Pushing the Envelope of Indoor Wireless Spatial Reuse using Directional Access Points and Clients

Xi Liu[†], Anmol Sheth[‡], Michael Kaminsky[§], Konstantina Papagiannaki[§],
Srinivasan Seshan[†], Peter Steenkiste[†]
[†]Carnegie Mellon University, [‡]Intel Labs Seattle, [§]Intel Labs Pittsburgh

ABSTRACT

Recent work demonstrates that directional antennas have significant potential to improve wireless network capacity in indoor environments. This paper provides a broader exploration of the design space of indoor directional antenna systems along two main dimensions: antenna configuration and antenna control. Studying a number of alternative configurations, we find that directionality on APs and clients can significantly improve performance, even over other configurations with stronger directionality. Moreover, it is sufficient to have a small number of narrow beam antennas to achieve such gains, thus making such a solution practical for actual deployment. Designing systems with directional APs and clients for increased spatial reuse comes, however, with a number of challenges in the way the directional antennas are controlled. Antenna control needs to encompass antenna orientation algorithms, an appropriate MAC layer protocol, and novel client-AP association solutions. To overcome these challenges, we propose Speed, a distributed directional antenna control system that is easy to deploy and significantly improves network capacity over existing solutions.

Categories and Subject Descriptors

C.2.2 [**Computer System Organization**]: Computer-Communication Networks

General Terms

Algorithm, Design, Performance

Keywords

Directional Antenna, Indoor Wireless Capacity

1. INTRODUCTION

Wireless technologies have gained tremendous popularity in recent years, resulting in the dense deployment of wireless devices in indoor environments. Such high density leads to increased interference which can greatly impact wireless performance. Among a number of different research directions targeted at addressing this problem is using directional antennas to increase capacity through spatial reuse. Recent work [13] has started this exploration by examining the relatively simple scenario in which steerable directional antennas on APs are controlled by a centralized MAC protocol to schedule concurrent transmissions from the APs to clients.

In this paper, we broaden the exploration of indoor directional antenna systems to understand the larger design space. We examine this design space in two basic dimensions: *antenna configuration* and *antenna control*. We summarize the two design space dimensions in Table 1.

By antenna configuration, we refer to two system design decisions. First, where do we put directional antennas–only on the APs, only on the clients, on both, or neither? And second, what type of directional antennas are most appropriate (phased-array, patch antennas, etc.)? Although it is expected that directional APs and clients can provide higher spatial reuse than other configurations, we show that even with overall weaker directionality over the link, partitioning directionality across both APs and clients can provide a significant increase in network capacity, especially when clients are clustered. With respect to antenna type, we surprisingly find that due to the rich multipath in indoor environments, only a small number of narrow beam antennas are needed on directional clients, enabling a practical and less cumbersome directional antenna deployment on clients.

The second dimension of the design space, antenna control, consists of three design decisions. First, how should individuals nodes decide what antenna orientation to use? Although pointing in the direction of maximum signal strength works well in outdoor environments [15], multipath rich indoor environments are decidedly more complex. The second, related, decision is what type of MAC protocol to use. Should it be, for example, centralized and coordinated, or distributed? The third decision is regarding client-AP association. Should the association be based on existing metrics of the closest AP or the most idle AP, or is there a better choice? Based on our measurements of settings with both directional APs and directional clients, we propose a new, completely distributed antenna orientation algorithm and associated MAC protocol in which each AP chooses its antenna orientation based on existing transmissions in the network. In addition, we present a novel AP association algorithm that considers both AP load and available non-interfering antenna orientations.

Based on our exploration of the design space, we describe the design and implementation of Speed. Speed is the first

MobiCom'10, September 20–24, 2010, Chicago, Illinois, USA.

Dimension		Design Choices
Antenna Configuration	– ant. location	APs/clients/both/neither
	– ant. type	phased-array/patch
Antenna Control	– orientation algo	*MaxSNR*/*MaxCAP*/etc.
	– MAC protocol	centralized/distributed
	– association	closest/most idle/other

Table 1: Design Space for Indoor Directional Antenna Systems

completely distributed directional antenna system that optimizes indoor wireless spatial reuse using directional APs and directional clients. While Speed can be applied to a variety of antenna configurations, in this paper, we focus on one particular setup where Speed APs use phased-array antennas, and Speed clients are equipped with four patch antennas with an additional omni-directional antenna. In contrast to MIMO's multi-radio setup, Speed clients leverage antenna selection diversity and use a single radio. Speed has three components: 1) a timeslot based MAC protocol where each AP can reserve any timeslot that does not interfere with existing reservations on that timeslot, 2) an antenna orientation algorithm that each Speed AP runs to ensure non-interfering operation while transmitting directionally, and 3) a client-AP association algorithm that each Speed client runs periodically to determine which AP to associate with. Our experimental evaluation in two indoor testbeds shows that in practice Speed improves network capacity by up to 100% over existing solutions.

The rest of the paper is structured as follows. In Section 2, we explore the choices of the two design dimensions. Based on the experimental results, in Section 3 and Section 4, we propose the design, implementation, and evaluation of Speed, a new indoor directional antenna system. We summarize the related work in Section 5 and conclude in Section 6.

2. EXPLORING DESIGN SPACE

In this section, we use empirical measurements collected across two indoor testbeds to explore the design space of indoor directional antenna systems: antenna configuration and antenna control. Our exploration of antenna configuration shows that placing directional antennas on both APs and clients leads to a significant improvement in spatial reuse, especially when the clients are located close to each other. To balance the deployment overhead of adding directional antennas to clients, we find that a small number of narrow beam antennas are sufficient. The exploration of antenna control highlights several challenges in developing a practical directional antenna system. Specifically, we show that existing solutions in antenna orientation, MAC protocol, and association mechanism do not work well in the directional setting where both APs and clients are directional.

2.1 Antenna Configuration

Benefits of Directional APs and Clients. There are four choices of where to deploy directional antennas: on APs, on clients, on both, or neither. One would expect that stronger directionality can lead to better performance, but there are two important questions that need to be answered here: 1) how much improvement can be obtained from adding directionality on clients over APs only? And 2) how does additional directionality on the client compare with stronger directionality (e.g., narrower beams) on the APs? We answer these two questions based on measurements taken in two indoor scenarios.

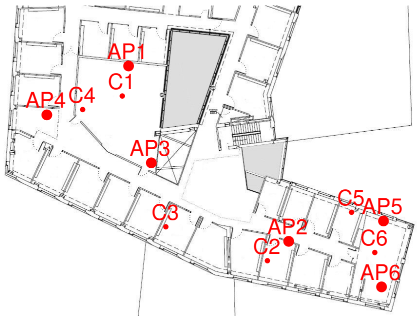

(a) Campus

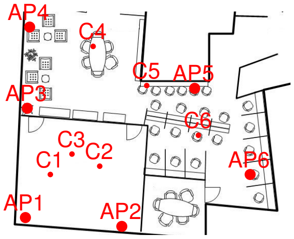

(b) Lab

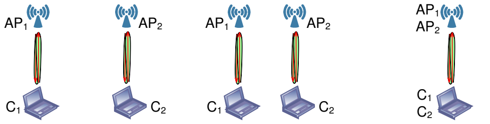

(c) Topo 1 (d) Topo 2 (e) Topo 3

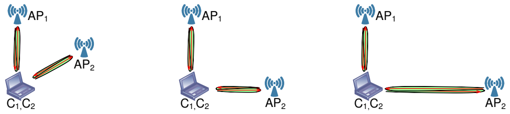

(f) Topo 4 (g) Topo 5 (h) Topo 6

Figure 1: Experimental Map

The primary metrics that we use to measure the benefits of placing directional antennas on APs and/or clients are *network capacity* and the total *directionality cost*. For network capacity, we focus on the best capacity that can be achieved from each configuration by using a centralized algorithm called *MaxCAP* [13] that does an exhaustive search over all possible orientations to find the orientations that provide the maximum capacity. The second metric, the directionality cost, is used to compare directional antenna configurations with different beamwidths and antenna placement. The directionality cost for a particular directional antenna is measured as the minimum number of antenna elements required to implement a phased-array antenna system with the same beamwidth as the directional antenna, i.e., $360/beamwidth$. For example, the directionality cost of a 16° antenna is $360/16 = 22.5$ antenna elements, and the directionality cost of a 35° antenna is $360/35 = 10.3$ antenna elements. The directionality cost of a link is the sum of the directionality cost across the AP and the client.

