

The Myth of Spatial Reuse with Directional Antennas in Indoor Wireless Networks

Sriram Lakshmanan¹, Karthikeyan Sundaresan², Sampath Rangarajan² and Raghupathy Sivakumar¹

¹ Georgia Institute of Technology, Atlanta, GA, U.S.A

² NEC Labs. America, Princeton, NJ, U.S.A

Abstract. Interference among co-channel users is a fundamental problem in wireless networks, which prevents nearby links from operating concurrently. Directional antennas allow the radiation patterns of wireless transmitters to be shaped to form directed beams. Conventionally, such beams are assumed to improve the spatial reuse (i.e. concurrency) in indoor wireless networks. In this paper, we use experiments in an indoor office setting of Wifi Access points equipped with directional antennas, to study their potential for interference mitigation and spatial reuse. In contrast to conventional wisdom, we observe that the interference mitigation benefits of directional antennas are minimal. On analyzing our experimental traces we observe that directional links do *not* reduce interference to nearby links due to the lack of signal confinement due to indoor multipath fading. We then use the insights derived from our study to develop an alternative approach that provides better interference reduction in indoor networks compared to directional links.

Key words: Indoor wireless networks, directional antennas, spatial reuse

1 Introduction

The growing density of wireless deployments and limited spectrum availability have motivated the need for techniques that provide increased capacity using the same spectrum. This makes smart antennas a natural solution for improving wireless network performance. Offered by several commercial vendors [4, 5], smart antennas are an array of multiple antenna elements, available in different form factors. Depending on the sophistication of signal processing involved, these can be broadly classified as *multiple-input multiple-output* (MIMO) and *beamforming* antennas. In MIMO, multiple antenna elements at both the transmitter (Tx) and receiver (Rx) along with the multipath nature of the environment, help create several virtual pipes between the Tx and Rx, over which multiple independent data streams (spatial multiplexing) or dependent data streams (space-time coding) can be sent [7]. On the other hand, beamforming is a process whereby the Tx forms a beam pattern to reinforce signal reception at the Rx [7]; allowing for operation even with omni-directional clients. Using a predetermined set of fixed beam patterns is referred to as switched beamforming; forming arbitrary

beam patterns on the fly in the signal space is referred to as adaptive beamforming, whose sophistication comes at the expense of channel estimation and large feedback overhead from Rx to Tx. Given the tradeoff between complexity, performance and ease of deployment, switched beamforming provides a nice balance and is consequently popular in practical WLANs [3, 4].

Switched beamforming is realized using an antenna array and a set of beams each of which typically focuses signal energy in a specific direction. Hence they are also called directional antennas. In addition to improving the *link quality* (SNR), they also enable multiple links to operate concurrently i.e. *spatial reuse*, by suppressing energy in unwanted directions. While their benefits are well understood in outdoor environments [8–10], when it comes to indoors, conventional wisdom appears to be that the multipath nature of the environment has a detrimental effect on their link quality benefits [7]. In our earlier work [2], through experimental measurements, we showed that although the effectiveness of directional antennas is reduced indoors, they still hold promise to improve the link quality. This has prompted researchers to design practical solutions [3, 6] for leveraging their benefits indoors. However, given the growing density of wireless networks, improving spatial reuse becomes the more critical problem and existing works [3, 6, 9] assume the occurrence of spatial reuse when using directional antennas without establishing how much of it is actually available in practical indoor scenarios. Thus, it becomes critical to understand and address the following basic questions. *(i) What is the potential of directional antennas to improve spatial reuse in indoor wireless networks? (ii) Are there simple yet effective strategies that provide improved spatial reuse in indoor environments?*

To understand the potential of directional antennas for interference reduction, we perform measurements with the help of Wifi APs with eight element arrays in an indoor office. Our studies reveal the following key inferences. (1) Directional antennas *do not* hold much promise for improved spatial reuse (unlike for link quality) due to multipath scattering in the environment which cause signal leakage in unwanted directions. (2) Indoor multi-antenna channel exhibits significant channel gain variation across antennas thereby affecting directional antennas which typically split the power equally across antenna elements. (3) Selecting only a subset of the available antenna elements and splitting power across them helps strike a more efficient balance between improving the desired link quality and reducing the interference caused to other concurrent links.

