Maximizing Spatial Reuse In Indoor Environments

Xi Liu†

Thesis Committee:
Srinivasan Seshan (co-chair)†
Peter Steenkiste (co-chair)†
David Anderson†
Konstantina Papagiannaki‡

†Carnegie Mellon University
‡Intel Labs Pittsburgh

Abstract

Wireless technologies have gained tremendous popularity in recent years, resulting in a dense deployment of wireless devices in every indoor environment. High densities of wireless devices dramatically increase the severity of the interference problem in these indoor environments. Existing solutions to this problem, including frequency division, time division, and MIMO spatial multiplexing, all have limitations and cannot fully solve the problem.

In this thesis, we show the effectiveness of another solution in mitigating wireless interference in indoor environments: increasing spatial reuse. We explore two techniques: power control and directional transmission. We show in this thesis that both techniques can significantly improve network capacity by allowing simultaneous transmissions.

While both power control and directional transmission have been widely used in outdoor scenarios, different solutions are required to meet the new challenges offered by the unique characteristics of the indoor environments, i.e., both APs and clients are chaotically deployed and there are usually multiple RF paths between the AP and the client. There are two challenges in system design, choosing the appropriate power levels / antenna orientations and choosing the right MAC protocol.

In this thesis, we first evaluate the benefits, i.e., spatial reuse opportunities, of both directional transmission and power control. Based on this evaluation, we present the motivation, challenges, and design of three systems: 1) Speed, a distributed directional antenna system with both directional APs and clients designed to solve the interference problem in the future due to the highest level of spatial reuse, 2) DIRC, a centralized directional antenna system with directional APs designed for enterprise networks with omnidirectional clients, and 3) an omnidirectional antenna system with power control, designed to improve spatial reuse of existing omnidirectional nodes.

1 Introduction

Wireless technologies have gained tremendous popularity in recent years, resulting in a dense deployment of wireless devices in every indoor environment such as conference rooms, enterprise networks and homes. For example, it has been envisioned that in the future, every gadget and every appliance in home will have a wireless radio, which may add up to hundreds of wireless radios in future homes. High densities of wireless devices dramatically increase the severity of the interference problem in these indoor environments.

From a physics point of view, there are three orthogonal approaches to isolate wireless devices and to prevent wireless interference: frequency, time, and space. First, multiple wireless devices that operate on orthogonal frequencies do not interfere with each other. However, this approach is limited by the fact that the amount of unlicensed frequency is not infinite. Second, interfering devices that operate in the same wireless channel can avoid interference by not transmitting at the same time, by using CSMA and TDMA. This approach does not suffice either because it does not improve network capacity, i.e., the link throughput of each node is inversely proportional to the number of interfering nodes. Finally, existing space domain solutions focus on spatial diversity, e.g., spatial multiplexing with MIMO antenna configurations. The network capacity of such MIMO networks, however, is limited by multipath channel conditions, i.e., the rank of the channel matrix between the sender and the receiver.

In this thesis, we show the effectiveness of another solution in the space domain in mitigating wireless interference in indoor environments: increasing spatial reuse. We explore two techniques: power control and directional transmission. We show in this thesis that both techniques can significantly improve network capacity by allowing simultaneous transmissions.

Directional transmission can confine the signal at both senders and receivers to a narrow region, which allows the senders to transmit simultaneously without interfering. Power control can also achieve simultaneous transmissions by tuning the power levels on wireless devices. While both power control and directional transmission have been widely used in outdoor scenarios, different solutions are required to meet the new challenges offered by the unique characteristics of the indoor environments. There are two challenges in system design, choosing the appropriate power levels / antenna orientations and choosing the right MAC protocol. Power control has been widely used in cellular phone networks to improve spatial reuse, where power levels are reduced to the minimum levels required to decode the frames. While such power reduction is optimal for celluar networks, i.e., honey-grids of base stations, we show in this thesis that it fails to maximize spatial reuse in indoor environments where both APs and clients are deployed more chaotically. Directional antennas have been widely used in outdoor applications to extend communication range. However, unlike in outdoor scenarios where the sender and the receiver orient directly towards each other, in indoor environments, choosing the right antenna orientations is challenging since there are usually multiple paths between the sender and the receiver and the optimal antenna orientations depend on interfering transmissions. The second challenge for both power control and directional transmission is the MAC protocol, which plays an important role in exploiting spatial reuse opportunities. Unfortunately, existing carrier sensing based solutions, e.g., that tune CCA thresholds or that use directional network allocation vectors (DNAV), interact poorly with power control and directional transmission and thus perform poorly.

The rest of the thesis proposal is organized as follows. Next, we evaluate the benefits, i.e., spatial reuse opportunities, of both directional transmission and power control (Section 2). Based on this evaluation, we present the motivation, challenges, and design of three systems in the next three sections: 1) DIRC, a centralized directional antenna system with directional APs (Section 3), 2) Speed, a distributed directional antenna system with both directional APs and clients (Section 4), and 3) an omnidirectional antenna system with power control (Section 5). The Speed system, which has the highest level of spatial reuse but involves the most deployment efforts, is what we envision to be the solution to the interference problem in the future. The DIRC system is primarily designed for enterprise wireless networks where directional APs are deployed to provide wireless services to omnidirectional users. The power control protocol is designed to improve spatial reuse for existing omnidirectional nodes.

2 Benefiting from Power Control and Directional Antennas

In this section, we use measurements collected in two indoors scenarios to evaluate the spatial reuse opportunites from power control and directional antennas. For power control, we compare no power control (*NoPC*) versus optimal power control (*OptPC*). For directional antennas, we compare four options of directional antenna locations, i.e., on APs (*dir-tx,omn-rx*), on clients (*omn-tx,dir-rx*), on both (*dir-tx,dir-rx*), or neither (*omn-tx,omn-rx*). Since power control is orthogonal to directionality, we consider all combinations, i.e., four directional antenna configurations both with and without power control. In this section, we focus on spatial reuse opportunities, i.e., the best capacity that can be achieved from each combination, by using centralized algorithms called *MaxCAP* and *OptPC* to steer directional antennas and to choose power levels respectively. Both algorithms do an exhaustive search over the entire space.

2.1 Separation Metric

In order to understand the difference between various directional antenna configurations, we propose to use a separation metric that has a strong correlation to the wireless network capacity. Note that currently the separation metric is defined without considering power control, which we plan to address in future work. The separation metric for a network of transmissions is the average of the pairwise SINR of different transmissions. A higher separation indicates a higher chance for simultaneous transmissions. We assume that there are N transmissions in the network, i.e., N APs and N clients. Let $S(AP_i, C_j, k_{AP_i}, k_{C_j})$ denotes the signal strength from AP_i to C_j when the AP orients towards the k_{AP_i} direction, and the client orients towards k_{C_j}.