Directional antennas can usually provide stronger signal strength than omni-directional antennas, which, in turn, can increase the data rate and improve individual link throughput. However, in this paper, we wish to focus only on the improvements achieved by spatial reuse. To eliminate the effects of stronger signal strength from our measurements, we normalize the RSSI readings across the different antennas. For example, if the directional AP has a 5 dBi gain and the directional client has an 8 dBi gain, we subtract 13 dB from the RSSI readings for that link. Thus, the signal strength from a directional antenna will always be the same or more likely, lower than that from an omni-directional antenna. Additionally, it is difficult to have directional antennas with different beamwidths with exactly the same location and size of side-lobes relative to the main lobe. To minimize the impact of side-lobes, we explicitly selected antennas that have similar side-lobe patterns.

First, we compare four different antenna configurations in terms of downlink capacity. Since the majority of the traffic in infrastructure wireless networks tends to be downlink, in this paper, we use the terms AP and tx interchangeably, and we use rx for clients. Thus the four antenna configurations are denoted as *omn-tx&omn-rx*, *dir-tx&omn-rx*, *omn-tx&dir-rx*, and *dir-tx&dir-rx*.

The experimental setup for this evaluation is as follows: Each directional client is equipped with two 35° fan beam patch antennas (135° in the other plane), two 65° patch antennas, and one omni-directional antenna. The reason for this setup will be discussed later in this section. Each directional AP is equipped either with 16 sectors of 35° antennas, 32 sectors of 16° antennas or an omni-directional antenna. The sectors are emulated by rotating the antenna using a mechanical turn table.

We take measurements in three different setups. The first setup is an office scenario in a campus building, as shown in Figure 1(a). The second setup is in a research lab that has more open space than the office scenario, as shown in Figure 1(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 1(c)-(h), i.e., Topologies 1–6. From Topology 1 to 3, the two transmissions are getting physically closer to each other and in Topology 3, both APs and clients are co-located at the exact same location. For Topology 4 to 6, we assume the two clients are co-located and we change the location of the APs. These topologies are constructed to illustrate the fundamental problems of the directional AP only configuration.

We first quantify the benefits of of the *dir-tx&dir-rx* configuration over other configurations. Figure 2(a) shows the network capacity for four different antenna configurations along with the upper bound (dotted line); the results for both the *MaxCAP* and *MaxSNR* algorithms [13] are shown, but here we focus on the *MaxCAP* algorithm (we compare the two algorithms in the next section). The directionality costs for the four configurations are 16.3 (35° APs and effective 60° clients), 10.3 (35° APs), 6 (effective 60° clients), and 0, respectively. Note that our fan beam antennas used on clients have an effective beamwidth of 60° and has a directionality cost of 6. The details of the experiment are described in Section 4.2.

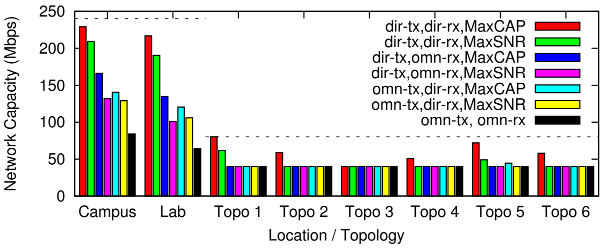

(a) Network Capacity

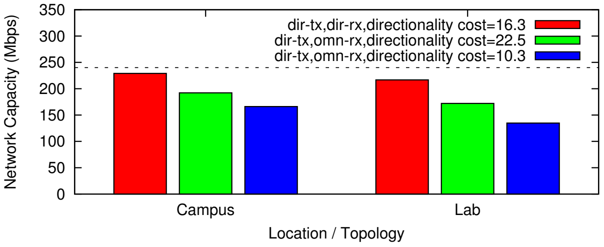

(b) Capacity with Various Directionality Cost

Figure 2: Network Capacity with Various Antenna Configurations and Orientation Algorithms

Let us first look at the different antenna configurations for the campus and the research lab testbeds. For both testbeds, the capacity of *dir-tx&dir-rx* (i.e., with *MaxCAP*) is close to the upper bound, i.e., all APs can concurrently transmit at highest physical data rate with reasonable frame loss rate.[1] Also, *dir-tx&dir-rx* improves over *dir-tx&omn-rx* by 38% in the campus scenario, and 61% in the lab scenario, and improves over *omn-tx&omn-rx* by 170% and 242% respectively. Across the simple topologies (Figure 1(c)-(h)), the *dir-tx&omn-rx* configuration never improves spatial reuse and the network capacity is maximized only with the *dir-tx&dir-rx* configuration. We will further explain these results in Section 4.

Second, while it is expected that additional directionality on clients would improve spatial reuse over using directional APs only, our results show that even with stronger directionality, *dir-tx&omn-rx* still performs worse than *dir-tx&dir-rx*. Figure 2(b) shows the network capacity and directionality cost for three different antenna configurations for the campus and research lab testbeds. Due to hardware limitations we are not able to compare different antenna configurations with the same directionality cost. Comparing the two *dir-tx&omn-rx* configurations, we find that even when the total cost is doubled from 10.3 (35° APs) to 22.5 (16° APs), the network capacity only improves by 15% and 27% for the campus and research lab testbeds respectively. However, distributing the cost across the APs and the clients has a higher payoff. Even with a lower directionality cost of 16.3, the *dir-tx&dir-rx* configuration performs better than the *dir-tx&omn-rx* configuration with a directionality cost of 22.5. The network capacity improves by 37% and 60% over the *dir-tx&omn-rx* configuration with a cost of 16.3. In Section 4, we will show that such configurations are especially important in

---

[1] Note that though the measurements in these testbeds suggest that all links can communicate at the highest rate, our antenna control system can accommodate rate diversity in the network.

improving spatial reuse when clients are located close to each other.

Antenna Type. We now consider what type of directional antennas to use on APs and clients. There are two types of directional antennas that are commonly used in indoor environments. *Phased-array antennas* [13, 14, 17, 21] 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° antennas. In fact, to ensure full coverage, wider beams can also be used, e.g., 4 sectors of 120° antennas. Such provisioning ensures that the signal strength across the directional antennas will never be lower than the signal strength of a single omni-directional antenna configuration. This provisioning, however, turns out to be unnecessary in indoor environments. The rich scattering in indoor environments allows 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 a large outdoor space, (2) in a campus building, and (3) in a research lab. The maps of the campus and the research lab are shown in Figure 1 (though the number of nodes and their locations are different). Ideally we would deploy directional antennas on both senders and receivers, however, in order to obtain an extensive set of locations and physical orientations of the client (we obtained a total of 10624 data points), we deploy 6 omni-directional APs and 6 directional clients in the campus, and 10 omni-directional APs and 13 directional clients in the lab. For the directional clients, we use antennas with different beamwidths (35°, 60°, 75°, and 135°) and number of sectors (1 to 16). In each experiment, we measure the signal strength from all APs to all clients, across the 64 possible orientations of the client.

