Slotted Symmetric Power Control in Managed Wireless LANs

Vishnu Navda[‡][†] Ravi Kokku[‡] Samrat Ganguly[‡] Samir Das[†]
NEC Laboratories America, Inc.[‡] SUNY at Stony Brook[†]

Abstract

We present Contour, a novel dynamic power management technique for improving spatial reuse in managed wireless LAN (WLAN) deployments. WLAN deployments are often characterized by dense placement of access points (APs) to provide maximal coverage to clients. The capacity of such dense deployments is sub-optimal due to interference. In this paper, we focus on dynamic transmit power control for maximizing spatial reuse and improving network capacity. While dynamic power control of packet transmissions can mitigate interference, it can also introduce link asymmetry, thereby leading to unfair allocation of network resources. To avoid such asymmetry and maximize spatial reuse, we propose a slotted symmetric power control framework that forms the basis of Contour.

Through simulations on real topologies and a prototype implementation, we demonstrate that Contour significantly improves spatial reuse and fairness of resource allocation. For instance, in a real deployment of two APs with eight clients located randomly in different offices in the building, our experiments show a cumulative throughput improvement of 40%, and improved fairness.

1 Introduction

This paper explores the impact of dynamic power management for spatial reuse in *managed* wireless LAN deployments— the word *managed* refers to deployments in which all access points (APs) are under the same administrative domain. Such managed deployments are common today at multiple scales, from small-range corporate networks [21] and campus-wide networks [4], to hot-spot networks such as T-mobile [10], and city-wide muni-wifi networks such as Google's Wifi [5], MIT's roofnet [16], and Houston's urban network [19].

To provide maximal coverage, these wireless networks often contain multiple APs with overlapping transmission and interference ranges. Greater overlap leads to increased contention to transmit and increased packet collisions, thereby reducing the overall network throughput.

As WLANs become ubiquitous and users adopt the wireless medium as their first-class last-mile access network, network administrators are forced to increase the capacity of their wireless networks. Unfortunately, the capacity of an 802.11 network is limited due to the interference among network nodes and the availability of limited number of channels. Consequently, several research efforts are being focused on mitigating interference as much as possible to get maximum throughput out of these networks. Two major directions for capacity improvement have been toward exploiting channel diversity through efficient channel allocation [13] and spatial reuse through power control [12] or use of directional antennas.

In this paper, we propose fine-grained transmit power control for increasing spatial reuse by tuning the transmit power of each AP on a per-client basis. Our motivation is that clients closer to an AP can successfully receive packets even when the AP transmits at a lower power than that needed by a farther client. From a practical standpoint, such a proposition is made possible by the availability of devices with multiple power levels that can achieve per-packet power control with minimum-overhead.

Unfortunately, we observe that while such fine-grained control promises to substantially increase spatial reuse, it can lead to the formation of *asymmetric links* that can degrade the throughput of some clients. The result is an unfair allocation of network resources with the possibility of complete starvation of certain clients. For instance, consider two APs AP1 and AP2 each with one client C1 and C2 respectively. Let AP1 choose a smaller power to transmit to C1 than AP2, which is simultaneously transmitting to C2. While AP1 can carrier-sense AP2, AP2 might not be able to do so. As a result, the link from AP1 to C1 starves in the presence of AP2 to C2.

To this end, the goal of this paper is to devise a power control mechanism that is able to increase spatial reuse without sacrificing fairness; we propose a novel *slotted symmetric* power control framework. The main idea underlying this framework is to have all access points operate at the same power level at any given instant of time and have all access points follow a sequence of power levels synchronously. Time is divided into slots, and in each slot all access points operate at same power level, thereby avoiding any link asymmetry. Through hopping over different power levels across successive slots, all access points are able to serve the clients at their corresponding minimum power level requirements. Both simulation based evaluation as well as experimentation with prototype implementation shows that the slotted symmetric power control simultaneously improves the network throughput and provides better fairness to clients.

In summary, our key contributions are as follows:

- We develop a power control technique, Contour, based on the slotted symmetric power control framework. Contour is designed to adapt to traffic load and arrival/departure of clients. Contour's design ensures that there is no inter-AP messaging and that packet level scheduling is done at APs without involving any centralized entity.

- We compare Contour using ns simulations with (a) baseline CSMA/CA without power control, (b) static power control, and (c) unslotted dynamic per-packet power control. Our results show that removing link asymmetry in Contour leads substantially improved fairness compared to the above schemes. For instance, whereas schemes (b) and (c) above improve the throughput of 50-60% of the clients over CSMA and *hurt* the throughput of the rest of the 40-50% clients, Contour improves the throughput of 95% of the clients and hurts only 5%.

- We implement the proposed Contour dynamic power control technique in Click [27]. We evaluate the prototype for throughput improvement and fairness on a real testbed. Our evaluation shows that a capacity improvement of up to 60% is achievable by employing Contour. To the best of our knowledge, this is the first such implementation and evaluation of a fine-grained synchronous slotted framework in WLANs.

The rest of the paper is organized as follows. Section 2 presents background and motivation for dynamic power control. Section 3 discusses the challenges in realizing dynamic power control, and our conceptual approach to address the challenges. Section 4 discusses the design of Contour. Section 5 presents extensive simulation results to demonstrate the efficacy of Contour. Section 6 discusses important implementation and evaluation details. Section 7 identifies the outstanding issues and future directions of this work. Section 8 places Contour in the context of related work, and Section 9 concludes.

2 Background and Motivation

The performance of any 802.11-based wireless network is heavily dependent on the power at which senders transmit packets. A sender using higher transmission power can make greater number of neighboring senders defer their transmissions, thereby reducing the aggregate sending throughput of the network. Higher transmission power also results in higher probability of packet collisions at the receivers; for a given topology, fewer packets will be recovered by *capture effect*. In essence, if all access points (APs) operate at the highest available power level, the overall throughput of the network can be lower than what can be achieved. However, reducing the transmit power level of each AP in the interest of reducing interference at the neighboring APs can compromise coverage and degrade the delivery ratio achieved by some the clients. Here, the delivery ratio of a client is defined as the ratio of successfully received packets to the total number of packets sent in air. To this end, the objective of any power control mechanism is to maximize spatial reuse while meeting the target delivery ratio at the highest bitrate for each client.

The need for practical and efficient power control mechanisms is becoming increasingly critical in the future WLAN deployments for two reasons. First, the AP deployments are becoming more dense in order to handle the increasing client density. Higher density results in greater number of APs and clients getting affected by the transmission power of a given AP. For example, our analysis of the Google WiFi network [5] with clients placed at uniformly random locations shows the following: with an interference range of 400m, 85% of the clients are in the interfering range of four or more APs. With reduction in transmit power, and hence the interference range to 200m, the possibility of interference reduces to 15%. Second, the next generation of WLANs is targeted to provide a higher level of performance *predictability* as being defined by the IEEE 802.11T standard. The requirements suggest achieving uniform throughput and high delivery ratio for each client independent of the client's position and traffic profile.

Power control approaches can be categorized into (1) per-AP coarse granularity power control, and (2) per-client fine granularity power control. Most existing WLAN production networks (such as Cisco airespace [2]) provide per-AP power control. Per-AP power control mechanism uses the lowest power level for each AP that is needed to maintain a good delivery ratio for all the associated clients. The power level is adapted by observing each client's performance. Per-AP power control may not provide the best spatial reuse as the power level at which an AP operates is determined by the client with the worst link.