The separation metric depends on the antenna orientations, but here we present the separation metric for *MaxCAP*. For a pair of transmissions $AP_i \to C_i$ and $AP_j \to C_j$, we define the separation metric $SEP(i, j)$ for *MaxCAP* approach as follows:

$$SEP(i,j) = \max_{k_{AP_i},k_{AP_j},k_{C_i},k_{C_j}} (S(AP_i,C_i,k_{AP_i},k_{C_i}) + S(AP_j,C_j,k_{AP_j},k_{C_j}) - S(AP_i,C_j,k_{AP_i},k_{C_j}) - S(AP_j,C_i,k_{AP_j},k_{C_i}))$$

For a network of N transmissions, the separation metric is $\frac{\sum_{i,j\neq i} SEP(i,j)}{2*N*(N-1)}$. Intuitively, the separation metric is effectively the mean pair-wise SINR values, which is the reason the separation metric has

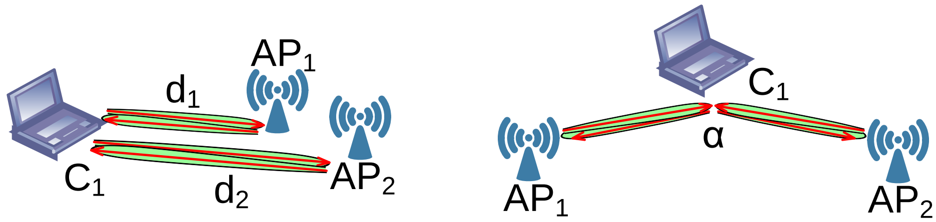

(a) Distance Separation (b) Angular Separation

Figure 1: Two Components of Separation Metric

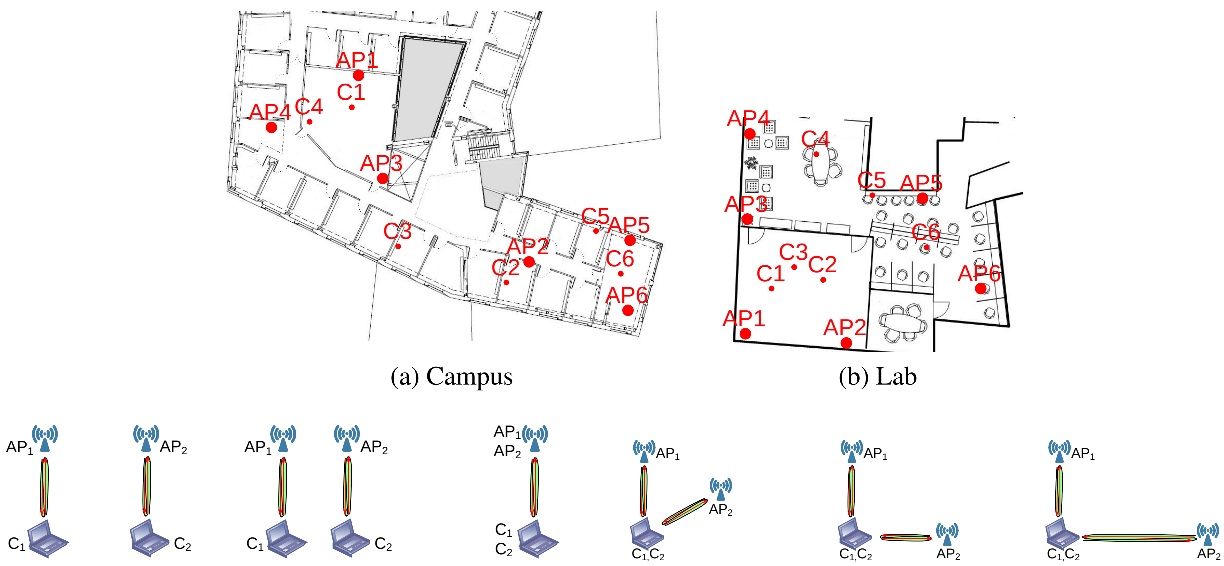

(a) Campus (b) Lab

(c) Topo 1 (d) Topo 2 (e) Topo 3 (f) Topo 4 (g) Topo 5 (h) Topo 6

Figure 2: Experimental Map

a strong correlation with the network capacity. So given that the SINR threshold for 54Mbps is 25dB, for all transmissions to happen simultaneously at 54Mbps, the separation metric need to be at least 25dB.

Distance and Angular Separation In essence, the separation comes from two sources, distance separation and angular separation. The distance separation is due to the difference between the distance (or more accurately, pathloss) from the client to its own AP and to the interfering AP. Figure 1(a) shows an example of distance separation, where the client is closer to its AP than the interfering AP, i.e., $d_1 < d_2$. The angular separation is due to the ability of the directional antennas to focus its energy on a particular direction. Figure 1(b) shows an example of angular separation, where the orientations from the client to two different APs are separated by α. Note that angular separation only applies to directional antennas, and for omnidirectional antennas, the angular separation will always be 0.

The separation is the sum of both angular and distance separation. In order to evaluate the separation of a particular node setup and scenario, measurements need to be taken in that scenario.

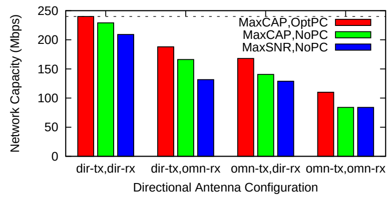

(a) Network Capacity (Campus)

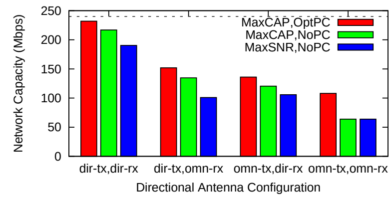

(b) Network Capacity (Lab)

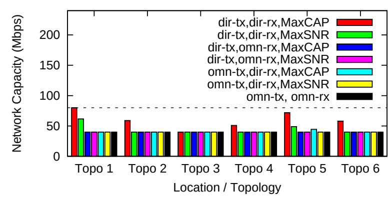

(c) Network Capacity (Topologies)

Figure 3: Network Capacity with Various Antenna Configurations, Power Control, and Orientation Algorithms

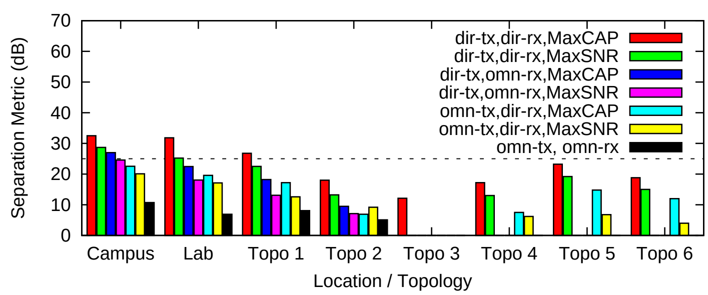

Figure 4: Separation Metric with Various Antenna Configurations and Orientation Algorithms

2.2 Characterizing Spatial Reuse

Now we use measurements in two indoor testbeds to understand the benefits of power control and directional antennas. The experimental setup for this evaluation is as follows: Each client is equipped with two 35 degree fan beam, two 65 degree patch antennas, and one omnidirectional antenna. The reason for such setup will be discussed in Section 4.2. We emulate a directional AP that has 16 directions of 35 degree beams by steering a 35 degree patch antenna using a turn table, and we also use an omnidirectional antenna on the AP to measure the performance of omnidirectional APs.

We take measurements in three different setups. The first setup is an office scenario in a campus building, as shown in Figure 2(a). And the second setup is in a research lab that has more open space than the office scenario, as shown in Figure 2(b). In both scenarios, we place the clients in six locations and the APs in six other locations across the floor. The third setup is a more controlled environment, a large room with tables and machines around, where we construct the different topologies of transmissions as shown in Figure 2(c)-(h), i.e., Topologies 1–6, to illustrate the fundamental problems of the directional APs only configuration (thus we do not evaluate power control in these constructed topologies). We compare different antenna configurations and power control in terms of downlink capacity.