The rest of the paper is organized as follows. In Section 2, we present some background on directional antennas, followed by the experimental setup and methodology in Section 3. We evaluate the potential of directional antennas for spatial reuse in Section 4. We analyze the reasons for the observed performance and propose an alternate reuse strategy in Section 5. We conclude with some discussions in Section 6.




Fig. 1. Antenna array and directional beams

2 Background

Conventional antennas are referred to as omni-directional antennas since they radiate power equally in all directions. Beamforming is an advancement in antenna technology that adjusts the shape of the radiation pattern using an array of antenna elements as shown in Figure 1(a). It is typically achieved by applying complex weights (containing a magnitude and phase) to each of the antenna elements. By applying an appropriate set of weights, the signals can be reinforced in a required direction to cause a high Signal To Noise Ratio (over an omnidirectional pattern) at the receiver in that direction. The beam patterns with such a high gain direction (main-lobe) are called directional beam patterns (e.g. Figure 1(b)). Such beam patterns also incur a spillover of energy in unwanted directions (referred to as the side-lobes). The weights can be optimized to reduce sidelobes (as illustrated in Figure 1(c)) although at the cost of reduced main lobe gain or increased beamwidth.

3 Measurement methodology

Setup: Our experimental setup, shown in Figure 2, comprises six access points and eight clients distributed in an indoor office building. Each of the access points is a commercial 802.11g device fitted with an eight element antenna array from Fidelity-Comtech [1]. Each client is a laptop running Ubuntu 8.10 equipped with a D-Link 802.11g card based on the Atheros AR5212 chipset. The nodes run Linux kernel v2.4.26 and the MadWiFi driver and their WLAN radios connect to external omnidirectional antennas with a gain of 6dBi.

Methodology: Each access point is pre-loaded with sixteen directional beams provided by the manufacturer that together span 360° . These have also been used in related work [3]. For reference, we use a fixed element (element 0) for the omni-directional pattern. Throughout the experiments, the total power transmitted from each AP is fixed at 10 dBm. All experiments are performed on




Fig. 2. Layout of Access Points and Clients.

channel 6 in the 2.4 GHz band in a building with no external Wifi interference. Using a channel scan, it is ensured that there is no extraneous interference before each measurement. Multiple time-spaced experiments are performed for increased confidence. We use Iperf to generate traffic from each AP to its client and the ‘athstats’ madwifi utility on each laptop to obtain fine grained statistics from the card. In each run, 128 Byte UDP packets are sent from the APs to the intended clients. The antenna pattern at AP_i is fixed at 1 and that at AP_j is consecutively changed from 1 to 16 after which the pattern at AP_i is changed to 2 and so on. In all, this yields a total of 256 pattern combinations for a two AP two client scenario. This is repeated for other AP pairs. The received signal strength (RSSI) at clients is the metric of interest and is a measure of the received SNR reported by the card. We compute the aggregate rate of two concurrent links, using the signal and interference powers received at each client from the desired and interfering APs respectively, along with the noise floor of the radio. We consider the default association model that is prevalent in current WLANs where a client uses the strongest AP for association. We use all unique AP-client pairs (equal to 60 after taking into account the association) for the aggregate results and a representative set of two APs, AP_1 and AP_2 with clients C_1 through C_8 , for the microscopic results.

Strategies: When considering the operation of multiple links in a single-hop wireless network, the following four strategies are possible: (1) Omni-Joint: Two links operate concurrently using omni-directional antennas with potential interference; (2) Omni-TDMA: Two links operate in a time division manner using omni antennas without interference; (3) Dir-Joint: Two links operate concurrently using one directional beam each from the available set with potential interference; (4) Dir-TDMA: Two links operate using one directional beam each in a time division manner without interference.