In per-client fine-grained power control, each packet transmission uses the lowest power level necessary to meet the delivery ratio of individual clients. Evidently, per-client power control results in higher spatial reuse, but requires control of power level on a per-packet basis. Such dynamic power control was infeasible in the past for two reasons: (1) the number of power levels at which an interface can transmit were limited, and (2) changing the power level incurred significant latency.

Recently, per-client power control has become feasible with the availability of wireless interfaces that provide greater number of power levels and per-packet fine-grained power control with minimal or no overhead. In order to demonstrate the viability of fine-grained power control, we conduct simple experiments. Figure 1 depicts the average signal to noise ratio (SNR) values of packets received by clients positioned at different distances (either in line-of-sight or not) from an AP that is transmitting packets at different power levels. The AP uses an Atheros [1] AR5212 chipset PCMCIA card for the top line and an Atheros miniPCI card with the same chipset for the bottom three lines; we use the MadWifi [8] device

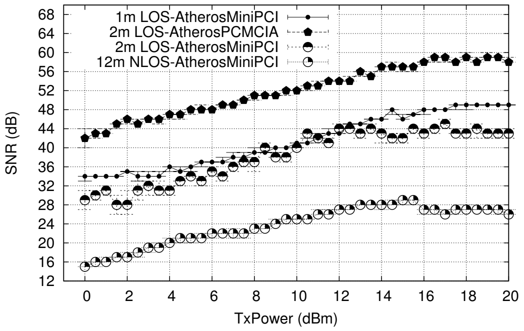

Figure 1: Average received packet SNR values at various distances from an AP transmitting at different power levels.

driver, version 0.9.2. To achieve transmit power control (TPC) on these cards, we recompiled the driver with TPC feature enabled.

The graph indicates that these cards achieve fine grained transmit power control; a 2dBm change in the AP's transmit power results in a perceptible change in the SNR value at the client. Further, our measurements show that each packet can be transmitted at a different power level without any impact on the sending rate (packets sent per unit time). These observations form the basis of our work that employs per-client fine grained power control.

3 Basic concepts

In this work, we take the approach of *dynamically controlling the transmit power of each AP on a per-client basis*. This approach is motivated by the observation that an AP can transmit a packet at a power level just sufficient for the client to receive it successfully; thus for close by clients the APs can use a lower transmit power, thereby increasing spatial reuse. APs can still send broadcast and management frames at the maximum allowed power level to ensure all clients within the coverage region receive it.

Although fine-grained power control improves spatial reuse, it also introduces link asymmetry. We discuss this problem next.

3.1 Link Asymmetry

If each AP independently chooses a power level, there will be frequent occurrence of situations where two or more neighboring APs are transmitting at different power levels. In such cases, *asymmetric links* are considered to be formed when among two APs using different power levels, the lower-power AP carrier-senses the higher power AP, but not vice versa. This situation leads to two problems : (1) unfairness in terms of sending throughput to the client, and (2) degradation in client delivery ratio.

To illustrate these two problems better, consider two APs (AP1 and AP2) having the same carrier sense threshold. Let AP1 and AP2 transmit at power levels P1 and P2 respectively with P1 $\geq$ P2. Under this setup, we discuss the above two problems.

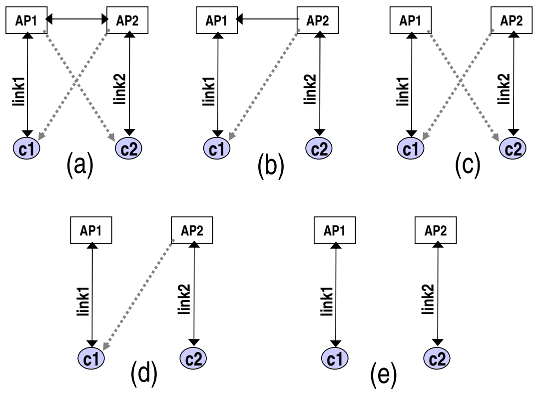

Figure 2: Scenarios of interaction between two AP→ client transmissions.

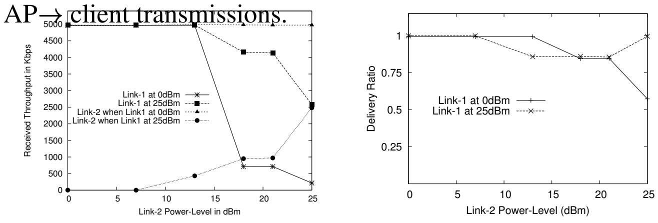

(a) (b)

Figure 3: Scenarios of interaction between two AP→ client transmissions.

Unfairness: With the above power levels, a situation can arise where AP2 can sense AP1 transmissions, but not vice versa. Every time AP2 senses AP1's transmission, it defers its own transmission; whereas AP1 never senses AP2's transmission and hence never defers its own transmissions. As a result, AP2 gets very few chances, if any, to transmit. Therefore, if both APs have data to send, AP2's sending throughput reduces drastically, thereby leading to starvation of the clients associated with AP2. Hence, link asymmetry may potentially lead to unfairness at AP-level in terms of sending throughput. Note that CSMA/CA already suffers from unfairness [35]; a naive dynamic power control scheme can worsen the fairness.

Delivery ratio degradation: Recall that delivery ratio is defined as the ratio of successfully received packets to the total number of packets sent in air. Packet collisions at the receiver can lead to lower delivery ratio. Again, with the above power levels, consider a scenario in which both AP1 and AP2 can not sense each other's transmission, but the client of AP2 can perceive packet transmissions from AP1. In this case, AP2 continues to transmit packets, while not being aware that they are being "destroyed" by transmissions from AP1. Since AP2 is using lower power level to transmit, the chance of recovering packets through capture effect at the client decreases further. This will degrade the delivery ratio of the client associated with AP2.

From both the above problems, we note the inherent tradeoffs in the objectives of a power control mechanism. While increasing spatial reuse can result in higher aggregate throughput, the throughput as perceived by clients is not uniform. In other words, increased spatial reuse may result in poor coverage (non-uniform client throughput).

Case studies: We illustrate the above two problems through experimental results. For ease of exposition, we refer to Figure 2 for identifying the cases. In Figure 2, solid dark arrow from node A to node B indicates that node A can carrier sense transmission from node B. In the same figure, the dotted arrow from AP to client shows interference that results in packet collision at client.

Figure 3(a) shows the sending throughput of two APs on link1 and link2 respectively. The sending throughput is determined by the number of packets transmitted in unit time (regardless of successful reception at client). On Y-axis, we vary the transmit power level for link2. If link1 operates at low power level (0dbm), the sending throughput is high when link2 is at 0dbm (Figure 2(e)) and falls drastically as link2 power level increases beyond 12 dbm (Figure 2(b)). Figure 3(a) further shows that when both links operate at same power level (0dbm and 25dbm), the throughput achieved by both links is equal. This corresponds to scenario shown in Figure 2(a) for high power level and scenario shown in Figure 2(e) for low power level.

Figure 3(b) shows the delivery ratios for two links. On Y-axis, we vary the transmit power level for link2. If the link1 operates at low power level (0dBm), we observe that as link2 power increases, the delivery ratio on link1 degrades from almost 1.0 to 0.6 (scenario shown in Figure 2(b)).