Figure 3(a) and (b) show the network capacity of different power control and directional antenna location combinations in the campus and the lab scenario and compares it with the the upper bound (dotted line). Figure 3(c) shows the network capacity in constructed topologies, where we only focus on directional antenna configuration, i.e., without power control. Figure 4 shows the separation metric—the distance and angular separation—for each directional antenna location. The dotted line at 25 dB shows the separation needed for concurrent 54 Mbps transmissions. The separation for the *omn-tx&omn-rx* case is the distance separation for each scenario, and for other antenna configurations, the angular separation can be calculated by subtracting the distance separation from the measured separation in that case.

Question: What are the benefits of power control? By enabling power control, the network capacity for *omn-tx,omn-rx* case can be improved by 31% in the campus scenario and 70% in the lab scenario. The benefits of power control for directional antennas systems decrease with stronger directionality. For example, for *dir-tx,dir-rx* case, power control only improves capacity by 4% and 7% in the two scenarios. This is because the network capacity of *dir-tx,dir-rx* case is already very close to the upper bound, and leaves little space for improvement.

Question: What are the benefits of directional APs or directional clients? By deploying directional antennas on the APs, the network capacity can be improved over omnidirectional antenna networks, even with power control, i.e., 30% and 51% in the two scenarios. When power control is also enabled for wireless networks with directional APs, the improvement is even higher, i.e., 41% and 71% respectively.

In fact, in Section 5, we illustrate the limitation of power control, i.e., power control essentially redistributes SINR values (or separation) among various links. On the other hand, directional antennas can increase separation by the amount of angular separation.

Since the directional antennas on the clients have weaker directionality than that on the APs, the netowrk capacity of *omn-tx,dir-rx* case is a bit worse than that of *dir-tx,omn-rx* case.

Question: What are the benefis of directional APs and clients? Let us first look at the campus and the research lab scenarios. The capacity of *dir-tx&dir-rx* is close to the upper bound, i.e., all APs can transmit at 54 Mbps with reasonable frame loss rate. The reason is that the separation for the *dir-tx&dir-rx* case are higher than the SINR threshold to decode the 54Mbps frames. Also *dir-tx&dir-rx* improves over *dir-tx&omn-rx* by 38% in campus scenario, and 61% in lab scenario.

In the simple topologies (Figure 2(c)-(h)), the *dir-tx&omn-rx* configuration never improves spatial reuse. Since the clients are located close to each other, the distance separation is very small, independent of where the APs are. The angular separation is similarly small when the clients are close (topo 1-2) to each other or co-located (topo 3-6), where it is 0, no matter how narrow the beams on the directional APs are. The results show that when the clients are located close to each other, both distance and angular separation will be very small without directionality on clients. Since such settings are especially common in settings such as meeting rooms, we believe that deploying directional antennas on both the APs and the clients is critical for maximizing spatial reuse in various locations.

2.3 Motivation for Three Systems

Even though our measurement study indicates that directional APs and clients can maximize spatial reuse in various locations, such antenna configuration requires significant deployment efforts. Thus we also need to consider scenarios where some/all nodes use omnidirectional antennas. In this thesis, we pick three different scenarios that require different levels of directional antenna deployment, and we present three systems that maximize spatial reuse in each of the scenario, as follows:

- A distributed system with directional APs and clients: This is the system which we envision to be the solution to the interference problem in the future, i.e., with directional APs and clients, spatial reuse can be maximized. Also, a distributed MAC protocol is required for it to be deployed in all indoor scenarios.
- A centralized system with directional APs: This is the system designed for enterprise networks, which do not require the clients to equip directional antennas. Also APs are usually connected through a wired connection in such scenarios, thus a centralized MAC protocol can be used to maximize spatial reuse in such scenarios.
- A distributed system with power control and omnidirectional antennas: This is the system that can improve spatial reuse of existing networks without additional hardware. Similarly, a distributed MAC protocol is required for it to be deployed in all indoor scenarios.

We first present DIRC, a centralized system with directional APs, in Section 3. Then we present Speed, a distributed system with directional APs and clients, in Section 4. And finally we present the power control protocol in Section 5.

2.4 Status and Proposed Work

We made several observations from our measurement study, but there are still several limitations and we propose to address them as future work.

Dynamics and Beamwidth The measurements collected in this section are limited in terms of dynamics and various beamwidths. One downside of having higher level of spatial reuse, either from stronger directionality or from power control, is that such systems may be more susceptible to dynamics in the environments, i.e., it may require the nodes to reconfigure when dynamic events happen. Thus in order to better evaluate the tradeoffs of power control and antenna configurations, we need to understand how different dynamic events would affect their performance.

Framework to Characterize Spatial Reuse In this section, we proposed the separation metric for directional antennas as a means of understanding their benefits in several scenarios, but this separation metric only focuses on links, i.e., N APs and N clients with fixed associations. Also the separation does not consider power control. Thus, the separation metric defined in this section have many limitations and cannot answer the following questions:

- What if the association is not fixed or there are N APs and M clients where $M >> N$? Can the separation metric capture the best association in these case?
- How to use the separation metric to guide the placement of APs?

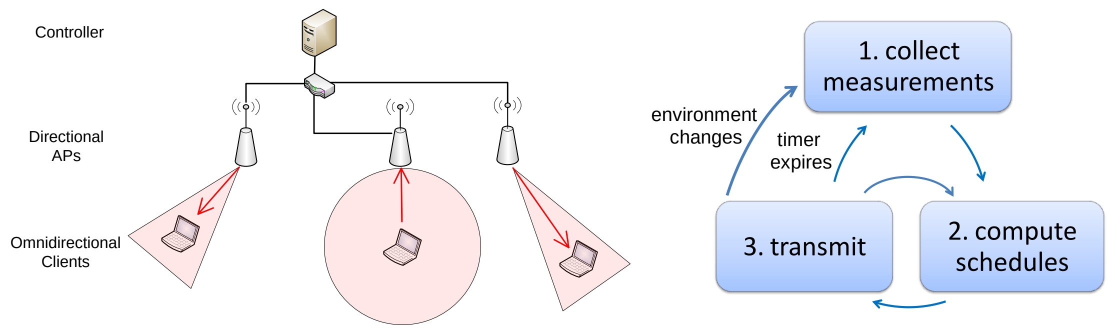

(a) DIRC System Overview (b) DIRC Operations

Figure 5: DIRC Overview and Operations

- How to modify the separation metric for both with and without power control?

The plan is to come up with a framework that can characterize the spatial reuse of both power control and directional antennas. The core of the framework would be a modified version of the separation metric defined in this section.

3 DIRC

In [14], we presented DIRC, an indoor directional antenna system with directional APs that are controlled by a centralized MAC protocol. The DIRC system has been designed, implemented, and evaluated.

DIRC is designed for the following scenario—an infrastructure wireless network where there are n APs and m clients. We assume that all of the APs use phased-array directional antennas, and all of the clients are omnidirectional. Each directional AP has k directions, or, more precisely, antenna patterns. The APs are connected to each other through a separate, independent channel such as wired ethernet that can be used for coordination. Examples of such a scenario can be commonly found in enterprise wireless networks.

3.1 Related Work and Challenges

Directional antennas have been primarily used in outdoor space to extend communication range [23, 28] and to improve connectivity for vehicular wireless networks [21]. In these outdoor systems, the simple strategy of orienting the antennas in the direction of maximum signal strength (SNR) is best, and we call this the *MaxSNR* approach. Choosing antenna orientations for indoor directional antennas is harder because there are multiple paths between the AP and the client.