Figure 3 shows the coverage for the two indoor and one outdoor environments. We define coverage as the percentage of cases where the signal strength is stronger than that of omni-directional antennas. The result is averaged over all directional client/omni-directional AP pairs, and all 64 physical orientations of the client (thus a total of 10624 location pairs).

Figure 3(a) shows that in indoor environments, even with only 4 sectors, 35° antennas can provide as good a coverage as wider beam antennas. As a comparison, Figure 3(b) shows that in the outdoor scenario, it requires 16 sectors of 35° 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° 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 multiple narrow beam antennas: 35° or 65° patch antennas, along with an omni-directional antenna to handle the scenario where none of the directional sectors can provide strong signal strength. Note that a typical 35° patch antenna is large in size, but antennas with a fan beam, e.g., 35° in one plane and 135° in the other plane, can be very small. In our measurement-based study, each client is equipped with two 35° fan beam antennas and two 65° patch antennas. This configuration would be practical for a laptop, but the engineering of new types of patch antennas and antenna placement is out of scope for this paper.

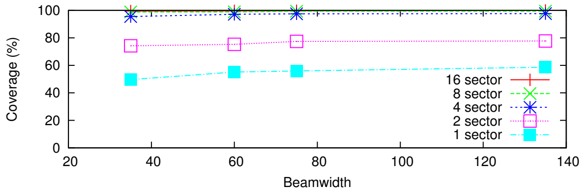

(a) Indoor

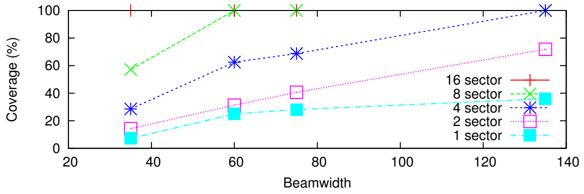

(b) Outdoor

Figure 3: Coverage of Directional Antennas

Since APs do not have strict size constraints, APs can use either sectorized patch antennas or phased-array antennas. In our measurements, we mechanically steer a patch antenna to emulate a directional AP; while in end-to-end protocol evaluation, we use phased array antennas on APs.

2.2 Antenna Control

We have shown that *dir-tx&dir-rx* antenna configuration can provide better performance than *dir-tx&omn-rx* even when using wider beamwidth antennas, especially in cases where the clients are physically located close to each other as in a conference room. We now explore the design choices for antenna control for the *dir-tx&dir-rx* antenna configuration. Antenna control consists of three main sub problems: the design of an antenna orientation algorithm, an appropriate MAC protocol, and a mechanism for the association of directional clients to APs. We summarize the limitations of existing approaches and highlight the challenges involved in developing a solution.

Antenna Orientation and MAC Protocol. In outdoor systems, the simple strategy of orienting in the direction of maximum signal strength (SNR) is best. Choosing antenna orientations for indoor directional antennas is harder because there are multiple paths between the AP and the client, and the optimal choice of antenna orientation depends on interfering transmissions. Below, we present the evaluation of two algorithms [13]: *MaxSNR*, an algorithm that chooses antenna orientation such that the AP-client signal strength is maximized, and *MaxCAP*, an algorithm that picks the optimal antenna orientations by doing an exhaustive search over all antenna orientations.

Figure 2(a) shows the network capacity from different configurations and orientations in each scenario. Let us focus on the *dir-tx&dir-rx* case. The results show that *MaxCAP* outperforms *MaxSNR* by about 15% in the campus and research lab scenarios, and by much more, up to 47%, for the constructed topologies.

A closely related design decision is the choice of MAC protocol that coordinates the APs to choose antenna orientations for the transmissions and to identify non-interfering transmissions. Past systems, e.g. DIRC [13], have used a centralized scheduler-based MAC using *MaxCAP*. However, such designs are intrusive and require coordination among APs, making it unsuitable for some deployments, e.g., in wireless hotspots and neighborhood wireless networks. *MaxSNR* can be implemented in a distributed manner, but, as shown above, the performance is much worse. Thus, the challenge is to design a distributed antenna orientation algorithm and MAC protocol that can perform as well as its centralized counterpart.

Client-AP Association. In omni-directional 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. For the *dir-tx&dir-rx* setup, even within the same channel, the directional client may have multiple APs that it can associate with due to the angular separation of the directional antennas. Here, naively associating a client with the closest AP may be a suboptimal choice. One example is shown in Figure 1(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 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 omni-directional. 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 needs to consider both the traffic conditions and available antenna orientations that minimize interference. The key challenge is to design an association mechanism that considers both factors.

3. SPEED DESIGN

In this section, we present the design of Speed, a distributed directional antenna control system that optimizes spatial reuse. Speed's antenna setup is based on both APs and clients equipped with directional antennas. The antenna control system uses conflict graphs and traffic load estimates to overcome the unique challenges of antenna orientation, MAC protocol, and client-AP association posed by a network of directional APs and clients.

3.1 Background and Assumptions

Since the majority of traffic is downlink, Speed is specifically designed to optimize downlink performance. For uplink traffic, clients use carrier sensing to send data traffic to APs. We also assume that APs and clients are equipped with a single radio that provides antenna selection diversity, which existing commodity Wi-Fi radios already support. Thus, RSSI readings can be measured on all antennas for each incoming frame.

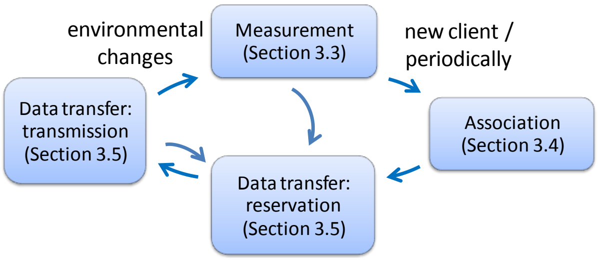

Figure 4: Speed Operation

For notation, we assume that there are N APs and M clients in the network, that have K_{AP} and K_C orientations, respectively. Let $S(AP_i, C_j, k_{AP_i}, k_{C_j})$ denote the signal strength from AP_i to C_j when the AP point in the k_{AP_i} direction and the client orients towards k_{C_j}.

The design of Speed is based on the conflict graph and the SINR model, similar to those used in DIRC [13]. The SINR model is a widely used model which states that whether a frame reception can be decoded at the receiver depends on the SINR, or signal to interference plus noise ratio, of the received frame. If the SINR level is larger than a threshold, then the frame can be decoded; otherwise, the receiver is unable to decode the frame. We also use a generalized conflict graph to encode interference information of the wireless network. First, each vertex in the generalized conflict graph denotes a quadruplet of AP, client, orientations of the AP and the client, i.e., $(AP_i, C_j, k_{AP_i}, k_{C_j})$. Thus, there are $M * N * K_{AP} * K_C$ vertices in the conflict graph. Second, in the generalized conflict graph, an edge between two vertices is annotated by their two pair-wise SINR levels. For any two vertices $(AP_i, C_j, k_{AP_i}, k_{C_j})$ and $(AP_m, C_n, k_{AP_m}, k_{C_n})$, the pair-wise SINR levels are

$$\begin{aligned} PWS(i,j,m,n) = &\quad S(AP_i, C_j, k_{AP_i}, k_{C_j}) \\ &-S(AP_m, C_j, k_{AP_m}, k_{C_j}) \\ PWS(m,n,i,j) = &\quad S(AP_m, C_n, k_{AP_m}, k_{C_n}) \\ &-S(AP_i, C_n, k_{AP_i}, k_{C_n}) \end{aligned}$$

Note that given the pair-wise SINR, we can determine the data rate to be used based on the SINR thresholds of each data rate. Thus in the generalized conflict graph, an annotated edge shows the data rate that can be sustained on both transmissions if they happen simultaneously. This way, the generalized conflict graph does not rely on a particular data rate being used on the link. In order to generate such conflict graphs, each wireless node needs to maintain a table of signal strengths from all APs to all clients with all orientation combinations. The SINR thresholds for each rate are computed offline.