3.2 Slotted Symmetric Power Control

To overcome the above problems, it is required that all APs send at same power level. However, this does not achieve the full potential of per-client power control for maximum spatial reuse. In order to allow for both per-client power control and maintaining link symmetry, we propose a slotted symmetric power control technique—Contour. In this approach, each AP transitions through a sequence of power levels over time. However, at any instant of time, all APs in the network operate at same power level to avoid link asymmetry. Over time, by using different power levels, the system achieves per-client power control to maximize spatial reuse. We next discuss this approach in detail.

The basic slotted symmetric power control technique works as follows. Time is divided into slots, each slot is assigned a power level, and all APs use the same power level in a slot. At each slot boundary, all APs change the power level and stay in that allocated power level for the entire slot duration. This requires all APs to be synchronized in time. In a given slot, each AP decides the set of clients it should transmit to. Assigning clients to a slot is done based on the traffic load and the AP to client link characteristics. The sequence of power levels that each AP uses across successive slots defines the *envelope* that is repeated over time.

To illustrate the process, Figure 4(a) shows the mapping of clients of each AP to minimum power levels at which they can successfully receive packets. The AP associates c_1 at level P_3, c_2 and c_3 at level P_2, and c_3, c_4 and c_5 at level P_1. Figure 4(b) shows an example envelope with a sequence of three power levels that is followed by all APs in a synchronized manner. $\mathcal{T}$ is the envelope width or the periodicity of the power level sequence followed by each AP. In this example τ_1, τ_2 and τ_3 denote the slot width of power levels P1, P2 and P3. The amount of time τ_i the AP stays in power level P_i depends upon the total traffic load from clients mapped to the power level.

Properties: The slotted symmetric power control achieves better control of both power allocation and client traffic scheduling. With efficient refinement of the envelope, the system can adapt to changing workload, changing link characteristics, and new arrival and departure of clients. We discuss the adaptive strategies in the next section.

Having a synchronized slotted structure on top of CSMA/CA also enables us to leverage its benefits of providing good utilization in a distributed fashion. Since traffic fluctuates significantly in a WLAN, and often is unpredictable, a pure TDMA-like approach requires centralized packet scheduling for maximizing the utilization; such a per-packet scheduling solution is difficult to realize and does not scale well. In our proposed slotted structure, a centralized entity will only be required to create and refine the envelope at a coarse timescale. This entity is usually the WLAN controller that exists in most enterprise WLAN systems built on the thin-AP architecture model [3, 7].

Contour also has the property of alleviating hidden terminal effects. For example, suppose simultaneous transmissions from AP1 → C1 and AP2 → C2 are experiencing hidden-terminal problem and result in packet collisions at C2. In Contour, these simultaneous transmissions can be easily avoided by mapping C1 to the next higher power level. This is because transmissions at different power levels never happen at same time. Upgrading to a higher power level to address the hidden node problem happens implicitly in Contour as the mapping of client to power levels (that is dynamically adapted) is based on monitoring the delivery ratio of each client.

The above property also allows us to make the following corollary: Contour ensures that if a client C_i experiences the hidden node problem in steady state, the client C_i experiences the problem even in the network of APs that do not use Contour and all nodes transmit at their maximum power level. This is because, through upgrading the clients, the APs—in the worst case—operate at

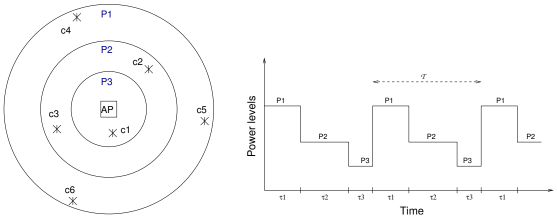

(a) (b)

Figure 4: (a) Mapping of clients to power levels for an AP (b) Envelope followed by each AP in cycles.

maximum power level, which is equivalent to having no power control.

4 Contour

Given a set of power levels $P_j, j \in [1...n]$, that each AP supports, the design of Contour is governed by our observation that an AP can transmit to each of its clients C_i at the lowest power level P_j that minimizes interference to transmissions of other APs, while not affecting the performance perceived by the client C_i. In what follows, we provide an overview of Contour's architecture and functionality, and then present the design of its individual components.

4.1 Overview

In designing Contour, we make the following assumptions.

Deployment Assumptions: First, each AP that uses Contour is connected to a central controller either directly through a wired or a wireless medium, or indirectly through other APs. Contour does not require any other connectivity between APs. Second, our client modifications are restricted to what can already be supported by current standards. For this reason, we do not advocate slotted symmetric power control at the client. Clients just transmit at the lowest power level sufficient for packets to reach the AP successfully. We assume adherence to IEEE 802.11h by the clients and APs for the AP to set the client power level.

Hardware and MAC Assumptions: First, Contour cannot modify the MAC layer for compatibility with commodity hardware. Second, each AP allows several power levels at which packets can be transmitted; the granularity of power levels is similar on all APs and is known to the central controller.

Functionality: We outline the steps taken by Contour to perform dynamic slotted symmetric power management:

- The controller and all APs are synchronized to a global real-time clock. In Section 6 we discuss a light-weight technique to synchronize the clocks of all APs to within 200μs using a commodity GPS receiver.
- An AP maps each of its clients to a minimum power level (in mW) at which the delivery ratio is greater than a threshold value (see Figure 4(a)).
- The controller determines an envelope of power levels (in mW) and the time τ_k to spend in each power level P_k and directs the APs to start following the envelope at a specific real-time. The envelope is a sequence of tuples of the form $[(P_1,\tau_1),(P_2,\tau_2),...,(P_n,\tau_n)]$, where $\sum_{k=1}^{n} \tau_k = \mathcal{T}$. $\mathcal{T}$ represents the period of the envelope. The times τ_k are initially set to a default allocation of $\mathcal{T}/n$.

  Each AP repeats the envelope starting at the specified real-time, as shown in Figure 4(b). Since APs are synchronized to the same global real-time clock, following the envelope ensures that all APs transmit at the same power at any instant of time till a new envelope is determined by the controller.
- Each AP determines refinements to the envelope to better adjust to its clients' requirements (including their position, traffic, etc). The AP sends refinement hints to the controller. The controller derives a new global envelope based on all the suggested hints and directs the APs to start following the new envelope starting at a given real-time in the future.
- If the AP detects that a particular client is observing a low delivery ratio, it upgrades or maps the client to the next power level.

Observe that the behavior above decouples synchronization of APs from communication of new envelopes between the controller and APs, thereby ensuring that the communication delay has minimum impact on the system behavior.

4.2 Packet Scheduling with Contour

Figure 5 shows the scheduling architecture of each AP using Contour for packets on the downlink(AP→Client). Contour employs multiple queues in the downlink direction to enable easy packet scheduling; each queue corresponds to a power level at which the packets can be transmitted to reach the clients successfully. On a packet arrival from the network, an AP identifies the power level of the client for which the packet is destined, and enqueues the packet in the corresponding queue. Packets are dequeued by the pull switch when the output device is ready to send packets.