Simple Antenna Orientation Strategy Does Not Work Figures 6(a) & (b) illustrate that the simple *MaxSNR* approach does not work well because the optimal choice of antenna orientations depend on interfering transmissions. Nodes $S1$ and $S2$ are two directional senders that wish to transmit data to omni-directional receivers $R1$ and $R2$ respectively. Given that there are no obstructions between senders and receivers, the max SNR direction is the same as the LOS direction

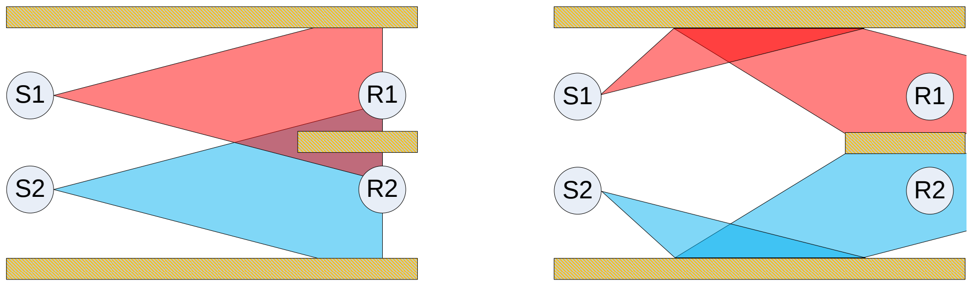

(a) Using max SNR/LOS directions (b) Exploiting reflected paths

Figure 6: Example of exploiting multiple paths using directional antennas

(Figure 6(a)). Unfortunately, the LOS/max SNR directions lead to high interference at the receivers. In this configuration, the MAC protocol must ensure that the two senders never transmit at the same time. In contrast, if the two senders select the orientations shown in Figure 6(b), then both senders could transmit simultaneously. Interference would still exist at the receivers, but it would be weaker, leading to a higher SINR at $R1$ and $R2$, and potentially successful packet receptions.

Figure 3 shows the network capacity from the two algorithms, i.e., *MaxCAP* and *MaxSNR*. Here we focus on *dir-tx,omn-rx* case. In this case, *MaxCAP* outperforms *MaxSNR* by 27% and 34% in the two scenarios. We also did some measurements in [14], which shows more significant improvement in two other indoor testbeds.

Enormous Search Space The greatest drawback of the *MaxCAP* approach is that it requires exploration of all possible orientations of every sender. The size of this search space grows exponentially when potential interferers are also directional. Assuming that all n directional senders in the network can choose any of k directions, the search space to identify the optimal orientation for each one of the senders to their respective receivers is k^n. As directional antenna technology improves, beam widths are likely to become smaller [1], thus increasing k. This will render a brute force approach even more impractical.

Directional MAC Protocol Choosing the correct MAC protocol is crucial to realizing the performance benefits described above. The most important task of any MAC protocol is to identify the set of non-interfering transmissions in an area and to coordinate the activities of the various senders. As we discussed above, the notion of non-interfering transmissions depends on the antenna orientations of the senders. Thus, an indoor directional MAC protocol must not only identify the set of possible concurrent transmissions but also determine their orientations.

One possible choice for a MAC protocol is to use CSMA like in 802.11. However, as earlier research point out, while CSMA works well in networks with omni-directional transmitters, it has several problems in networks with directional transmissions and performs poorly. Past research has proposed a wide range of MAC protocols for directional wireless networks [11, 6, 12, 35, 5, 24, 37, 3]. Much of this work uses RTS/CTS, and directional virtual carrier sensing (DVCS) or directional network allocation vector (DNAV). The basic idea of these solutions is that a direction will be reserved if RTS or CTS is received from that direction, and that direction is marked as unusable during the DVCS/DNAV. One way in which these designs fail to meet our needs is that they assume that the antenna orientation for any sender-receiver pair is fixed (it is a function only of the receiver). This is a reasonable assumption for outdoor settings, where there is a single reasonable orientation (the LOS direction) for any transmission. Another weakness of past MAC designs is that they largely ignore packet capture, i.e., CSMA, DNAV and DVCA approaches tend to be too conservative in scheduling transmissions. Although this observation is not specific to directional networks or indoor environments, we found that taking advantage of capture can significantly improve spatial reuse.

3.2 Design Overview

The core of DIRC's design is to use the SINR model to reduce the number of measurements needed to orient the antennas. In DIRC, only $n * k$ measurements are necessary to implement a heuristics of the *MaxCAP* algorithm, which is the same number of measurements needed to implement the *MaxSNR* algorithm.

Our design of DIRC is based on a central controller that leverages the wired network infrastructure to coordinate the access points (Figure 5(a)). DIRC uses a TDMA based centralized MAC protocol, and in order to prevent the interference from omnidirectional transmissions, each timeslot is split into two phases, dirc-tx phase where only directional APs can transmit, and omni-tx phase where omnidirectional clients can initiate transmissions. In dirc-tx phase, the centralized controller computes a schedule for the directional APs; while in omni-tx phase, the default CSMA MAC is used to coordinate the omnidirectional transmitters. As shown in Figure 5(b), DIRC operates in three stages, 1) collecting measurement, 2) scheduling, and 3) transmitting. The measurements need to be updated on two conditions, 1) the measurements need to be updated periodically, and 2) when the environment changes dramatically.

In the measurement stage, the centralized controller instructs each AP to send a number of frames in each direction, and all receivers would record the signal strength from that AP with a particular direction. During this stage, the set of measurements collected is the full table of $S(i, j, k)$ where i is the AP, j is the client, and k is the direction on the AP.

In the scheduling stage, the centralized controller would apply the SINR model on the collected measurements to determine which APs can transmit and the best antenna orientations for those APs. In DIRC, the centralized controller keeps a FIFO queue of backlogged transmissions, and then try to include as many transmissions to the next timeslot as possible. For each combination of antenna orientations, the controllers computes the SINR value for each transmission link and infers the link throughput achieved by that link, then it computes the network capacity as the sum of all link throughputs. Finally, it picks the best combination of antenna orientations and use that as the schedule for the next timeslot.

In the transmitting stage, the backlogged APs would inform the controller that they have frames to be transmitted to particular clients. Then the controller computes the schedule for the next timeslot and send the schedules to the APs. Then in dirc-tx phase, APs would set their directions and send frames to the clients. In omni-tx phase, the clients can send link-layer ACKs and data frames to the APs.

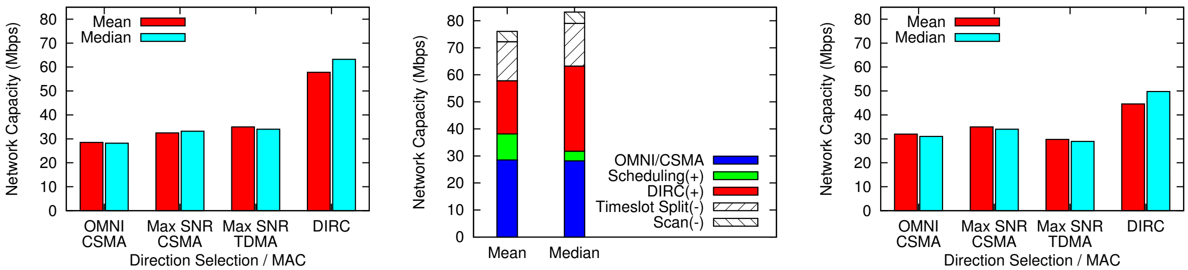

(a) UDP Performance (b) Overhead and gain for UDP (c) TCP Performance

Figure 7: Evalution of DIRC' End-to-End Performance

3.3 Selected Evaluation Results

We present some selected evaluation results of DIRC protocol and implementation. We measure UDP and TCP throughput using the standard *iperf* utility. The packet generation rate for UDP is set to 30 Mbps (TCP manages the rate itself).