3.2 Overview

The Speed system operates in three phases: 1) measurement phase, 2) association phase, and 3) data transfer phase. Figure 4 shows these three phases, with two sub-phases in the data transfer phase.

The goal of the measurement phase is to construct the conflict graph at every node in the network. This is done by each AP sending probe messages across all its antenna orientations and clients recording the RSSI across all their antennas and reporting the results back to all APs. The measurement process is described in detail in Section 3.3.

In the association phase, the client uses the conflict graph obtained from the measurement phase and traffic load estimates embedded in the AP beacons to determine which AP to associate with. The association process is presented in Section 3.4.

Finally, APs use the antenna orientation algorithm and the supporting MAC protocol for the data transfer on the downlink. As mentioned in previous section, the distributed *MaxSNR* algorithm does not perform as well as the centralized *MaxCAP* algorithm. To manage this tradeoff, Speed uses a new distributed algorithm which is a compromise between the *MaxSNR* and *MaxCAP* approaches. Speed's MAC protocol is based on timeslots (20 ms each with a beacon interval of 1 sec) and timeslot reservations. The idea is that a Speed AP can reserve a timeslot if it can find antenna orientations that do not interfere with existing reservations in that timeslot. This way, unlike *MaxSNR*, Speed does consider the orientations of other transmissions in the network to determine antenna orientations; and unlike *MaxCAP*, it does not rely on exhaustive search and can be implemented without requiring centralized coordination. Speed APs disable random backoff on their data traffic transmit queues [8, 19], which effectively disables carrier sensing for data traffic and relies on carrier sensing only for other traffic such as timeslot reservations, traffic from clients to APs, traffic from external APs, etc. Thus, during each timeslot, APs concurrently transmit non-interfering data traffic 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 such as timeslot reservations, may be transmitted anytime, but these frames rely on CSMA mechanism to avoid collisions. The data transfer phase is presented in Section 3.5.

3.3 Measurement Collection

The goal of this phase is to collect the signal strength information from all APs to all clients with all possible antenna orientations on APs and clients, i.e., the complete table of $S(AP_i, C_j, k_{AP_i}, k_{C_j})$. This information is needed to construct the generalized conflict graph on each node. To obtain this information, each AP sends a number of frames using each of its orientation, and the clients record the RSSI readings of all received frames from each of its sectors. We assume the received signal strength at each receive sector can be measured simultaneously for any incoming frame, which is a valid assumption because existing techniques such as selection diversity already exploit this capability. Note that the measurement process in Speed is a distributed version of the measurement process in DIRC [13].

Measurements must be taken when a client joins the network, and the association process is initiated after the measurement process. The new client first sends a *request-to-scan* frame to the network, notifying all the APs of the new client. Upon receiving such a frame, each AP can initiate its scanning process, which will take five timeslots, or 100ms. Thus, before the scanning process, each AP needs to reserve the next five timeslots by sending its *request-to-scan* frame. If a second AP also needs to initiate its scanning process, it can reserve the subsequent five timeslots by sending its *request-to-scan* frame in the last timeslot of the existing scanning process. In order to avoid frequent measurements, each AP postpones the scanning process if its last scanning finishes in less than one second.

During scanning, the AP transmits a number of scanning probe frames (5ms for each orientation, thus 80ms or four timeslots for 16 orientations) to the network using the lowest data rate. The scanning probe includes the AP and the orientation used. All clients that receive these scanning probes record the mean RSSI readings, the AP identifier, the AP orientation, and the receive sector. During the fifth timeslot, all the clients send the scan response frames to all other nodes in the network. If another AP would like to initiate scanning next, it sends a *request-to-scan* frame in this timeslot. Since these message are critical for Speed to work, all the scan response frames are sent twice to reduce the probability of message loss.

After the measurement phase, all nodes in the network will have all the information necessary to generate the generalized conflict graph, which will be used for downlink data transfer and client-AP association.

3.4 Association

In Speed, clients determine which AP to associate with. The goal of the association phase is for the client to associate with the AP that can provide the highest estimated throughput. We assume that clients have already selected their channel of operation.

The estimated throughput is based on the conflict graph and the traffic load on each AP. The conflict graph is constructed based on the measurements collected. To estimate the traffic load at the client, each AP embeds its traffic load in its beacon messages, which includes the number of timeslots already allocated to clients associated with that AP $(t_1, ..., t_n)$, and the number of remaining idle timeslots t_{idle} on that AP.

The throughput estimation involves two parts: 1) the estimated maximum number of timeslots that can be allocated to the client, t_{exp}; and 2) the expected link throughput for each timeslot, rt_{exp}. Then the estimated link throughput can be calculated as $t_{exp} * rt_{exp}$.

First, suppose the total number of timeslots between beacon intervals is t_{total}, where $t_{total} = \sum_i t_i + t_{idle}$. The client estimates the maximum number of timeslots t_{exp} that can be allocated to it if it associates with that AP as follows: 1) If $t_{idle} \geq \max_i(t_i)$, then $t_{exp} = t_{idle}$, 2) otherwise, $t_{exp} = \max_i(t_i) \times \frac{t_{total}}{\max_i(t_i)+\sum_i(t_i)}$.

Second, the client estimates the link throughput rt_{exp} in each usable timeslot using the conflict graph. Algorithm 1 shows how a client C_j estimates link throughput if it associates with AP_i. The algorithm assumes that the traffic condition is stable from last beacon interval. The idea is to calculate the link throughput from the best antenna orientations that can coexist with transmissions from other APs (lines 6–21). The client first calculates the link throughput if that link operates on its own, *rts* (line 4), which is used to calculate link throughput when there is no interference. Then for each antenna orientation combination (line 6), and for every other AP (line 8), it computes the expected link throughput with regard to that interfering AP, *currt*. This computation accounts for the probability of transmission from the interfering AP (lines 10–16) and the probability of that AP being idle (line 17). If there are multiple APs, the expected throughput is determined by the strongest interfering AP (line 18). Finally, the expected throughput is chosen

Algorithm 1: For association: C_j calculates expected throughput per timeslot from AP_i

```
   Output: expected throughput per timeslot $rt$
1  $rt \leftarrow 0$
2  $rts \leftarrow 0$
3  foreach $k_{AP_i}, k_{C_j}$ do
4      $rts \leftarrow \max(sinr\_to\_thp(S(AP_i, C_j, k_{AP_i}, k_{C_j})), rts)$
5  end
6  foreach $k_{AP_i}, k_{C_j}$ do
7      $rtdir \leftarrow 54$
8      foreach $AP_m \neq AP_i$ do
9          $currt \leftarrow 0$
10         foreach $C_n$ that associates with $AP_m$ do
11             foreach $k_{AP_m}, k_{C_n}$ do
12                 $sinr_1 \leftarrow PWS(i, j, m, n)$
13                 $sinr_2 \leftarrow PWS(m, n, i, j)$
14                 if $sinr\_to\_thp(sinr_1) + sinr\_to\_thp(sinr_2) > thresh$ then $currt \leftarrow currt + (sinr\_to\_thp(sinr_2)) * t_{C_n}/t_{total}$
15             end
16         end
17         $currt \leftarrow currt + t_{idle}/t_{total} * rts$
18         $rtdir \leftarrow \min(rtdir, currt)$
19     end
20     $rt \leftarrow \max(rt, rtdir)$
21 end
22 return $rt$
```