The envelope provided by the controller defines the transmit power levels to use at each instant of time. Based on the envelope, the envelope tracker runs a timer and sets the *allowed* power level in the pull switch at each transition point. The output device, when ready to send a

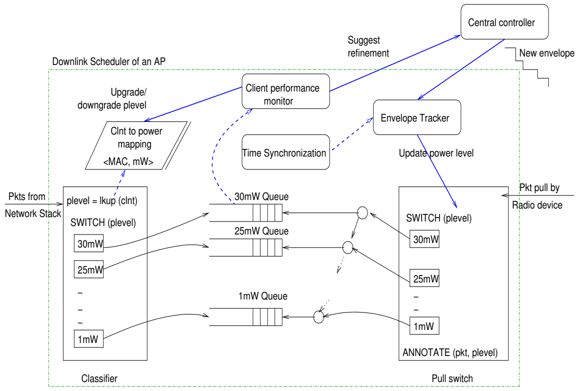

Figure 5: Scheduling architecture of an AP using Contour for downlink traffic.

packet, pulls a packet through the pull switch from the non-empty queue with highest level equal to or below the allowed level. Allowing packets to be pulled from not just the queue with the allowed level, but also from all the queues *below* the allowed level. This is to ensures better utilization. Note, however, that packets from queues at lower levels are still transmitted at the current allowed level to avoid link asymmetry.

4.3 Envelope Refinement

The performance monitor maintains an exponentially weighted moving average of arrival rate, λ_i, for each queue. Based on this rate, the monitor periodically generates refinement hints such that the time spent at each level is proportional to the arrival rate at that level, i.e., the refinement hints are generated as a sequence of tuples

$$\left[\left(P_1, \frac{\mathcal{T}}{n}.\frac{\lambda_1}{\rho_1}\right), \left(P_2, \frac{\mathcal{T}}{n}.\frac{\lambda_2}{\rho_2}\right), ..., \left(P_n, \frac{\mathcal{T}}{n}.\frac{\lambda_n}{\rho_n}\right)\right]$$

where ρ_i represents the maximum transmission rate at power level P_i and $\mathcal{T}$ represents the period of the envelope. The monitor sends the refinement hints to the controller.

In refining the envelope, the controller ensures fair sharing of the channel at each power level across all APs. To achieve this, the controller constructs a new global envelope in two steps. First, each level is assigned a minimum time allocation to avoid starvation. The remaining time out of the period $\mathcal{T}$ is then distributed among the power levels that require more than the minimum allocation. In particular, let t_k^j denote for brevity the time $\frac{\mathcal{T}}{n}.\frac{\lambda_k}{\rho_k}$ at power level P_k requested by AP j in the refinement hints.

1. The controller first calculates the minimum time spent by all APs at each power level as

$$\tau_k = \texttt{MIN}\left[\left(\forall j \ \ \texttt{MAX}\ (t_k^j)\right), \frac{\mathcal{T}}{n}\right]$$

2. Let $\mathcal{S} = \mathcal{T} - \sum_k \tau_k$ represent the total slack out of the period $\mathcal{T}$, where $\sum_k \tau_k$ denotes the total time used up in the first step. Further, let $\mathcal{D}_k = (\forall j \ \texttt{MAX}(t_k^j)) - \tau_k$ denote the deficit at each power level. We distribute this total slack by taking a greedy approach; we iterate over the power levels with non-zero deficit and fulfill the deficit of the highest power level in each iteration until all the residual slack is allocated or there exist no more power levels with non-zero deficit. The remaining slack, if any, is distributed among all the power levels proportional to their current allocation.

The controller sends the new envelope to all the APs with a new real-time far enough in the future such that every AP receives the new envelope before their clocks reach the real-time.

4.4 Adapting Client Power Level

When a new client associates with an AP, the AP chooses the minimum power level (marked as the default level) at which the delivery ratio exceeds a threshold. The AP monitors the delivery ratio of each of its clients at regular intervals and upgrades the client to the next power level if the delivery ratio in the interval is below the threshold. The client is downgraded to its default level every few seconds to ensure that clients will not get permanently upgraded to higher power levels due to transient problems in delivery ratio. Adapting the client power level as above addresses two problems—(1) it addresses the hidden node problem, and (2) it makes the network robust to fluctuating channel conditions. In our prototype, the delivery ratio threshold is set to 90%, the monitoring interval is set to 1 second, and the client is downgraded to its default level 30seconds after it was last upgraded. However, each of the parameters is configurable in our prototype.

To estimate the delivery ratio of each client without assistance from the client, each AP gathers statistics of number of retransmissions, number of packets dropped due to excessive retransmissions, and the total number of packets transmitted to each client. Delivery ratio is then calculated as the fraction of the number of successful transmissions to the total number of transmissions.

5 NS Simulations

In this section, we evaluate Contour through packet-level simulations in ns2.29 over several topologies generated using subsections of the Google Wifi network [5]. Each topology contains APs from a 1x1 kilometer subsection of the Google network, and clients generated at random locations in the same subsection. Each AP has between one and ten clients, with five clients on an average. For brevity, we show results using three topologies; similar results were obtained with other topologies we used. Each AP is considered to have a maximum transmission range of 250m (as in ns2 by default), and clients that are not

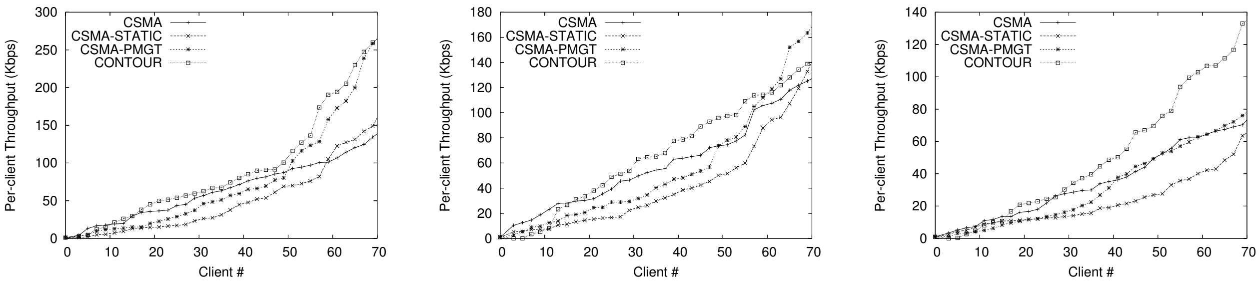

(a) Topology 1 (c) Topology 2 (e) Topology 3

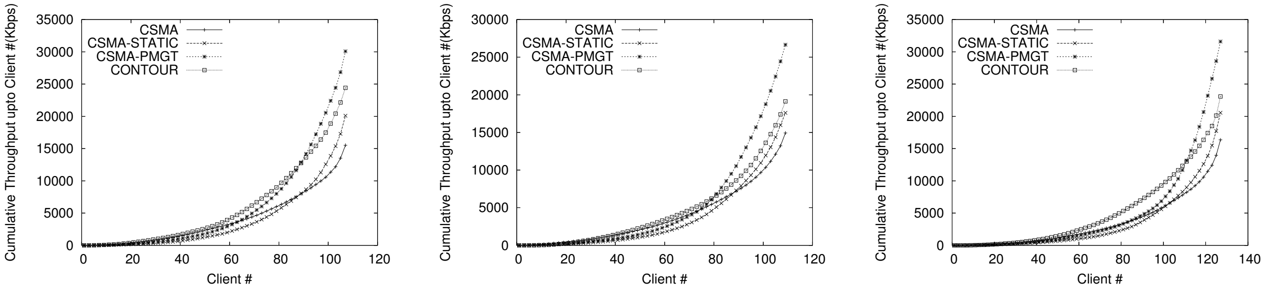

(b) Topology 1 (d) Topology 2 (f) Topology 3

Figure 7: Graphs (a),(c) and (e) show the individual client throughput sorted in the increasing order from left to right, and (b),(d),(f) show the aggregate throughput of the lowest X clients.