Figure 7 shows the UDP and TCP performance of DIRC. In these experiments, the directional APs run *iperf* to each of their clients for 10 seconds, and we repeat this end-to-end measurement for each of the AP's clients. We report mean and median capacity. Although the UDP packet generation rate is 30 Mbps, the maximum actual throughput that DIRC can achieve is only 27 Mbps. This is because by default DIRC reserves 20% of the airtime for client transmissions, though it recovers some of that lost throughput by reducing the inter-frame spacing (IFS); as the maximum effective throughput for 54 Mbps is approximately 32 Mbps, the maximum throughput DIRC can achieve is approximately 27 Mbps. Consequently, the maximum network capacity for 3 directional transmitters is approximately 81 Mbps. Our results show that the median UDP capacity of DIRC is about 76% of this upper bound. Note that CSMA MAC does not have the 20% loss from airtime reserved for the client (nor the gain from reduced IFS), allowing it to achieve close to 30 Mbps (out of 90 Mbps). Figure 7(b) shows the breakdown of DIRC's gain and overhead, where the solid part shows the gain and the patterned part shows the overhead. It shows that much of the benefits come from directionality of the antennas, and the major overhead for UDP traffic is the reserved 4 ms omni-tx phase. The TCP performance of DIRC is approximately 45 Mbps, which is a 40% improvement over *MaxSNR* and 42% over the default omni-directional antennas.

3.4 Status and Proposed Work

DIRC system has been designed, implemented, and evaluated. There is no proposed work for DIRC.

4 Speed

In [13], we present Speed, a distributed directional antenna system with both directional APs and clients. Speed system has been designed, implemented, and evaluated.

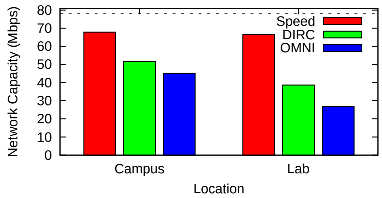

(a) UDP

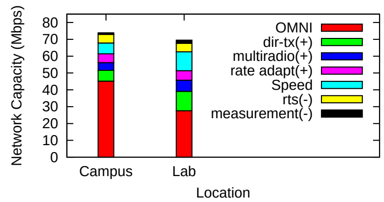

(b) UDP Breakdown

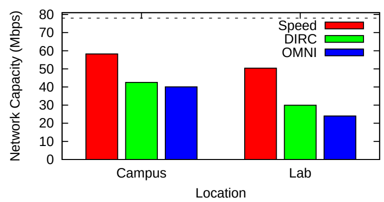

(c) TCP

Figure 8: Evaluation of Speed's End-to-End Performance

Speed is designed for any indoor wireless application scenarios, where 1) both Speed APs and clients use directional antennas to maximize spatial reuse, and 2) Speed nodes use a distributed antenna orientation and MAC protocol in order to be deployable in various application scenarios.

4.1 Related Work and Challenges

As presented in DIRC system (Section 3), the simple strategy of orienting the antennas in the direction of maximum SNR is best for outdoor scenarios, but suboptimal for indoor scenarios.

Figure 3 shows the comparison of the two algorithms. Here we focus on the *dir-tx,dir-rx* case. For *dir-tx,dir-rx* case, *MaxCAP* outperforms *MaxSNR* by about 15% in the two scenarios, and by much more, up to 47%, for the contructed topologies. The results indicate that the improvement is still significantly, especially when clients are located close to each other.

In Section 3, we present a centralized MAC protocol that implements a heuristics of *MaxCAP* orientation algorithm. While such MAC protocol can provide highest level of spatial reuse and may be suitable for certains scenarios such as enterprise networks, such design is too intrusive and unsuitable for some other deployments, e.g., in wireless hotspots and neighborhood wireless networks. Since Speed is designed for any indoor wireless scenario, such centralized approach is unsuitable. On the other hand, *MaxSNR* can be implemented in a distributed manner, but, as shown above, the performance is much worse. The challenge is to design a distributed orientation algorithm and MAC protocol that works as well as their centralized counterparts.

Another challenge for Speed is the problem of client-AP association. In omnidirectional antenna networks, the AP selection problem can be considered as a channel selection problem. When the client chooses a particular channel, it always associates with the closest AP within that channel. But for a directional client, even within the same channel, it may have multiple APs that it can associate with. Here, naively associating a client with the closest AP may be a suboptimal choice. One example is shown in Figure 2(g), Topology 5. Since both clients are co-located, they will always associate with the same AP with this naive association. A better choice is for each client

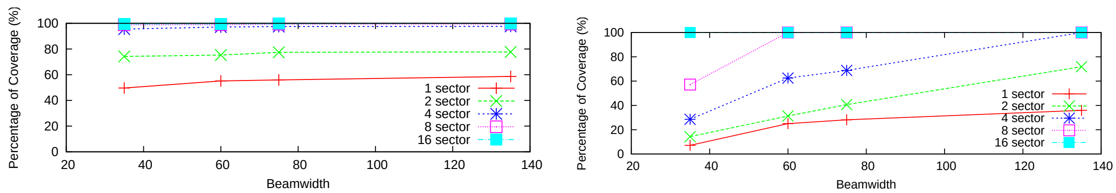

(a) Indoor (b) Outdoor

Figure 9: Coverage of Directional Antennas

to associate with a different AP to improve spatial reuse. Note that this is possible because of the angular separation at the clients, and it would not apply if the clients were omnidirectional. Though associating the client with the most idle AP works in this scenario, it may be suboptimal in other scenarios. For example, consider Topology 5, but this time assume that there is a third AP AP_3 at the exact location of AP_1. In this case, the second client should still associate with AP_2, even if AP_2 has more traffic load than AP_3. This indicates that client-AP association need to consider both the traffic condition and the location of the APs. The challenge is to design an association mechanism that take both factors into consideration.

4.2 How to Put Directional Antennas on APs and Clients?

Let us first consider what type of directional antennas to use on APs and clients. There are two candidates, i.e., phased-array and patch antennas. *Phased-array antennas* [14, 21, 27, 34] use multiple antenna elements to form a directional antenna pattern, which can be oriented electronically in any direction. Although phased array antennas are suitable for deployment on APs, their relatively high price and size prevents them from being deployed on smaller wireless clients. *Patch antennas* are smaller and cheaper, thus we consider deploying them on wireless clients. The naive way to deploy patch antennas is to use antennas with a beamwidth that can cover the whole 360 degrees, e.g. four 90 degree antennas. This provisioning, however, turns out to be unnecessary in indoor environments. The rich scattering in indoor environments allow narrow beam antennas to provide similar coverage to that of wide beam antennas, even with a small number of sectors.

To demonstrate this result, we take measurements in three different locations: (1) in the middle of an outside space, (2) in a campus building, and (3) in a research lab. For the directional clients, we use antennas with different beamwidth, i.e., 35, 60, 75, and 135 degrees, and there are 64 directions at each directional client.