Metrics: In addition to aggregate rate of the links, we introduce two new metrics to characterize the spatial reuse.

i. Aggregate rate: We use both the 802.11g SINR table and the Shannon Capacity expression ($C = \log_2(1 + SINR)$) to study the aggregate rate with practical and ideal rate adaptation. The aggregate rate for *joint* strategies is the sum of rates of the concurrent links taking interference into account ($SINR$). For

TDMA strategies, it is given by the sum of rates on the individual links without interference (*SNR*) and scaled by half (due to time multiplexing).

ii. Interference power ratio: It is the factor by which the interference power at a client from an interfering AP is reduced with a directional beam compared to Omni and indirectly contributes to spatial reuse. Ideally, a highly directional beam in one direction at an AP should cause very low interference power in other beam directions. In practice, the effect of sidelobes and indoor scattering might result in several directions receiving strong interference from an AP.

iii. Spatial reuse factor: Aggregate rate by itself does not directly capture spatial reuse since it is impacted by two factors namely, link (array) gain and interference reduction (spatial reuse) gain. While the link gain improves the signal power S , spatial reuse gain comes from the reduction of interference I to jointly impact the SINR of the receiver, making it hard to isolate their contributions. However, we note that the link gain is common to both TDMA and joint strategies, while spatial reuse gain is specific to joint strategies, and hence define the following metric to better capture spatial reuse:

$$\beta = \frac{\text{Sum rate of concurrent links}}{\text{Average rate of isolated links}} \quad (1)$$

Thus, when $\beta > 1$, we have throughput benefits that are directly contributed due to interference reduction and hence spatial reuse.

4 Spatial reuse with directional antennas

4.1 Aggregate rate of directional links




Fig. 3. Aggregate rate of directional links

We first study the performance improvement that concurrent directional links provide over other competing strategies. We perform experiments over all possible two link pairs (sixty pairs) in our testbed. We plot the aggregate rate CDF results (using 802.11g SINR-rate table) for each of the four strategies: Omni-Joint,

Omni-TDMA, Dir-Joint, Dir-TDMA in Figure 4.1 along with the ideal two link rate. While directional strategies provide good throughput improvements compared to Omni strategies as expected, there are two important insights that the figure reveals: (1) Dir-Joint performs worse than Dir-TDMA for around 20% of the link pairs, indicating that joint use of links is not always good. (2) For around 40% of the cases, Dir-TDMA performs better than Dir-TDMA. However, only in 10% of the cases does Dir-Joint achieve the ideal two link rate (108 Mbps). This indicates that a only a small fraction of the gains of directional beams is contributed by interference suppression. This is also confirmed by the fact that the median gain of both Dir-Joint and Dir-TDMA is the same (54 Mbps).

4.2 Analyzing the performance degradation




Fig. 4. Interference power reduction of directional beams

Interference power reduction: To understand the reasons behind marginal spatial reuse, we focus on the interference powers at the receivers. We first present the CDF of the interference power ratio over different two link configurations. By configuration, we mean each of the two APs communicates simultaneously with one of its clients using one of 16 directional beams. In the ideal case, all the beams (except that beam which points from the interfering AP to the client) should lead to very low interference power at the client. Consequently, one would expect that most of the configurations yield an interference power ratio less than 1 (0 dB). Results in Figure 4(a) reveal the surprising observation that for more than 40 % of the configurations, the use of directional beams *increases* the amount of interference caused (compared to omni) rather than decreasing it. This clearly shows that unlike expected in an ideal free space environment, most beams cause significant interference powers to be radiated towards unintended clients. This is shown in greater detail in Figure 4(c) which plots the interference power reduction for a single two-link pair across beam configurations alone. This non-ideal interference reduction occurs due to a combination of factors such

as side lobes, scattering, varying channel gains across different elements, etc. However, the net impact is that *directional antennas do not provide the assumed interference power reduction as expected.*