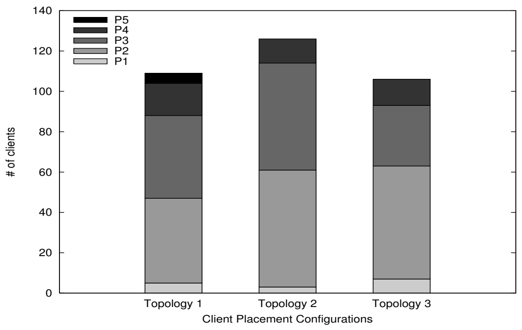

Figure 6: Client power level distribution.

within 250m of any AP are discarded. The APs can transmit at six power levels numbered from 1 to 6 with maximum transmission ranges of 30, 90, 130, 170, 210 and 250 meters respectively. Our ns2 modifications for power control are based on Kawadia's ns2.26 code [24]. At each level, the carrier sense and interference ranges are double the transmission range. Each client associates with the AP closest in distance. Figure 6 shows the number of clients associated at each power level in each topology, and the topologies have 24 APs.

5.1 Algorithms and parameter settings

To demonstrate the efficacy of Contour, we compare its performance with three algorithms.

1. CSMA: 802.11's base-line technique of using CSMA/CA with no power management. APs and clients use power level 6 for transmission.
2. CSMA-STATIC: a static power management technique in which each AP transmits at the power level required to reach its *farthest* client. Clients transmit at the power level used by the associated AP.
3. CSMA-PMGT: per-client dynamic power management technique in which each sender chooses the optimum power for each packet to reach the receiver.

For simulations with Contour, we set the period $\mathcal{T}$ of the schedule to 60ms, and the envelope is refined every 6 seconds. The MAC data rate is set to 11Mbps for all the experiments. Since our measurements with a real system showed no overhead of dynamic transmit power control (TPC), we do not account for any overhead of TPC in our simulations. RTS/CTS is turned off, and we use two-ray path loss radio propagation model. MAC-level synchronous acks for each packet (from both APs and clients) are transmitted at the same power level at which the packet was transmitted.

5.2 Results

For the algorithms above, we study the aggregate network throughput and per-client fairness. Through these simulations, we demonstrate that Contour provides (1) increased network throughput compared to CSMA, and (2) provides better fairness compared to CSMA-STATIC and CSMA-PMGT, and even greater network throughput in some cases.

Fairness: We first demonstrate that Contour has better fairness than CSMA-STATIC and CSMA-PMGT. To do so, in Figure 7(a), we setup backlogged downstream (from APs to clients) CBR traffic in one topology, and plot the individual throughputs of clients using all the four algorithms. The throughputs are sorted in the increasing order. The graph in Figure 7(a) shows that both CSMA-PMGT and CSMA-STATIC have 50 and 60 clients respectively out of 106 whose throughputs are lower than with CSMA. This behavior is because of link asymmetry introduced by the two algorithms. Whereas, Contour has only about 8 clients whose throughputs are below that of CSMA; these clients have lower throughput because of using a global schedule across all APs. This graph demonstrates that Contour makes the throughput of most of the clients better unlike CSMA-PMGT and CSMA-STATIC, thereby exhibiting better fairness properties.

Network throughput: Figure 7(b) shows the aggregated throughput obtained from the same data in Figure 7(a); for instance, the aggregated throughput at x-value 60 represents the sum of the lowest 60 client throughputs. We make two observations. First, Contour gets higher network aggregate throughput than CSMA, as do the other two power management algorithms. Second, the shapes of CSMA-PMGT and CSMA-STATIC curves demonstrate that the throughputs of a few clients dominate and lead to a higher aggregate improvement over CSMA, while most clients have lower throughput than CSMA. We observe similar results in Figures 7(c),(d),(e) and (f) for the other two topologies.

Effect of upstream traffic: One key concern with Contour being applied only at the APs is that the upstream traffic from clients can introduce asymmetry, thereby reducing the benefits of Contour. To study the effect, we setup various amounts of upstream traffic and compare the aggregate network throughput with both CSMA and Contour. Figure 8 varies the amount of upstream traffic generated by each client. The graph shows that Contour improves the downstream aggregate throughput over CSMA even at high upstream traffic. The total throughput (upstream and downstream), however, decreases with increasing upstream traffic. Note that for upstream traffic, clients with Contour use the power level at which they are associated.

Contour's effect on delay: Finally, we study the effect of slotted scheduling in Contour on packet delays in air. Any increased delay can have adverse impact on applications. For instance, if the WLAN hosts VOIP applications, the quality of voice calls can degrade because of increased delays induced by Contour. Similarly, the round-trip time (RTT) for TCP sessions can increase, thereby reducing the throughput of applications (Recall that the throughput of a TCP connection is inversely proportional to its RTT).

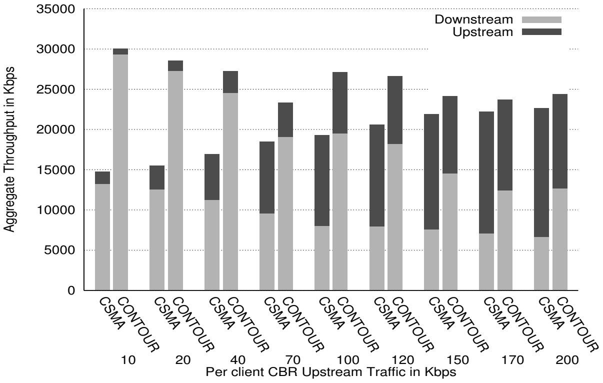

Figure 8: Impact of upstream traffic.

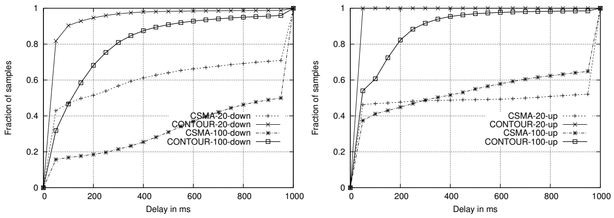

(a) Downstream (b) Upstream

Figure 9: Delay profile for voice-like traffic.

Hence, we perform two experiments—one with voice-like CBR sessions, and one with TCP sessions. The voice-like sessions are two-way and setup between each client and its associated AP; each session generates 50 packets per second in either direction (client to AP and vice versa), each packet is of 100 bytes. We plot in Figure 9(a) and (b) the CDF of the delay between the time when the packet was sent by the sender-side application, and the time when the receiver application received it. The total number of packets successfully sent are normalized to those sent by Contour. We show the results for two cases: with voice sessions setup at 20 randomly selected clients, and with sessions at 100 clients. In both cases, for a given delay budget, Contour is actually able to send greater number of packets than CSMA. Contour also causes minimal impact on TCP traffic; for brevity, we do not show the graph here.

Summary: Although Contour's slotted scheduling causes certain APs to idle during certain periods, the results show that the technique globally provides better fairness and increased throughput.

6 Prototype

To explore the benefits of Contourin a realistic environment, we build a prototype. In this section, we briefly describe our prototype implementation.

6.1 Implementation

We realize the architecture shown in Figure 5 using the Click Modular Router [27]. We implement a new element *SlottedPull*, which is a modified *PullSwitch* element supported in Click. The SlottedPull element implements the packet scheduling logic and the envelope tracker. The envelope tracker uses the *gettimeofday* function for accessing the real-time and runs a timer to update the power levels. We patched the Linux kernel 2.6.16.13 with the high resolution timer subsystem [6] patch, which provides microsecond accurate timers with minimal cpu overhead. The central controller is implemented as a module running on a separate machine that communicates with the SlottedPull element using the Click control socket [1]. APs are connected to the central controller via a wired ethernet connection.