Figure 9 shows the percentage of coverage for the two indoor and one outdoor environments. We define coverage as the percentage of cases where the signal strength is stronger than omnidirectional antennas. The result is averaged over all directional client/omnidirectional AP pairs, and all 64 physical orientations of the client (a total of 10624 location pairs).

Figure 9(a) shows that in indoor environments, even with only 4 sectors, 35 degree antennas can provide as good a coverage as wider beam antennas. As a comparison, Figure 9(b) shows that in the outdoor scenario, it requires 16 sectors of 35 degree antennas to provide good coverage. The reason for the difference is the rich scattering observed in indoor environments that provides multiple antenna orientations with comparable or higher signal strength than the omni-directional antenna. Given the number of sectors, we should choose the minimum beam width that has good coverage (e.g., for four sectors, 35 degree antennas should be used) since this offers the stronger angular separation.

Thus, we propose that directional clients that are primarily used in indoor environments should be equipped with a small number of narrow beam antennas: 35 or 65 degree patch antennas, along with an omnidirectional antenna to handle the scenario where none of the directional sectors can provide strong signal strength. Since APs do not have a strict size constraint, APs can use either sectorized patch antennas or phased-array antennas.

4.3 Design Overview

The MAC protocol in Speed is based on timeslots and timeslot reservations. During each timeslot, non-interfering data traffic from multiple APs will transmit at the same time, regardless of the carrier sensing, to achieve spatial reuse. Other traffic, such as uplink traffic from clients to APs, traffic from other non-protocol compliant transmissions, and all management frames, may also be transmitted anytime, but these frames rely on CSMA mechanism to avoid collisions.

To manage the tradeoff between the *MaxSNR* and the *MaxCAP* algorithms, Speed uses a new distributed algorithm which is a compromise between the *MaxSNR* and the *MaxCAP* approaches. Unlike *MaxSNR*, Speed does consider other transmissions in the network to determine antenna orientations; and unlike *MaxCAP*, it does not rely on exhaustive search and can be implemented distributedly. The idea of Speed's algorithm is for each AP to choose its antenna orientation depending on existing reservations of the timeslot, i.e., the APs can reserve the timeslot with certain antenna orientations if the new reservation of the timeslot do not interfere with any existing reservations on the same timeslot.

When sending data frames, Speed takes the approach presented in [7, 32], to allow clients to send data packets to the APs (e.g., TCP ACKs), and to allow the existence of other non-protocol compliant transmissions. In Speed, APs disable random backoff but keep carrier sensing, for data frames. All management frames are sent to a separate hardware txqueue where default CSMA is used.

In Speed, directional clients determine which AP to associate with. We assume that clients have already chosen the channel it operates in. To simplify presentation, we ignore AP backbone capacity, which is also an important metric in AP selection [22, 33] and assume that the wireless link is the bottleneck. To determine association, the idea is to let the client predict the link throughput if the client is going to associate with a given AP. The predicted throughput is based on the conflict graph and the traffic load on each AP. In Speed, this prediction involves two pieces of information: 1) the number of timeslots that can be allocated to the client; and 2) the expected link throughput for each timeslot. Then the predicted link throughput is the product of the two. This requires traffic load information, which the APs include in their beacon messages. The load information includes the number of timeslots allocated to each client that associates with that AP, and the number of idle timeslots on that AP.

4.4 Selected Evaluation Results

We evaluate both the UDP and TCP performance of Speed in the campus and the lab scenarios (Figures 2(a) and (b)). We placed six clients C_1 to C_6 as indicated on the map, and placed three directional APs in various locations. For each AP location, we activate all possible client combinations, and present the mean capacity from all combinations. For each setup, the experiment runs for 1 minute, and the results are averaged over 3 runs. Figure 8 shows the UDP and TCP performance for Speed, DIRC [14], and OMNI (omnidirectional APs and clients). The results show that in the lab senario, Speed improves UDP performance over DIRC by 100% and over OMNI by 127%, and Speed improves TCP performance over DIRC by 56% and over OMNI by 93%. In campus scenario, Speed improves UDP performance over DIRC by 31% and over OMNI by 50%, and improves TCP performance over DIRC by 36% and over OMNI by 45%. The reason that the improvement in the lab scenario is much higher than that in the campus scenario is that the distance separation in the campus scenario is higher, thus the performance of DIRC and OMNI in the campus scenario is much better.

Figure 8 shows the breakdown of gains and losses for UDP performance for the lab scenario. Compared to the performance of omnidirectional transmissions, the gain of Speed primarily comes from directionality on APs (42%) and additional directionality on clients (85%). But at the same time, the gain will be lower if there is no data rate adaptation (11%), or there is only one radio on each directional client (9%). The overhead of Speed includes reservation traffic (8%) and measurement updates (3%).

4.5 Status and Proposed Work

Speed system has been designed, implemented, and evaluated. However, there are still questions regarding the directional antenna placement on APs. First, there is a question of what type of directional antennas should be deployed on APs. We show in Section 4.2 that both phased-array and patch antennas can be used on APs, but the two technologies still differ in certain ways. For example, phased array antennas are more expensive, but they can change beamwidth and can orient to any desired direction. Such capability may be useful in multicast applications and during dynamic events. Thus a better understanding of dynamics and beamwidth is necessary to determine which type of directional antennas should be deployed in practice. Second, the problem is further complicated in Speed since Speed clients are equipped with a small number of narrow beams antennas. We only show that when the APs are omnidirectional (or similarly when APs have phased-array antennas that can steer to any particular direction), a small number of narrow beam antennas can provide as good coverage as wide beam antennas. However, it is unclear if that still holds when the APs also use a small number of narrow beam antennas.

5 Power Control

In this section, we present a power control protocol for dense wireless networks. This protocol is designed for existing omnidirectional antenna networks that do not require deploying additional hardware. The power control algorithm has been designed, and we plan to design its MAC protocol, implement the protocol, and evaluate it.

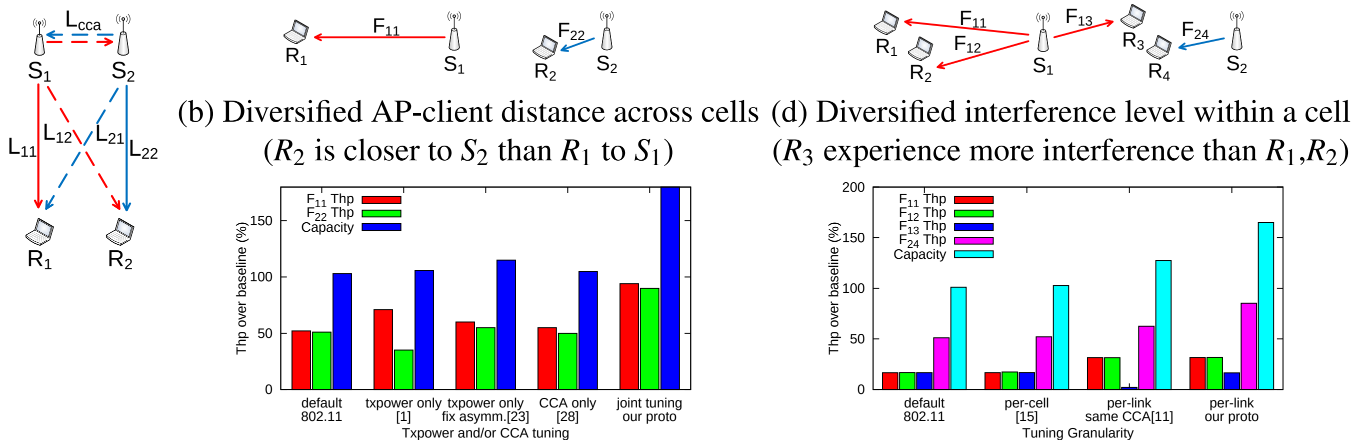

(a)

(b) Diversified AP-client distance across cells (R_2 is closer to S_2 than R_1 to S_1)

(c) Evaluation of (b)

(d) Diversified interference level within a cell (R_3 experience more interference than R_1,R_2)

(e) Evaluation of (d)

Figure 10: For power control: Topologies (a) two flows in general, (b)(c)(d)(e) two example scenarios and evaluation.