We further analyze the interference power reduction assuming that an ideal beam selection algorithm is used. i.e the best combination (pair) of beams that maximizes the aggregate throughput of the two concurrent links is used. The resulting interference power for this case is plotted for link 1 and link 2 for 16 pairs of client locations in Figure 4(b). From the figure, we observe that while some of the link pairs do indeed have a reduced interference power compared to omni, most of the pairs have increased (> 0 dB) interference power. The mean interference is increased by 3 dB and reduced by -3 dB for link 1 and link 2 respectively indicating that the *interference suppression is minimal even with ideal beam adaptation.*




Fig. 5. Spatial reuse factor β

Distribution of spatial reuse factor β : For each pair of links, the signal and interference powers are measured at the clients while changing the beam pattern on the two APs chosen as described in Section 3 for the Joint and TDMA directional strategies. The CDF of β over all link pairs and beam configurations is presented in Figure 5(a). The results reveal that: (1) Unlike a free space ideal environment, where the expected improvement is $\beta = 2$, the maximum improvement is much less indoors (i.e around $\beta = 1.45$). (2) In an ideal environment where directionality is preserved, the CDF of β is expected to indicate a sharp rise near 2. However, in practice, the CDF is distributed from 0.3 to 1.45 indicating that the expected directionality and consequent interference suppression are not obtained in practice for a majority of link and beam pair configurations. (3) While the median improvement is expected to be close to $\beta = 2$ with good interference suppression, the actual median improvement is 0.9! Further, more than 70% of the configurations have $\beta < 1$, indicating that for a majority of scenarios, there is no benefit from spatial reuse with directional antennas, where TDMA using link gain performs better.

Performance with optimal beam selection: We also fix two APs, AP1 and AP2, and determine the best beam combination (pair) for concurrent operation (Dir-Joint) and best directional beams for isolated operation (Dir-TDMA) for the resulting 16 link pairs, results of which are presented in Figure 5(b). Note that this Dir-Joint is an exhaustive version of the greedy algorithm in [3] and thus upperbounds the gains from spatial reuse. The maximum value of β for the 16 link pairs is plotted in Figure 5(b). The figure indicates that the best improvement using concurrent directional beams is distributed from 0.6 to 1.45. Additionally, there are link pairs (namely 2,5,6) for which Dir-Joint performs worse than Dir-TDMA. More interestingly, the average β across link pairs is just 1.17, meaning that the contribution from spatial reuse improves performance by only around 17%.

Thus, our experiments highlight that *while directional antennas provide gains over omni in indoor environments, the benefits are contributed largely by link gain and not by spatial reuse.*

5 Alternate strategy for spatial reuse

To gain a deeper understanding, we study the indoor multiple antenna channel. Our experiments reveal that the channel gain varies drastically (over 10 dB) even for antennas which are closely spaced due to multipath fading. This drastic variation in channel gain motivates adapting the power transmitted from each antenna. This unequal gain across antennas also makes directional beams sub-optimal because directional beams are typically created by splitting the transmit power equally across the antenna elements. To verify this proposition, we consider a strategy called *antenna selection*, where the antenna element with strongest gain is alone chosen for transmission and allocated all the power.

We experiment with fifty links (each activated in isolation) and measure the Signal to Noise Ratio at the receiver for both directional and selection of the best of eight elements. We observe that for more than 80% of the locations, antenna selection *outperforms* directional beamforming. Given the highly varying nature of the channel, a strategy that uses the higher gain channels (antennas at the Tx) would lead to higher SNR at the receiver. However, in a multi-link scenario, the transmitter must also minimize the interference it causes to other receivers. Hence, we study how antenna selection performs in a two link configuration with AP1 and AP2 in our testbed. We try out all possible combinations of single element selection. i.e AP1 and AP2 select antenna elements i and j respectively, $\forall i, j \in [1, 8]$. We then compute the pair of antenna elements (one for each link) which yields the highest aggregate throughput for each link pair. The corresponding interference powers at the two clients is noted. We then compute the difference in interference power reduction compared to Dir-Joint for each link and plot the results in Figure 6(a). It can be seen that, while selection performs well from a single link standpoint, it's performance in a multi-link context is not as significant. i.e it does not provide a large reduction in interference power compared to Dir-Joint.