Each AP in our prototype setup is a Dell laptop running Linux 2.6.16.13, and uses an Atheros [1] chipset AR5212 802.11a/b/g PCMCIA card; we use the Linux MadWifi [8] device driver, version 0.9.2. To enable per-packet transmit power control, we recompiled the MadWifi driver with TPC feature enabled. The card supports specifying 60 power levels at a granularity of 0.5dBm; thus a value of 30 corresponds to a transmit power of 15dBm.

6.2 Synchronization with GPS

Each AP in the network using Contour is synchronized closely to a real-time clock to enable all APs to transmit at the same power at any instant of time. While there are several techniques to achieve close synchronization [41], we use a simple and light-weight technique using a GPS receiver at each AP. However, Contour is independent of the specific synchronization mechanism used as long as the synchronization is accurate within an acceptable range.

Our technique is derived from White's [11] proposal. The technique uses clock pulses provided per second by a Garmin LVC-18 GPS receiver [22] through one of the six bare wires it supports. These pulses are used as a timing source by the NTP daemon. The raising edge of this pulse is synchronized to within 1μs of the second boundary of UTC clock time. Another bare wire of the GPS receiver carries NMEA [9] complaint GPS data, which is used for obtaining the absolute value of UTC clock times.

Our microbenchmarks show that the clock skew between the APs using our technique is always below 200μs, and is within 100μs 80% of the time. Since every AP synchronizes to the UTC clock obtained from the GPS receiver, this synchronization technique scales to arbitrary number of APs without the need of a complex coordination protocol.

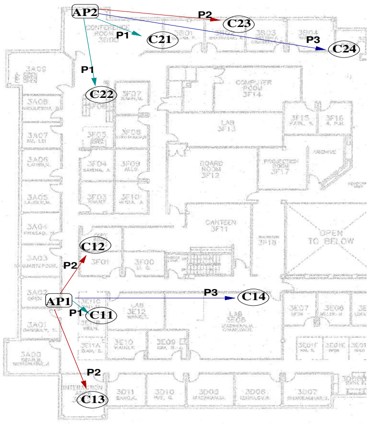

Figure 10: Experimental setup.

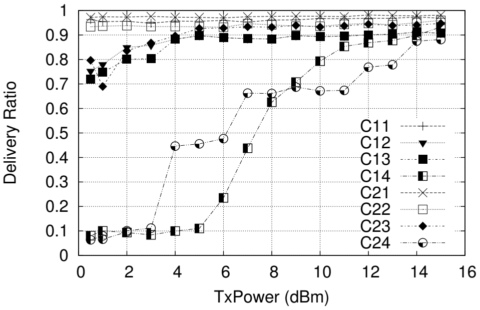

Figure 11: Determining minimum power level for clients.

6.3 Evaluation

In this section, we evaluate different aspects of Contour using our prototype implementation. In what follows, we describe the experimental setup, microbenchmarks and an experiment showing the overall benefits of Contour.

6.3.1 Setup

Figure 10 shows our experimental setup. Our setup contains ten nodes—two 802.11b APs, each with four associated clients. The clients are placed randomly in different office cubicles to emulate a realistic scenario. The APs are marked as AP# and the clients are marked as C#. The APs and clients are Dell laptops with the Atheros PCMCIA cards. The central controller program runs on another Dell laptop that is connected to all the APs through a wired ethernet network. Each AP and the controller are connected to a GPS receiver of their own.

Figure 10 also shows the minimum transmit power levels for each client required by an AP to ensure a delivery ratio of 90%. To determine the minimum transmit power for each client, each AP sends a series of UDP packets at different power levels to each client and measures the delivery ratio. Note that this same technique can be used in a real system when a new client associates with an AP to determine the transmit power level for the client. Fig-

---

[1] The communication protocol between the controller and APs can make use of standard interfaces defined in the CAPWAP protocol [18].

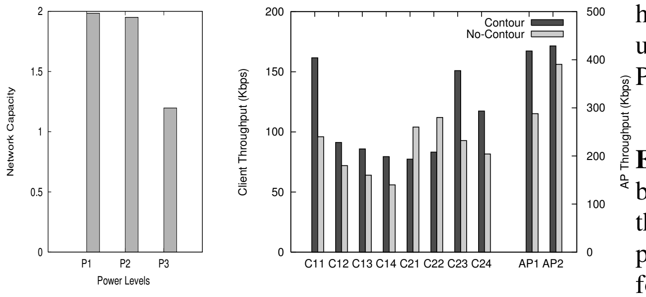

(a) Network capacity (b) Overall Benefits

Figure 14: Benefits in a setup with 2 APs and 8 Clients.

ure 11 shows the graph of delivery ratio with changing transmit power (in dBm) for different clients. Based on the measured delivery ratios, each AP chooses the minimum power level at which the delivery ratio exceeds 90%. In this case, C11, C21, C22 are placed at a power level of P1 (0dBm), C12, C13, C23 are at P2 (=6dBm), and C14, C24 are at P3 (=14dBm).

6.3.2 Microbenchmarks

Scheduling with Contour: The goal of this experiment is to demonstrate that two APs can adhere to the slotted symmetric framework. We choose four clients C11, C14, C21 and C24 for this experiment. C11 and C21 are at the lower power level P1, and C14 and C24 are at the higher power level P3. At P3, the APs can sense each other's transmissions and hence share the channel. At P1, neither of them can sense the other, and hence can transmit simultaneously. We setup four long-running UDP flows from the APs to each of their clients. Each AP uses the following envelope of power levels: $[(P1,20ms),(P2,0ms),(P3,20ms)]$. We plot in Figure 12(a) a point for each packet transmitted to the client using the transmit timestamps of the packet on x-axis. The graph shows that the timestamps of C11, C21 are together, and C14, C24 are together and different from C11, C21. The graph indicates that the two APs are closely following the envelope.

To demonstrate the benefits of following the envelope in this experiment, we perform two experiments. We first set both APs to always operate at the highest power level P3. Figure 12(b) plots the throughput obtained by each of the clients. In this case, since both APs share the channel, all clients should receive the same throughput. We observe, however, that the clients of AP1 get lower throughput compared to AP2 because AP1 is transmitting packets at a lower rate compared to AP2. We repeat the experiment with Contour that uses the above envelope. Figure 12(c) depicts the throughput obtained with Contour for each of the clients. The graph shows that the throughput of C11 and C21 improves substantially because APs transmit simultaneously to C11 and C21. On the other hand, C14 and C24 receive similar throughput as in Figure 12(b) since the APs share the channel at power level P3.

Envelope Refinement: We now demonstrate the feasibility of envelope refinement in a real system. We use the same topology of two APs and four clients as in the previous benchmark. We start with long-running flows for each client. The two power levels P1 and P3 initially have equal slot widths of 20ms. We initiate a series of changes to the flow rates during the course of the experiment and evaluate how the refinement algorithm in Contour is able to adapt to the change in packet arrival rate. Figure 13 shows the changes and adaptation in the envelope (The plot is derived similar to Figure 12(a) by using the packet timestamps). After the 200th millisecond, we halve the packet arrival rate for C11 from 400 pkts-per-second to 200 pkts-per-second. As a result AP1 suggests a refinement of 10ms at P1 and 30ms at P3 to the central controller after monitoring arrival rate in the queue corresponding to P1 for 80ms (two periods). AP2 on the other hand does not observe any change in the arrival rate and hence does not suggest any changes in the envelope. Thus the envelope remains unchanged. Next, flow C21 is halved to 200 pkts-per-second. As a result, both APs independently request for an envelope of 10ms at P1. The controller evaluates the new envelope using both the requests and initiates a change in the envelope at 500th millisecond. Next we stop flows for both C11 and C21 after 800th millisecond. This results in a new envelope that only uses P3 for all slots.