Algorithm 2: For data transfer: AP_i determines orientations for C_j

```
   Input: set of received request-to-send frames $R$, set of orientations on AP $K_a$, set of orientations on client $K_c$
   Output: orientations used on AP $dir_{AP_i}$, on client $dir_{C_j}$, and rate $rt_i$; $rt_i = 0$ indicates no concurrent transmissions possible
1  $rt_i \leftarrow 0$; $dir_{AP_i} \leftarrow -1$; $dir_{C_j} \leftarrow -1$
2  foreach $k_{AP_i}, k_{C_j}$ do
3      $currt \leftarrow 54$
4      foreach $(AP_m, C_n, k_{AP_m}, k_{C_n}, rt_m) \in R$ do
5          $sinr_1 \leftarrow PWS(m, n, i, j)$
6          $sinr_2 \leftarrow PWS(i, j, m, n)$
7          if $sinr_1 > thresh(rt_m)$ then $currt = \min(currt, sinr\_to\_drate(sinr_2))$
8          else $rt \leftarrow -1$; break
9      end
10     if $rt > rt_i$ then
11         $rt_i \leftarrow currt$; $dir_{AP_i} \leftarrow k_{AP_i}$; $dir_{C_j} \leftarrow k_{C_j}$
12     end
13 end
14 return $(rt_i, dir_{AP_i}, dir_{C_j})$
```

among all possible orientations (line 20). Then the expected link throughput from an AP is calculated as $t_{exp} * rt_{exp}$.

Since traffic from each client changes over time, to maximize performance, association should be determined periodically to best balance AP traffic. First, to avoid the client from jumping back and forth among APs, it only switches to a new AP when the expected improvement is above a threshold, whose optimal value depends on specific traffic patterns. We set this threshold to 30% in our experiments. Second, to avoid flash crowd, in which an AP becomes idle and becomes attractive to all the clients around, the reassociation decision should be probabilistic. Ideally, such a probability should depend on the number of clients in the network, but in our system, we simply set the probability to 0.5.

3.5 Data Transfer

After association, the AP may send data frames to the client, using the distributed antenna orientation algorithm and MAC protocol described in this section. Note that the directional MAC protocol in indoor environments for data transfer has two functions: selecting the antenna orientations and identifying non-interfering transmissions.

The MAC protocol ensures non-interfering operation across directional transmissions by allowing new reservations only if they do not interfere with existing reservations. The process can be described in the following steps. First, each AP records all received *request-to-send* reservation frames. Second, when an AP has traffic to send to a client, it checks whether there exists any antenna orientation combination (Algorithm 2, line 2) such that this transmission would not interfere with existing reservation of the timeslot (line 4-9). If no such orientation exists, the AP cannot reserve the timeslot for the client ($rt_1 = 0$ in Algo 2). Otherwise, the AP prepares the *request-to-send* frame that includes the sender, receiver, the computed antenna orientations, and data rate. Note that unlike *MaxSNR*, Speed's Algorithm 2 does consider antenna orientations of other active transmissions, and unlike *MaxCAP*, it involves only local decisions that do not require a global exhaustive search.

The *request-to-send* frame is a management frame and is sent using CSMA/CA. Thus, each competing *request-to-send* frame has equal opportunity to access the channel. When the frame has been successfully transmitted, the reservation has been confirmed, which means that all the other APs can transmit only if they can do so without interfering. If an AP receives a *request-to-send* from another AP while its own *request-to-send* is enqueued , then the AP will remove its own *request-to-send* from the transmit queue and will go back to the second step of Algorithm 2. Also, when a client receives a *request-to-send* intended for it, it would use the receive orientation specified in the *request-to-send* frame for reception. Note that since there may be multiple intended recipients of such *request-to-send* frames, there is no ACK generated. Instead, Speed relies on duplicate *request-to-send* frames and the use of low data rates.

A common technique to support concurrent transmissions is to disable physical layer carrier sensing at the APs. However, this would cause collisions with unscheduled uplink traffic (e.g., TCP ACKs) and other non-protocol compliant traffic. To mitigate this problem, downlink data frame transmissions are subject to two constraints. First, all transmissions occupy the same airtime independent of the data rate. And second, frame transmissions are synchronized. This maximizes spatial reuse achieved in a network with rate diversity. We satisfy these constraints by fragmenting and padding data frames to the airtime required for transmitting 1500 byte packets at 54 Mbps and disabling random backoff [8, 19]. The synchronized transmissions on the downlink lead to synchronized link-layer ACKs causing

ACK collisions. However, surprisingly we observe that the synchronized ACKs were rarely corrupted. The same observation has been made by previous work [19]. All management frames on the downlink are enqueued in a separate hardware transmit queue where default CSMA is used. In the case of the uplink traffic, the AP may not be at the optimal receiving orientation. However, since these frames are transmitted with regular CSMA and are non-colliding (i.e., with no or low external interference), APs are still able to successfully decode these frames.

Data Rate Adaptation. Since the SINR level is independent of the data rate used on the interfering link, Speed APs can support existing rate adaptation schemes during a timeslot. In Speed, each AP picks the data rate to be used at the beginning of the timeslot by comparing the current SINR level to the minimum required SINR thresholds for the different rates. However, during the timeslot the selected rate may not be supported due to an inaccurate SINR model, uncalibrated RSSI levels, short-term fading, or mobility in the environment. In such cases, Speed APs can adapt their data rate without impacting other simultaneously transmitting links. While Speed can accommodate a wide range of rate adaptation algorithms, we use the simple auto rate fallback mechanism in our system.

Handling Dynamicity. One potential problem of having strong directionality on both APs and clients is that performance can be much more sensitive to medium dynamics, i.e., current beams are blocked, or the client moves around. Dynamicity in such scenarios can be categorized as 1) changes that affect the current transmission, e.g., node mobility and environmental changes that happen in the sender and receiver beams, and 2) changes that do not directly degrade the performance of current transmission; instead, they make other antenna orientations more attractive. Since in Speed, the client can measure the signal strength of an incoming frame at different receive sectors simultaneously, both mobility and environmental changes may be detected at either the AP or the client. An AP can detect the performance degradation by observing a throughput drop due to frame losses or reduced data rates. A client can further observe changes in the AP's signal strength across different sectors.

4. EVALUATION

In this section, we evaluate the Speed antenna control system described in the previous section. Our evaluation has two parts: 1) a measurement-based evaluation, and 2) an end-to-end evaluation. The purpose of the measurement based evaluation is to perform controlled experiments and compare the performance of Speed's antenna control algorithm with other approaches across different network topologies. To better understand the measurement based evaluation of the different antenna control algorithms, we also propose a heuristic, called the separation metric. The separation metric captures the SINR difference achieved by both directional antennas and path loss based on client-AP locations. Finally, we present an end-to-end implementation and evaluation of Speed and address the challenges of implementing the system using commodity hardware.