Link gain - Interference reduction trade-off: The previous experiments highlight the fact that the multi-link scenario is more constrained and both the desired and interfered clients must be considered when choosing the transmission beam pattern. Consequently, from a multiple antenna usage standpoint, there is a trade-off between the link (diversity gain) and the interference suppression gain. Using single element selection leverages the link to the desired client best, because it transmits all the power on the highest gain channel. However, this may not correspond to least interference to other clients. Similarly, selecting all available elements with equal power split also provides no flexibility in interference suppression and reduces array gain. On the contrary, selecting a small subset of elements with equal power split (antenna subset selection), provides the flexibility to tradeoff some array gain to choose antenna elements that also provide good interference suppression. To verify whether the trade-off occurs in




Fig. 6. Antenna selection rate improvements

practice, we evaluate the aggregate throughput improvements of both one element and two element selection. The throughput gain of one element selection compared to Dir-Joint is plotted in Figure 6(b). The figure reveals that single element selection provides throughput gains over directional for some link pairs but not all. We then plot the throughput gains of two element selection in Figure 6(c). This figure reveals two important facts: (1) The maximum throughput gain with one element selection is around 23% whereas the gain of two element selection is 53%. (2) More locations have a positive gain over directional. These results indicate that antenna subset selection is a promising approach.

Further, we evaluate the spatial reuse factor β for the single element and two element selection strategies for each of the link pairs. The results plotted in Figure 7 indicate that the improvements are more pronounced with two element selection reaching up to a factor of 1.6 (which is closer to the maximum β of 2 for two links). Additionally, the spatial reuse factor averaged over links increases from 1.19 with directional to 1.26 with single element selection and 1.42 using two element selection, thus confirming that antenna subset selection is a *better spatial reuse strategy* compared to directional.




Fig. 7. Spatial reuse factor β for antenna selection

6 Conclusion

The ability of directional antennas to reduce interference and consequently improve the spatial reuse in indoor wireless networks is impacted significantly by multipath scattering and fading. Additionally, indoor wireless channels have large gain variations across multiple antennas which motivates strategies that adapt the power transmitted on each antenna element. Selection of a subset of antennas enables such adaptation with low complexity. Intelligent algorithms to identify the right subset of antenna elements and comparison of its performance with closed loop MIMO strategies is an interesting avenue for future work.

References

1. Phocus Array, www.fidelity-comtech.com.
2. Marc Blanco *et al.*, “On the Effectiveness of Switched Beam Antennas in Indoor Environments.”, PAM 2008.
3. Xi Liu *et al.*, “DIRC: Increasing Indoor Wireless Capacity using Directional Antennas”, ACM SIGCOMM 2009.
4. Ruckus Wireless Inc. www.ruckuswireless.com
5. Cisco Inc. www.cisco.com
6. A.P.Subramanian, H. Lundgren and T. Salonidis, “Experimental Characterization of Sectorized Antennas in Dense 802.11 Wireless Mesh Networks”, ACM Mobihoc 2009.
7. A.Paulraj, R.Nabar and D.Gore, “Introduction to Space-Time Wireless Communications,” Cambridge University Press, May 2003.
8. Vishnu Navda *et al.*, “MobiSteer: Using Steerable Beam Directional Antenna for Vehicular Network Access”, ACM MobiSys 2007.
9. R.R.Choudury, X.Yang, R.Ramanathan and N. Vaidya, “On designing MAC protocols for wireless networks with directional antennas”, IEEE Transactions on Mobile Computing, Vol.5, pp.477-491, 2006.
10. S. M. Das *et al.*, “DMesh: Incorporating practical directional antennas in multi-channel wireless mesh networks”, IEEE Journal on Selected Areas in Communication, vol.24, issue-11, pp. 2028-2039, 2006.