5.1 Related Work and Challenges

A wide range of work has explored techniques for tuning transmit power to improve performance. Most work that tunes txpower reduces the power level to the minimum necessary to decode frames at the receiver [30, 31, 2, 25, 16, 9, 17, 18, 29, 19, 26, 20, 38, 8, 4]. The minimum power is chosen for two reasons. First, it minimizes interference; however, this conclusion is based on a range-based model and this model is not accurate with existing hardware. Second, it minimizes energy consumption. In this thesis, we do not focus on energy consumption. It has also been widely observed [31, 2, 16, 9, 17, 18, 29, 26, 20, 4] that using minimum power can cause unfairness or link asymmetry when nodes are controlled by CSMA based MAC protocols, and many previously mentioned systems are designed to deal with this problem. For example, recent work [26] detects such asymmetry and triggers power increase. Protocols that tune CCA thresholds have also been proposed to maximize spatial reuse without power control [36, 39]. In ECHOS [36], wireless nodes set their CCA thresholds to avoid collisions at its own receiver, but at the same time it may cause link asymmetry and exposed terminal problems. In [39], nodes set their CCA thresholds to balance the effect of exposed and hidden terminal problems.

These approaches, however, do not deal well with diverse AP-client distance across cells, as shown in Figure 10(b), where R_1 is relatively farther away from S_1 than R_2 from S_2. Note that in all the examples, we only show transmissions that operate in the same channel, i.e., transmissions in Figure 10(b)&(d) are all in the same channel. In this example, suppose only txpower tuning is deployed and default CCA threshold is used, S_2 will always defer when F_{11} is active, preventing concurrent transmissions. Otherwise S_1 needs to use a very low power level to reduce the carrier sense level at S_2, but at the same time, the low power level prevents it from communicating with R_1. Suppose only CCA tuning is deployed, no matter what CCA thresholds are used on S_1 and S_2, no concurrent transmission is possible because S_1 will cause collisions at R_2.

We constructed the example scenario shown in Figure 10(b) on a testbed and used it to compare the performance of different approaches. We first obtained a baseline throughput by having only one of the sources transmit. We then ran experiments with both sources active, using default configuration, txpower tuning only (minimum power), txpower tuning only with link asymmetry fixed by power increase ([26]), CCA tuning only (ECHOS [36] with link asymmetry fixed), and the performance that can be achieved. The results, as a percentage of baseline, are shown in Figure 10(c). We see that tuning txpower only (minimum power) suffers from slight link asymmetry, which can be fixed by increasing the power on F_{22}. The capacity for both txpower tuning only and CCA tuning only are similar to that of default, i.e. 100% and no spatial reuse. Hoever, it is possible to enable concurrent transmissions in this scenario, increasing capacity to 185%.

Systems that jointly tune power levels and CCA thresholds have also been proposed [10, 15], but they tune at a coarse granularity. For example, [10] assigns the same CCA threshold to every node in the network, and [15] uses the same power/CCA configuration for all the nodes in the same cell. In [19], the authors conclude that coarse-grained approaches are asymptotically optimal, but spatial reuse can still be greatly limited by the "worst" client when the interference level within a cell is diverse.

One example of this limitation is shown in Figure 10(d). In the example, receiver R_3 is in a bad location since flow F_{13} interferes with F_{24}, but otherwise all flows can transmit simultaneously. In this example, concurrent transmissions can only happen without incurring starvation by a per-link protocol, because 1) without per-link txpower, as in [15], the same power level will be used on F_{11}, F_{12}, F_{13}, which imposes the same carrier sensing level on S_2, making S_2 unable to distinguish the three transmissions. S_2, in this case, can either use a high CCA that causes R_3 to starve, or it can use a low CCA that wastes the spatial reuse opportunities with F_{11}, F_{12}. 2) without per-link CCA, as in [10], S_1 can either use a high CCA that let F_{13} and F_{24} interfere with each other, causing F_{13} to starve, or it can use a low CCA which again wastes the spatial reuse opportunities. This example illustrutes that the performance of coarse-grained protocols is limited by the "worst" client.

We constructed this example scenario on our testbed, using the same setup as presented above. In this experiment, we ran the joint txpower and CCA tuning protocol on a per-cell [15], on a per-link basis with a fixed CCA threshold for all nodes [10], and the performance can be achieved. The results are shown in Figure 10(e) The per-cell approach prevents concurrent transmissions, i.e. capacity is around 100%. While the per-link approach with fixed CCA yields higher capacity (125%), one link F_{13} suffers from starvation. Lowering the fixed CCA threshold on both senders fixes the link asymmetry, but the capacity is then almost the same as that of the per-cell approach. In fact, all concurrent transmissions can be enabled in this scenario, yielding a capacity of 164%.

To gain insight into choosing the right power levels, we first consider a two-flow scenario. In Figure 10(a) S_1 transmits to R_1 and S_2 transmits to R_2. We use L_{ij} to denote the path loss from S_i to R_j $(i, j \in \{1,2\})$, and P_i to denote the transmit power level from S_i to R_i. Thus the SINR at receivers R_1 and R_2 are $SINR_1 = P_1 - L_{11} - P_2 + L_{21}$ and $SINR_2 = P_2 - L_{22} - P_1 + L_{12}$, respectively. Note that independent of the transmit power levels, we have $SINR_1 + SINR_2 = L_{12} + L_{21} - L_{11} - L_{22}$. Power control essentially allocates this sum between the two transmissions, i.e. increasing $SINR_1$ will decrease $SINR_2$. In order to enable concurrent transmission, we need both $SINR_1 \geq SINR_{thrsh}$ and $SINR_2 \geq SINR_{thrsh}$. Note that it may not be possible to satisfy the SINR constraints, indicating that concurrent transmissions are impossible. If there are multiple links, however, the choice made for one link may not be compatible with the choices for other links. In fact, finding the optimal configuration is conditional NP-hard, thus the running time to find the optimal power levels is $O(n^k)$, where n is the number of APs, and k is the number of power levels on each node.

5.2 Design Overview

The core of the power tuning algorithm is a greedy heuristic that iteratively tunes power levels to allow more concurrent transmissions by removing edges from the conflict graph. In each iteration, and for each link, the algorithm examines the power levels used on all other links and the topology to determine what power level would allow simultaneous transmission with the other links. It then picks the power level that allows the most concurrency. The new power level will be used if it allows more concurrent transmissions than that in the previous iteration. Note that after each iteration, the number of edges in the conflict graph decreases, and the algorithm converges when no more edges can be removed from the conflict graph. Also, by using this algorithm, the source that needs maximum power level will hit the power limit, and then all other sources will keep an appropriate ratio to that source.

