are more likely to be assigned to non-overlapping channels. Using this insight, we scale the channel separation between nodes C and D with the following *weight*:

$$W_{C,D} = S_C \times S_D + S_C \times R_D + S_D \times R_C$$

where S is the send demand, and R is the receive demand. Intuitively, if we abuse notation and let S_D denote the fraction of time D 's transmissions acquire the medium, the first term reflects the *probability* of C 's transmissions interfering with D 's transmissions. Similarly, the second term reflects the effect of C 's transmissions on D 's receptions, and the last term reflects the effect of D 's transmissions on C 's receptions.

3. Traffic-aware, client-agnostic channel separation. We augment the previous client-agnostic metric to: *Maximize* : $\sum_{i,j \in A} W_{i,j} \times Separation(i, j)$.

4. Traffic-aware, client-aware channel separation. Similarly, we obtain the following traffic-aware, client-aware metric: *Maximize* : $\sum_{i,j \in A \cup B, BSS(i) \neq BSS(j)} W_{i,j} \times Separation(i, j)$.

Optimizing channel assignment is a NP-hard problem. Therefore, we apply simulated annealing (SA) to search for a good channel assignment according to the above metrics. Since the performance of SA is sensitive to the initial solution, we develop an initialization

algorithm inspired by Chaitin’s approach to the register allocation problem [1]. We then apply simulated annealing to iteratively improve upon the initial assignment. Our evaluation uses 1000 iterations, which takes less than 1 second to compute. The output is the best assignment (among 1000 iterations) in terms of the given channel separation metric.

IV. Evaluation Results

In this section, we evaluate the effectiveness of traffic-aware channel assignment algorithms. We use both synthetic traffic and real traffic traces for our evaluation. First, we assume perfect knowledge of traffic demands. Later, we relax this assumption by developing traffic demand prediction algorithms. We use the total throughput over all network flows, as measured by ns-2, to quantify network performance.

First, we use synthetic traffic to understand when traffic-aware channel assignment is beneficial. Each topology in the evaluation has 50 APs and 200 clients, and a client has on average 4 APs in its communication range. We generate *hotspot traffic demands*. We randomly select an AP and all the other APs within its communication range as a hotspot. All APs inside the hotspot have traffic demands uniformly distributed between 0 and 3.6 Mbps, and all other APs have traffic demands uniformly distributed between 0 and 10 Kbps.

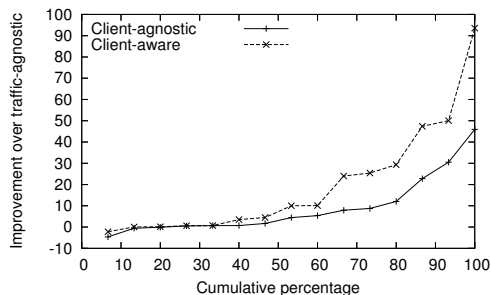


Figure 1: Percentage improvement of traffic-aware metrics in hotspot scenarios.

Figure 1 shows a CDF of the improvement of the traffic-aware, client-aware and traffic-aware, client-agnostic metrics over their traffic-agnostic counterparts. We see that being traffic-aware can improve performance by up to 93%.

Next we use trace-driven simulations to evaluate traffic-awareness under realistic traffic patterns. We use Dartmouth traces collected between Feb 1st and Feb 15th, 2004 to generate traffic demands. In the interest of brevity, we only report the performance results from the “ResBldg94” building, which contains 12 access points. Refer to [3] for other results. We assume that clients associated with an AP

are randomly distributed around the AP within a radius of 20m. We then compare the various channel assignment algorithms for every 5-minute interval. Throughout our simulations, we scale traffic demands upwards to study the effectiveness of different channel assignment schemes under high traffic load conditions.

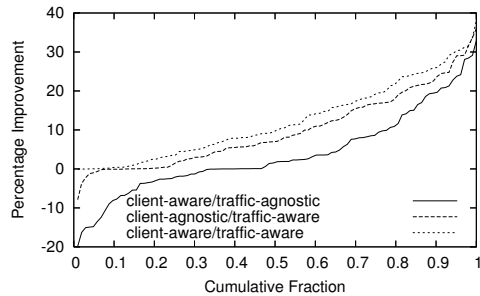


Figure 2: Improvement over the traffic-agnostic, client-agnostic metric in the Dartmouth traces.

Figure 2 shows a cumulative distribution function (CDF) of performance improvement of various channel assignments against the traffic-agnostic, client-agnostic baseline. We note that the average throughput improvement is 3.8% by incorporating client-side information alone; it raises to 9.9% by incorporating traffic-demands alone; and it becomes 12.3% by incorporating both traffic-demands and client-side information. We also note the amount of improvement is traffic-dependent. When traffic is more evenly distributed, we see little improvement from traffic-aware assignment. When traffic is more heterogeneous, the improvement is larger, as much as 40%.

Finally, we develop a series of traffic prediction algorithms to estimate future demands. One approach predicts demands at time t by using a simple weighted moving average of demands observed in previous intervals (EWMA). Another algorithm maximizes the total value of a metric for the traffic demands over the past N intervals (prev_N). While a more in-depth analysis is beyond the scope of this paper (see [3] for more details), we note that performance of the prediction algorithms is generally within 5% of the oracle. Therefore, we believe that traffic-aware channel assignment has significant potential in practical settings.

References

- [1] G. J. Chaitin. Register Allocation & Spilling via Graph Coloring. In *Proc. of SIGPLAN*, 1982.
- [2] Dartmouth campus-wide wireless traces. <http://www.cs.dartmouth.edu/campus/>.
- [3] E. Rozner, Y. Mehta, A. Akella, and L. Qiu. Traffic-Aware Channel Assignment in Wireless LANs. Technical Report TR-06-28, University of Texas at Austin, Computer Science, Austin, TX, May 2006.