4.1 Separation Metric

Intuitively, the performance of Speed is primarily determined by two factors: the network topology (the location of the APs and clients) and the capability of directional antennas and their orientations. In this paper, we call these two factors the distance separation and the 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. The angular separation is due to the ability of the directional antennas to focus its energy on a particular direction.

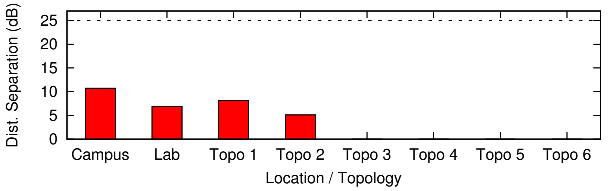

(a) Distance Separation

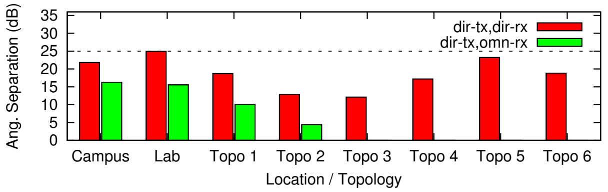

(b) Angular Separation

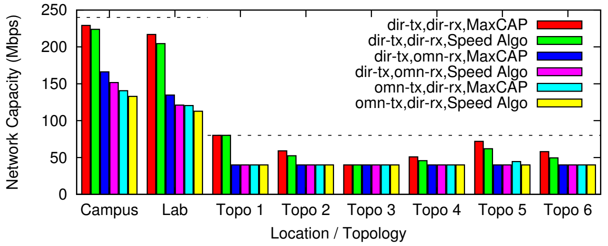

(c) Evaluating Speed's Algorithm

Figure 5: Measurement-based Evaluation

In order to discuss our measurement results, we propose to use a separation metric that captures both factors of network topology and directional antenna capability. Generally speaking, a higher separation metric identifies a network with a higher potential for concurrency.

We assume that there are N transmissions in the network, i.e., N APs and N clients. Let $S(AP_i, C_j, k_{AP_i}, k_{C_j})$ denote 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}.

First, we use the pair-wise SINR (*PWS*) defined in Section 3.1 to define the separation of a pair of links ($AP_i \rightarrow C_i$ and $AP_j \rightarrow C_j$), as $SEP(i,j)$:

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

Note that $SEP(i,j)$ is essentially the sum of SINR values at both receivers C_i and C_j, and for directional antennas, we pick the antenna orientations that can maximize such sum of SINR values.

The separation metric of a network of transmissions is defined as the mean pair-wise SINR of all transmission pairs, i.e., the separation metric for a network of N transmissions is $\frac{\sum_{i,j\neq i} SEP(i,j)}{2*N*(N-1)}$, where $SEP(i,j)$ is the separation or pair-wise SINR for each link pair. Since the separation metric is effectively the mean pair-wise SINR values, and given that the SINR threshold for 54 Mbps is 25 dB, the separation metric needs to be at least 25 dB for all transmissions to occur simultaneously at 54 Mbps.

Note that the separation defined here is the sum of both angular and distance separation, and in fact, the distance separation can be defined the same with omni-directional antennas on both APs and clients (where angular separation is 0).

Figure 5(a)&(b) show both the distance and the angular separation for each topology (we focus on *dir-tx,dir-rx* and *dir-tx,omn-rx* cases). Note that the separation shown in these two figures corresponds to the antenna configurations of *dir-tx,dir-rx*, *dir-tx,omn-rx*, and *omn-tx,omn-rx* with the *MaxCAP* antenna orientation algorithm, thus refer to Figure 5(c) or Figure 2(a) for the corresponding network capacity. The dotted line at 25 dB shows the separation needed for concurrent 54 Mbps transmissions. First, the distance separation is obtained through the setup with omni-directional APs and clients. Then, by measuring the separation for antenna configurations of *dir-tx,dir-rx* and *dir-tx,omn-rx*, we can identify the angular separation by subtracting the distance separation from the measured separation.

4.2 Measurement-based Evaluation

In this section, we present the measurement based evaluation and use the separation metric to explain our results.

Experimental Setup. This experiment consists of a campus scenario, a lab scenario, and six constructed topologies, i.e., Figure 1(a)-(h). In this experiment, we measure the signal strength from all APs to all clients, with all orientation combinations on the AP and the client. Except for our algorithm, which uses its own MAC protocol, all other approaches use an optimal scheduler that does an exhaustive search on all possible schedules. For each schedule, we use the SINR model to compute the link throughput. The network capacity is the sum of all link throughputs. Since our algorithm involves randomness in terms of which transmission sends the *request-to-send* frame first, we enumerate all possible *request-to-send* sequences and the network capacity is averaged over all sequences. Figure 2 and Figure 5 show the results.

Antenna Configuration. In Section 2.1, we presented the measurement results for four antenna configuration. Using the separation metric, we can explain why Speed performs close to optimal in the campus and lab scenarios. From Figure 5 we observe that the separation is higher than the SINR threshold to decode the 54 Mbps frames. This is achieved by a 10 dB angular separation from clients, about a 10 dB angular separation from APs, and about a 10 dB distance separation.

Next we use the separation metric to explain why *dir-tx&omn-rx* performs poorly when clients are clustered or closely located. First, when clients are clustered, the distance separation is small (the distance separation is 0 for co-located clients). In addition, the angular separation is similarly small when the clients are clustered (angular separation is also 0 for co-located clients if clients are omni-directional). 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. Such settings are especially common in settings such as meeting rooms.

Speed's Algorithm. In Section 2.2, we presented the measurement results for two existing antenna orientation algorithms. Next, we present the evaluation of Speed's antenna orientation algorithm. Figure 5(c) shows the performance of Speed's Algorithm, compared with *MaxCAP* and *MaxSNR* (refer to Figure 2(a)). The results show that for *dir-tx&dir-rx*, Speed's performance is very close to *MaxCAP*, i.e., 97% in the campus scenario and 94% in the lab scenario. Speed performs better than *MaxSNR* (refer to Figure 2(a)). The main reason is that the APs now choose antenna orientations considering the antenna orientations on other transmissions. Thus, the separation of Speed's algorithm will be higher than that of *MaxSNR*, achieving similar performance to *MaxCAP*. Note that this is especially true for *dir-tx&dir-rx*, because the separation is high enough and thus the difference in *MaxCAP* and Speed's algorithm is less apparent.

4.3 Implementation

In the implementation of Speed, each directional client is equipped with four 35° patch antennas and an omni-directional antenna. This configuration is different from that mentioned in Section 2.1 because we need to use antennas with the same gain in our end-to-end evaluation. This was not an issue in the measurement study since RSSI readings were normalized according to the antenna gain. Each antenna connects to a separate Atheros 5413 wireless card. Note that Speed assumes a single radio setup on both APs and clients, and we use a multi-radio setup in our prototype to emulate such a system using existing hardware (which allows for two antenna connectors per card, instead of our target five). Note that such a workaround is not fundamental, i.e., future wireless cards can have multiple antenna connectors, and a single radio is enough. Of course, such a setup is expected to improve network capacity over a single radio setup, but we will factor this effect out later by counting the number of data frames from antennas that are not specified in the *request-to-send* frames. All wireless cards are connected to the PCI slots of a desktop machine through PCIe-to-PCI bridges, and they are all set to the same MAC and IP addresses. We implemented the AP using a Phocus phased array antenna which has 16 directions of 45° beams, plus an omni-directional pattern. Phocus array antennas can switch their orientation in $100\mu s$. All APs and clients operate in ad hoc mode with synchronized clocks.