6.3.3 Overall Benefits

This experiment uses all the clients shown in Figure 10 to evaluate the throughput improvement obtained with Contour. In this experiment, we measure the throughput achieved by each client when run together with long-running UDP flows from the APs to the clients. In Figure 14(a), we first plot the network capacity, which is defined as the cumulative throughput obtained when both APs simultaneously transmit. The graph shows that the capacity is high at lower power levels because of greater number of simultaneous transmissions. At power level P3, both APs share the channel. Figure 14(b) shows the throughput achieved by each of the eight clients with and without Contour. The graph clearly demonstrates that cumulative throughput is increased with Contour. However, some clients receive lower throughput with Contour than the case without Contour. This is because, Contour currently provides fairness at a per-power-level granularity, and not on a per-client basis. We will explore the problem of providing fairness on a per-client basis as a part of future work.

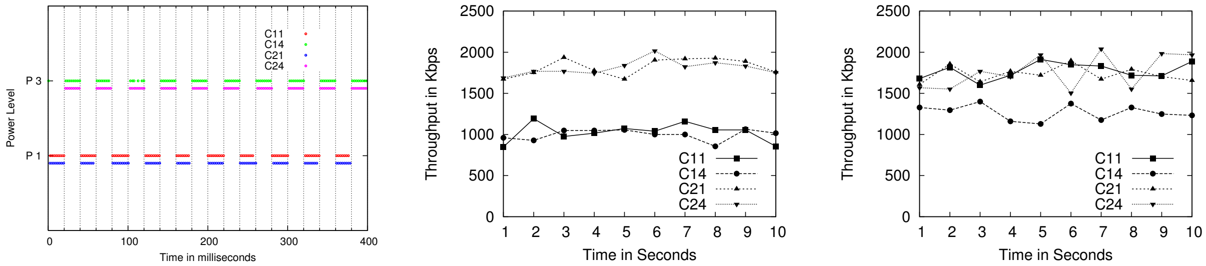

(a) APs following the envelope (b) AP1 and AP2 transmit at P3 (c) AP1 and AP2 use the envelope

Figure 12: Packet scheduling with Contour.

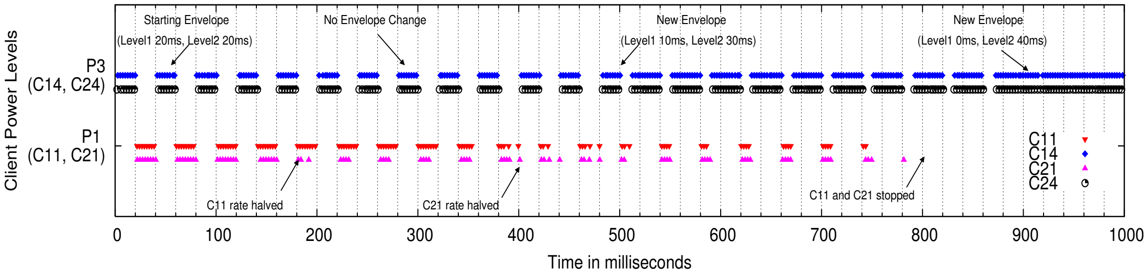

Figure 13: Envelope refinement.

7 Open Issues and Limitations

Although our evaluation demonstrates the benefits of per-client power control in a slotted symmetric framework, Contour is only a first step towards reaping such benefits. Contour can be enhanced in at least three fronts.

First, in this paper, we fix the transmission bit rate of each AP and study the benefits of power control. However, bit rate and transmit power are related; the lower the bit rate required to support the clients' throughput requirement, the lower can be the AP's transmit power. Hence, Contour can further reduce the AP's power level for each client, thereby increasing the benefits of power control.

Second, Contour currently provides fair sharing of resources across clients of the same power level. This can have a bias against power levels that have greater number of clients. As a result, some clients may receive lower throughput with Contour compared to a system without power control. Defining the right notion of client-level fairness and ensuring it in such a framework is an interesting topic that we plan to address in our future work.

Finally, since Contour makes each AP transmit at the same power level at any instant of time, the technique may be inefficient as the network of interfering APs grows in size; increasing number of APs may under-utilize their capacity. A more efficient approach would be to divide the network into smaller clusters such that each cluster includes APs with maximum interfering regions, and the number of interfering regions across the clusters is relatively few. These clusters can then use different envelopes that can be independently scheduled.

One limitation of Contour is that its benefits are maximum when all senders conform to the envelope. If some senders do not conform to the envelope, APs using Contour will eventually adapt their clients that achieve lower delivery ratio to be serviced at higher power levels, thereby reducing the benefits of power control. Nevertheless, Contour will reap the benefits of power control where there are no non-conforming senders or when the non-conforming senders do not contend for resources.

8 Related Work

In wireless networks, research involving transmit power control has focused on two issues – (i) energy saving and (ii) spatial reuse. For energy savings in mobile systems running on battery, a significant portion of research has focused on topology control in ad hoc and sensor networks [36]. The idea is to create useful network topologies that use only the minimum amount of energy in the network. Several other protocols mainly consider sleep and wakeup scheduling for conserving energy [20, 43, 40] or power-aware routing [42] without directly using transmit power control. The 802.11 protocol standard also considers energy savings by scheduling sleep and wakeup periods.

Our focus in this paper is spatial resuse. In [31, 17] the authors studied topology control aspects of transmit power control for interference reduction in a multihop network. Kawadia and Kumar developed the COMPOW protocol [34] in which all nodes have the same minimum transmit power level for the network to be connected. Their clusterPOW [24, 25] protocol addresses a problem with COMPOW when nodes are clustered (so low transmit power is possible within a cluster), but clusters are far apart; thus higher transmit power is needed to connect clusters. Their MINPOW protocol [25], on the other hand, finds a globally optimal solution. These protocols are essentially topology control solutions. The same authors also developed the LOADPOW protocol [25] where a cross-layer approach is proposed to influence routing decisions. Here, the next hop node is not fixed. The next hop is chosen based on the power level that can be safely used without interfering with ongoing communications. In the context of WLANs, Bejerano and Han developed cell breathing technique where power control is used to implement association control [15].

Transmit power control work has been extensively studied in the context of MAC protocols. Several papers developed new MAC protocols – often variations of 802.11 – for multihop networks targeting better spatial reuse [30, 23, 32, 33]. Inherent in these works is a concept of *interference margin* which is the amount of additional interference a receiver can tolerate without collision. Once potential senders know such a margin of the receivers in the neighborhood, they can adjust their transmit power to eliminate possibility of any collisions. However, these solutions require MAC modifications to advertise the margin on a per-packet basis. In several other works, power control is combined with rate control for better system throughput [44, 26, 29, 39]. This is due to the observation that with the same transmit power, i.e., same SINR, the packet capture probability increases with lower bit rate. Our work is limited so far to the same bit rate. But the same framework can be considered to do power and rate control jointly.