5.3 Preliminary Evaluation

In Figure 10, we evaluated an earlier version of the power control protocol. Figure 10(c)&(e) shows the link throughputs and network capacity of our earlier protocol in the two scenarios of Figure 10(b)&(d). The results show that our earlier protocol can outperform existing solutions, and we expect our new system to perform even better.

5.4 Status and Proposed Work

The iterative power control algorithm has been designed, and we propose 1) to design the MAC protocol, and 2) to implement and evaluate the system. The MAC protocol for power control would be based on the MAC protocol presented for Speed in Section 4, where Speed's MAC protocol needs to be adapted to support the iterative power control algorithm. For example, the iterative algorithm needs the set of active transmissions to determine the appropriate power level on each wireless node, and there is the question of how to determine the set of active transmissions. Timeslot reservations only provide partial information because potential senders may not be able to access the channel.

6 Timeline

We propose the following timeline to finish the thesis:

- Apr. - Jun. 2010: Finish evaluating the effects of dynamics and beamwith on performance, and the framework to characterize spatial reuse (Section 2.4)
- Jun. - Jul. 2010: Finish evaluating different antenna configurations on APs (Section 4.5)
- Jul. - Jan. 2011: Finish the power control protocol (Section 5.4)
- Feb. - May 2011: Tie up loose ends, write thesis

References

[1] *SiBEAM: Wireless Beyond Boundaries (www.sibeam.com).*

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

[3] L. Bao and J. J. Garcia-Luna-Aceves. Receiver-oriented multiple access in ad hoc networks with directional antennas. *Wirel. Netw.*, 11(1-2), 2005.

[4] Y.-J. Choi and K. Shin. Power-adjusted random access to a wireless channel. *INFOCOM*, 2008.

[5] R. R. Choudhury and N. H. Vaidya. Deafness: A MAC problem in ad hoc networks when using directional antennas. *ICNP*, 2004.

[6] R. R. Choudhury, X. Yang, R. Ramanathan, and N. H. Vaidya. Using directional antennas for medium access control in ad hoc networks. In *MobiCom*, 2002.

[7] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan. Understanding the real-world performance of carrier sense. In *E-WIND*, 2005.

[8] L. Jia, X. Liu, G. Noubir, and R. Rajaraman. Transmission power control for ad hoc wireless networks: throughput, energy and fairness. *WCNC*, 2005.

[9] E. Jung and N. Vaidya. A power control mac protocol for ad-hoc networks. *MOBICOM*, 2002.

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

[11] Y.-B. Ko, V. Shankarkumar, and N. H. Vaidya. Medium access control protocols using directional antennas in adhoc networks. In *INFOCOM*, 2000.

[12] T. Korakis, G. Jakllari, and L. Tassiulas. A MAC protocol for full exploitation of directional antennas in ad-hoc wireless networks. In *MobiHoc*, 2003.

[13] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste. Pushing the envelope of wireless indoor spatial reuse using directional access points and clients.

[14] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste. Dirc: increasing indoor wireless capacity using directional antennas. *SIGCOMM Comput. Commun. Rev.*, 39 (4):171–182, 2009.

[15] V. Mhatre, K. Papagiannaki, and F. Baccelli. Interference mitigation through power control in high density 802.11 wlans. In *Infocom*, 2007.

[16] J. Monks, V. Bharghavan, and W.-M. Hwu. A power controlled multiple access protocol for wireless packet networks. *INFOCOM*, 2001.

[17] A. Muqattash and M. Krunz. Power controlled dual channel (pcdc) medium access protocol for wireless ad hoc networks. *INFOCOM*, 2003.

[18] A. Muqattash and M. Krunz. A single-channel solution for transmission power control in wireless ad hoc networks. In *MobiHoc*, 2004.

[19] 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 *in European Wireless Conference*, 2002.

[20] V. Navda, R. Kokku, S. Ganguly, and S. Das. Slotted symmetric power control in managed wlans. *Technical report, NEC Laboratories America.*

[21] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das. MobiSteer: using steerable beam directional antenna for vehicular network access. In *MobiSys*, 2007.

[22] J. Pang, B. Greenstein, D. McCoy, M. Kaminsky, and S. Seshan. Wifi-reports: Improving wireless network selection with collaboration. In *MobiSys*, 2009.

[23] R. K. Patra, S. Nedevschi, S. Surana, A. Sheth, L. Subramanian, and E. A. Brewer. WiLDNet: Design and implementation of high performance WiFi based long distance networks. In *NSDI*, 2007.

[24] A. Prabhu and S. Das. Addressing deafness and hidden terninal problem in directional antenna based wireless multi-hop networks. *COMSWARE*, Jan. 2007.

[25] D. Qiao, S. Choi, A. Jain, and K. G. Shin. Miser: an optimal low-energy transmission strategy for ieee 802.11a/h. In *MobiCom*, 2003.

[26] K. Ramachandran, R. Kokku, H. Zhang, and M. Gruteser. Symphony: synchronous two-phase rate and power control in 802.11 wlans. In *MobiSys*, 2008.

[27] K. Ramachandran, R. Kokku, K. Sundaresan, M. Gruteser, and S. Rangarajan. R2d2: regulating beam shape and rate as directionality meets diversity. In *MobiSys '09: Proceedings of the 7th international conference on Mobile systems, applications, and services*, pages 235–248. ACM, 2009.

[28] B. Raman and K. Chebrolu. Design and evaluation of a new MAC protocol for long-distance 802.11 mesh networks. In *MobiCom*, 2005.

[29] V. Shah and S. Krishnamurthy. Handling asymmetry in power heterogeneous ad hoc networks: A cross layer approach. In *ICDCS*, 2005.

[30] A. Sheth and R. Han. Adaptive power control and selective radio activation for low-power infrastructure-mode 802.11 lans. In *ICDCSW*, 2003.

[31] A. Sheth and R. Han. Shush: reactive transmit power control for wireless mac protocols. *WICON*, July 2005.

[32] V. Shrivastava, N. Ahmed, S. Rayanchu, S. Banerjee, S. Keshav, K. Papagiannaki, and A. Mishra. Centaur: realizing the full potential of centralized wlans through a hybrid data path. In *MobiCom '09: Proceedings of the 15th annual international conference on Mobile computing and networking*, pages 297–308, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-702-8. doi: http://doi.acm.org/10.1145/1614320.1614353.

[33] K. Sundaresan and K. Papagiannaki. The need for cross-layer information in access point selection algorithms. In *IMC*, 2006.

[34] K. Sundaresan, K. Ramachandran, and S. Rangarajan. Optimal beam scheduling for multicasting in wireless networks. In *MobiCom '09: Proceedings of the 15th annual international conference on Mobile computing and networking*, pages 205–216, New York, NY, USA, 2009. ACM.

[35] M. Takai, J. Martin, R. Bagrodia, and A. Ren. Directional virtual carrier sensing for directional antennas in mobile ad hoc networks. In *MobiHoc*, 2002.

[36] A. Vasan, R. Ramjee, and T. Y. C. Woo. Echos - enhanced capacity 802.11 hotspots. In *Infocom*, 2005.

[37] Z. Zhang. DTRA: directional transmission and reception algorithms in WLANs with directional antennas for QoS support. *Network, IEEE*, 19(3), May-June 2005.

[38] R. Zheng and R. Kravets. On-demand power management for ad hoc networks. *INFOCOM*, 2003.

[39] Y. Zhu, Q. Zhang, Z. Niu, and J. Zhu. On optimal physical carrier sensing: Theoretical analysis and protocol design. *INFOCOM*, 2007.