In order to implement the MAC protocol, we had to work around several limitations of the driver. For example, in Speed's MAC protocol, ideally a node would remove the *request-to-send* frame from its transmit queue when it receives a *request-to-send* from another AP. Unfortunately, the driver does not support this, so we instead allow the frames to be transmitted. In order to ensure ordering, each AP inserts two new fields in the *request-to-send* frame: an order number and a random number. Ties are broken either by a lower order number or a lower random number and the AP resends a new *request-to-send* if it loses the tie. This work-around implementation does not work well with a large number of transmissions, but the limitation is not fundamental; we hope that future driver/firmware implementations provide the ability to drop packets from the transmit queue.

For the timeslot based MAC protocol to work, it is nec-

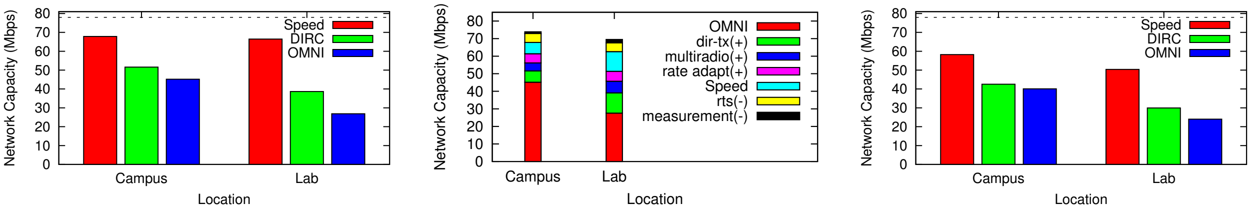

(a) UDP (b) UDP Breakdown (c) TCP

Figure 6: End-to-End Protocol Performance

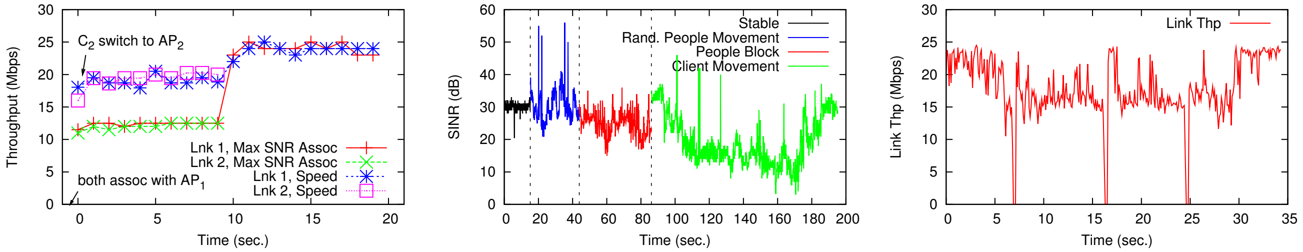

(a) Assocation (b) Impact of dynamic events (c) Handling Mobility

Figure 7: Protocol Behavior

essary that all frames have the same airtime as a 1500 byte frame transmitted at 54 Mbps frame. This approach incurs padding overhead which is dependent on the payload size and data rate. In the worst case, a frame would incur 1499 bytes of padding overhead at 54 Mbps. Fragmenting a full size 36 Mbps frame into two and a full size 24 Mbps frame into three reduces throughput by 25%. Though this limitation is not fundamental and can be reduced by aggregating frames and delaying transmission, we leave such implementation as future work. In UDP test, we simply truncate the data frames such that the duration of each frame is roughly $222\mu s$. For the TCP test, however, truncating the TCP packet would cause byte loss and retransmission; we work around this issue by using the same data rate on all TCP transmissions.

As mentioned in the previous section, channel dynamics may become more critical due to the strong directionality. The primary ways to detect dynamicity is either at APs, or at the clients. We found detecting throughput drop at the APs much more useful than detecting SINR changes at the clients, and this could deal with all dynamic events we introduced in our testbeds. Thus in our implementation, we only implemented detecting throughput drop at APs. A threshold is needed to trigger measurement updates, which is a tradeoff between responsiveness to dynamicity and measurement overhead. We pick a conservative threshold of 50% drop. We choose such a conservative threshold since it's good enough to handle events, such as people walking, and node mobility in our testbeds.

4.4 System Level Evaluation

The goal of this experiment is to evaluate the implementation and the end-to-end performance of the system.

UDP & TCP Performance. We evaluate both the UDP and TCP performance of Speed in the campus and lab scenarios (Figures 1(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. Note that in the measurement-based study, we are able to emulate six APs by putting one AP at six locations at different times, but in this evaluation, we are limited to the three phased-array antennas we have. In the campus scenario, the APs are in $(1, 2, 3), (1, 3, 5), (2, 3, 5)$, respectively; and in the lab scenario, the APs are in $(1, 4, 6), (1, 2, 5), (1, 2, 3), (2, 3, 5), (1, 4, 5)$, respectively. 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 6 shows the UDP and TCP performance for Speed, DIRC [13], and OMNI (omni-directional APs and clients). The results show that in the lab scenario, Speed improves UDP performance over DIRC by 100% and over OMNI by 127%, while Speed improves TCP performance over DIRC by 56% and over OMNI by 93%. In the campus scenario, Speed improves UDP performance over DIRC by 31% and over OMNI by 50%, while TCP performance is improved by 36% over DIRC and by 45% over OMNI. 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. TCP Performance of Speed degrades more than expected, especially in the lab scenario, because we fixed the transmit rate for the TCP experiments.

Figure 6 shows the breakdown of gains and overheads for UDP performance for the lab and campus scenario. In the lab scenario, directional antennas at the transmitter contribute 42% improvement in network capacity over omni-directional antennas. Adding directional antennas on clients increases the performance by 85% over the directional transmitter scenario. Enabling frames to be received across all the radios in Speed's setup improves performance by 9%. The impact of rate adaptation is also important where we see that disabling rate adaptation would lead to a 11% reduction in performance. The figure also shows the protocol performance overhead in the same stacked bar. The overhead of *request-to-send* control frames is 8% of the airtime and measurement updates contribute only 3% of the total airtime.

Association. To evaluate the association process, we use two directional transmissions, constructed as in Topology 5. We emulate the following 20-second scenario with the two transmissions: 1) in the first 10 seconds, both clients are active, and 2) in the next 10 seconds, C_2 leaves the network. APs would send UDP traffic to all associated clients. Figure 7(a) shows such process for Speed and for MaxSNR association. In the first 10 seconds, both clients associate with the same AP AP_1 for MaxSNR association. While in Speed, C_1 associates with AP_1 and C_2 associates with AP_2. Due to such association, the link throughputs from Speed are much improved over that from MaxSNR association. Also note that even though C_2 associates with an AP that does not have the strongest signal, the link throughputs of the two links are comparable. This is because of the randomness in the Speed MAC protocol, i.e., for each timeslot, both links can reserve the channel before the other. After C_2 leaves the network, the link throughputs from both mechanisms are comparable.

Dynamics. In this experiment, we first evaluate the effect of channel dynamics on Speed's performance in the case where Speed does not try to adapt. We enable two transmissions in the research lab, $AP_1 \rightarrow C_1$ and $AP_3 \rightarrow C_3$, and measure the SINR level at C_1. (Here we show SINR values instead of throughputs to portray the real time evolution in client performance.) We measure the following: 1) a baseline SINR where the environment is reasonably stable (i.e., no people movement or mobility); 2) SINR where a person circles around the room; 3) SINR where a person moves back and forth to block the LOS path between AP_1 and C_1; and 4) SINR when we move C_1 to the location of C_2 then to C_3 and finally back to C_1. Figure 7(b) shows the SINR levels on C_1 during this process; the dotted vertical lines indicate the transition from one experiment to the next. The results demonstrate the impact of dynamicity on the SINR level, below and even above the baseline. The SINR level rarely drops so low that the link cannot sustain some (lower) data rate.