In [14], the authors conclude that power control is infeasible with the 802.11 hardware and driver limitations that they experimented with. However, our measurements show that fine-grained power control is feasible now in commodity hardware. The underlying slotted scheduling framework we propose is similar to that proposed in [37, 28, 38]. Their work supports our observations that a packet scheduling framework over commodity 802.11 hardware is practical.

9 Conclusion

We present Contour, a novel dynamic power control technique for improving spatial reuse in managed wireless LAN deployments. Contour is compatible with the current 802.11 standards. In designing Contour, we identify the unfair resource allocation problems resulting from link asymmetry introduced due to dynamic power control, and propose a novel slotted symmetric power allocation framework that can avoid such asymmetry. By implementation and evaluation of Contour, we demonstrate the feasibility of dynamic power control on a slotted framework synchronized across all access points. We believe this feasibility study has wider applicability beyond power control in other wireless domains that require a synchronous slotted framework. Through simulations and prototype evaluation, we characterize the benefits of Contour in terms of spatial reuse and fairness. The results makes a first case for fine-grained power control in current WLANs for improving capacity.

References

[1] Atheros communications. `http://www.atheros.com`.

[2] Cisco Aironet 350 Series Access Points. `http://www.cisco.com/univercd/cc/td/doc/product/wireless/airo_350/accsspts/ap350hig/ap350ch2.htm`.

[3] Cisco Wireless LAN Controllers. `http://www.cisco.com/en/US/products/ps6366/`.

[4] CRAWDAD data on campus wireless deployments. `http://crawdad.cs.dartmouth.edu/data.php`.

[5] Google Wi-fi Coverage. `http://wifi.google.com/city/mv/apmap.html`.

[6] hrtimers - High-resloution timer subsystem. `http://www.tglx.de/hrtimers.html`.

[7] Meru Controller. `http://www.merunetworks.com/products/controllers.php`.

[8] Multiband Atheros Driver for Wifi. `http://madwifi.org`.

[9] National Marine Electronics Association. `http://www.nmea.org/`.

[10] T-Mobile Hotspots. `http://hotspot.t-mobile.com/`.

[11] Using a Garmin GPS 18 LVC as NTP stratum-0 on Linux 2.6. `http://time.qnan.org/`.

[12] P. S. A. Akella, G. Judd and S. Seshan. Self management in chaotic wireless deployments. In *ACM MobiCom 2005*.

[13] S. B. A. Mishra, V. Shrivastava and W. Arbaugh. Partially overlapped channels not considered harmful. In *SIGMETRICS '06*.

[14] F. B. Abdesslem, L. Iannone, M. D. de Amorim, K. Kabassanov, and S. Fdida. On the feasibility of power control in current ieee 802.11 devices. In *PERCOMW '06*.

[15] Y. Bejerano and S. Han. Cell breathing techniques for load balancing in wireless lans. In *INFOCOM '06*.

[16] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MobiCom '05*.

[17] M. Burkhart, P. von Rickenbach, R. Wattenhofer, and A. Zollinger. Does Topology Control Reduce Interference? In *MOBIHOC*, 2004.

[18] P. Calhoun, M. Montemurro, and D. Stanley. CAPWAP Protocol Specification. *Work in progress, draft-ietf-capwap-protocol-specification-01, IETF, May*, 2006.

[19] J. Camp, J. Robinson, C. Steger, and E. Knightly. Measurement driven deployment of a two-tier urban mesh access network. In *MobiSys 2006*.

[20] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *ACM MobiCom 2001*.

[21] J. Eriksson, S. Agarwal, P. Bahl, and J. Padhye. Feasibility study of mesh networks for all-wireless offices. In *MobiSys 2006*.

[22] Garmin. GPS-18 Technical Specifications. `http://www.garmin.com/manuals/425_TechnicalSpecification.pdf`.

[23] E.-S. Jung and N. Vaidya. A power control mac protocol for ad hoc networks. In *ACM MobiCom*, 2002.

[24] V. Kawadia and P. R. Kumar. Power control and clustering in ad hoc networks. In *INFOCOM*, 2003.

[25] V. Kawadia and P. R. Kumar. Principles and protocols for power control in ad hoc networks. *IEEE JSAC*, 2005.

[26] T.-S. Kim, H. Lim, and J. C. Hou. Improving spatial reuse through tuning transmit power, carrier sense threshold, and data rate in multihop wireless networks. In *Proc. ACM MobiCom Conference*, Sept. 2006.

[27] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The click modular router. *ACM Trans. Comput. Syst.*, 18(3):263–297, 2000.

[28] R. R. Kompella, S. Ramabhadran, I. Ramani, and A. C. Snoeren. Cooperative packet scheduling via pipelining in 802.11 wireless networks. In *E-WIND 2005*, pages 35–40, New York, NY, USA, 2005. ACM Press.

[29] R. R. Kompella and A. C. Snoeren. Practical lazy scheduling in sensor networks. In *SenSys '03*.

[30] J. P. Monks, V. Bharghavan, and W. mei W. Hwu. Power controlled multiple access protocol for wireless packet networks. In *IEEE Infocom*, 2002.

[31] T. Moscibroda and R. Wattenhofer. Minimizing Interference in Ad Hoc and Sensor Networks. In *DIALM-POMC, Cologne, Germany*.

[32] A. Muqattash and M. Krunz. A distributed transmission power control protocol for mobile ad hoc networks. *IEEE Transactions on Mobile Computing*, 3(2):113–128, April/June 2004.

[33] A. Muqattash and M. Krunz. POWMAC: A single-channel power-control protocol for throughput enhancement in wireless ad hoc networks. *IEEE JSAC*, 23(5), May 2005.

[34] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar. Power control in ad hoc networks : Theory, architecture, algorithm and implementation of the COMPOW protocol. In *Proc. European Wireless Conference*, 2002.

[35] D. Pong and T. Moors. Fairness and capacity trade-off in ieee 802.11 wlans. In *LCN '04: Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks (LCN'04)*, pages 310–317, Washington, DC, USA, 2004. IEEE Computer Society.

[36] R. Ramanathan and R. Rosales-Hain. Topology control in multihop wireless networks using transmit power adjustment. In *Proc. of the IEEE INFOCOM*, 2000.

[37] I. Ranani, R. R. Kompella, S. Ramabhadran, and A. Snoren. Efficient cooperative scheduling in 802.11 wireless networks. In *UCSD Technical Report CS2005-0830*, 2005.

[38] A. Rao and I. Stoica. An overlay mac layer for 802.11 networks. In *MobiSys 2005*.

[39] A. Sheth and R. Han. Shush : Reactive transmit power control for wireless mac protocols. In *WICON '05*.

[40] S. Singh and C. S. Raghavendra. PAMAS: Power aware multi-access protocol with signalling for ad hoc networks. *ACM Computer Communications Review*, July 1998.

[41] H.-S. W. So, G. Nguyen, and J. Walrand. Practical synchronization techniques for multi-channel mac. In *MobiCom '06*, 2006.

[42] M. Subbarao. Dynamic power-conscious routing for manets: An initial approach. In *IEEE Vehicular Technology Conference '99*.

[43] Y. Xu, J. S. Heidemann, and D. Estri. Geography-informed energy conservation for ad hoc routing. In *Proceedings of the Intl. Conference on MobiCom*, 2001.

[44] X. Yang and N. Vaidya. Spatial backoff contention resolution for wireless networks. In *IEEE WiMesh*, 2006.