Next, based on dynamicity case (4) above, we show how Speed can handle node mobility, which is the most dynamic scenario. Recall that in this case C_1 is moved from the location of $C_1 \rightarrow C_2 \rightarrow C_3 \rightarrow C_1$, while both transmissions $AP_1 \rightarrow C_1$ and $AP_3 \rightarrow C_3$ are active. Figure 7(c) shows the link throughput on $AP_1 \rightarrow C_1$ over time, and during this process, three conflict graph updates are triggered at $t = 6.9s, 16.3s, 24.6s$. While updating the measurements at $t = 6.9s, 16.3s$ significantly improve the link throughput, the update at $t = 24.6s$ seems to only slightly improve link throughput, which may be caused by noise.

5. RELATED WORK

Previous work has used directional antennas of various types to improve wireless network performance [13, 14, 17, 21, 15, 18, 9, 5, 10, 22, 4, 16, 23, 2]. Table 2 summarizes those systems in various dimensions.

In [14, 17], the authors propose to deploy directional antennas on vehicles to improve connectivity and packet reception between vehicles and APs. In [15, 18], the authors use parabolic directional antennas to extend the communication range in outdoor space. In these papers, the interactions between directional nodes are largely ignored. In papers that examine these interactions [9, 5, 10, 22, 4, 16, 23, 2], the authors present various problems with directional antenna systems: directional hidden terminal, directional exposed terminal, and deafness problems. To deal with these problems, they propose various directional MAC protocols, which rely on RTS/CTS, DNAV, and DVCS.

Unfortunately, all of the above systems focus on outdoor settings and do not handle indoor spaces well. Indoor use of directional antennas has received some attention recently. Measurements studies [20, 3] show that directional antennas are still useful even though directionality is weaker in indoor environments. These efforts have focused on measurements in indoor environments and not system or protocol design. While [12] observes that the majority of the benefits of directional antennas come from stronger signal strength instead of spatial reuse, their experiments are carried out in a testbed with relatively long AP-client distance; we focus on dense networks where APs are close to clients. In [11], the authors propose to use multiple default beams to form new desirable shapes. While Speed uses default antenna patterns, this interesting technique can augment Speed, with some added complexity in antenna orientation algorithm. Liu et al. [13] proposed a system called DIRC that does make use of directional antennas in an indoor environment. However, DIRC focuses on the simpler scenario where directional antennas are deployed on APs, and there is a fast, control channel between the APs (e.g., wired Ethernet). Though some of their observations still apply, our work considers a much broader design space.

In addition to spatial beamforming based on phased array antennas, MIMO beamforming can also be used to increase spatial reuse. Recent advances in multiuser-MIMO (MU-MIMO) techniques [7] enable a single AP to serve multiple clients simultaneously by leveraging independent multipath channels at each client. MU-MIMO techniques require channel state feedback across all the clients in the network to characterize the multipath channel [1] and use advanced signal processing techniques like dirty paper coding and block diagonalization [6] to simultaneously transmit to all or a subset of clients.

6. CONCLUSION AND FUTURE WORK

In this paper, we presented a detailed exploration of indoor directional antenna systems that maximize spatial reuse along two dimensions, antenna configuration and control. Based on this exploration, we designed and implemented Speed, an indoor directional antenna control system where both APs and clients are directional. The antenna control system uses conflict graphs and traffic load estimates to overcome the unique challenges of antenna orientation, MAC protocol, and client-AP association posed by a network of directional APs and clients. Our evaluation of Speed in two indoor testbeds show that Speed can indeed maximize spatial reuse by increasing network capacity by 31% and 100% over existing solutions

Speed also has several limitations, which we left as future work. First, while we show that Speed works well in two different indoor testbeds, it is difficult to generalize or pre-

Protocol	Scenario	Antenna Config.	MAC Protocol	Orientation Algo.	Goal
MobiSteer,R2D2 [14, 17]	outdoor	*omn-tx,dir-rx*	distributed	MaxSNR	connectivity
WildNet,RururalNet [15, 18]	outdoor	*dir-tx,dir-rx*	distributed	MaxSNR	comm. range
DMAC, etc. [9, 5, 10, 22, 4, 16, 23, 2]	outdoor	*dir-tx,dir-rx*	distributed	MaxSNR	spatial reuse
DIRC [13]	indoor	*dir-tx,omn-rx*	centralized	MaxCAP heuristics	spatial reuse
Speed	indoor	*dir-tx,dir-rx*	distributed	Algo 2	spatial reuse

Table 2: Summary of Directional Antenna Systems in Various Dimensions

dict performance across different indoor settings. Developing heuristics that can estimate the potential of directional antennas to exploit spatial reuse without detailed measurements is an open problem. Second, adapting the timeslot based MAC protocol to small TCP transfers, e.g., web browsing is also challenging. Several mechanisms may be used to improve the performance under such light and bursty network usage, including reducing timeslot size, allowing APs to send frames to multiple clients in a single timeslot, allowing APs to explicitly cancel a reservation if it is not currently utilized, etc. Third, while four patch antennas can be deployed on certain devices such as laptops, smaller devices like smartphones pose additional challenges in terms of antenna deployment. Also, Speed APs are equipped with phased-array antennas, but in applications such as data transfer between a camera and a laptop, phased-array antennas do not fit on either. We leave the exploration of smart antennas that can leverage spatial reuse on such devices as future work.

Acknowledgements: This research was funded in part by NSF under award numbers CNS-0520192 and CNS-0721857. We also thank our shepherd, Romit Roy Choudhury, and anonymous reviewers for their invaluable feedback.

7. REFERENCES

[1] IEEE Std. 802.11n-2009: Enhancements for higher throughput.

[2] 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.

[3] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan. On the effectiveness of switched beam antennas in indoor environments. In *PAM*, 2008.

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

[5] 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.

[6] T. Cover and J. Thomas. *Elements of Information Theory.* J Wiley and Sons Inc, 1991.

[7] D. Gesbert, M. Kountouris, R. Heath, C.-B. Chae, and T. Salzer. Shifting the MIMO paradigm. *Signal Processing Magazine, IEEE*, 2007.

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

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

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

[11] S. Lakshmanan, K. Sundaresan, S. Rangarajan, and R. Sivakumar. Practical beamforming based on rssi measurements using off-the-shelf wireless clients. In *IMC*, 2009.

[12] S. Lakshmanan, K. Sundaresan, S. Rangarajan, and R. Sivakumar. The myth of spatial reuse with directional antennas in indoor wireless networks. In *PAM*, 2010.

[13] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste. DIRC: increasing indoor wireless capacity using directional antennas. In *SIGCOMM*, 2009.

[14] 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.

[15] 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.

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

[17] K. Ramachandran, R. Kokku, K. Sundaresan, M. Gruteser, and S. Rangarajan. R2D2: regulating beam shape and rate as directionality meets diversity. In *MobiSys*, 2009.

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

[19] 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*, 2009.

[20] A. P. Subramanian, H. Lundgren, and T. Salonidis. Experimental characterization of sectorized antennas in dense 802.11 wireless mesh networks. In *MobiHoc*, 2009.

[21] K. Sundaresan, K. Ramachandran, and S. Rangarajan. Optimal beam scheduling for multicasting in wireless networks. In *MobiCom*, 2009.

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

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