Practical Beamforming based on RSSI Measurements using Off-the-shelf Wireless Clients

Sriram Lakshmanan, Karthik Sundaresan†, Sampath Rangarajan† and Raghupathy Sivakumar
{sriram,siva}@ece.gatech.edu, † {karthiks,sampath}@nec-labs.com

Georgia Institute of Technology
Atlanta, GA

†NEC-Laboratories America
Princeton, NJ

ABSTRACT

WLANs have become an important last-mile technology for providing internet access within homes and enterprises. In such indoor deployments, the wireless channel suffers from significant multipath scattering and fading that degrades performance. Beamforming is a smart antenna technology that adjusts the transmissions at the transmitter to reenforce the signals received through multiple paths at the receiver. However, doing this requires the accurate estimation of the channel coefficients at the receiver and its knowledge at the transmitter which off-the-shelf WiFi clients are incapable of doing. In this work, we develop a novel procedure that uses Received Signal Strength Indicator (RSSI) measurements at the receiver along with an intelligent estimation methodology at the transmitter to achieve beamforming benefits. Using experiments in an indoor office scenario with commercial WiFi clients, we show that the scheme achieves significant performance improvements across diverse scenarios.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and design**]: Wireless communication

General Terms

Experimentation, Measurement, Algorithms

Keywords

Beamforming, smart antennas, wireless link stability, throughput

1. INTRODUCTION

Indoor wireless networks operating in the 2.4 - 5 GHz spectrum have become popular last mile internet access networks using standards such as IEEE 802.11, WiMAX, etc. However, these networks are plagued by multipath scattering and fading [1] that severely undermine their performance. This is particularly significant in the context of high-bandwidth applications such as IPTV and home video distribution [2, 3]. Most works have focused on improving the link performance by coping with multipath through link layer mechanisms such as forward error correction codes, using multiple transmissions or using partially correct packets [4, 5, 6] while using a single antenna with an omni-directional radiation pattern (*Omni*). Smart antennas have emerged as an effective means to directly combat the negative effects of multipath propagation. To mitigate the effect of fading, multiple transmit antenna elements together with appropriate signal processing can be used to **a)** increase the Signal to Noise Ratio (SNR), and **b)** decrease error rate, at the receiver [1]. The ability to modify the transmitted signal in a way that signal components get *reinforced* by the multipath channel has led to the popularity of smart antennas and their adoption in several standards [7, 8, 9, 10]. More recently, several experimental works have started investigating the practical benefits of these smart antennas in indoor and outdoor wireless networks and are reporting promising results [11, 12, 13].

There are three main classes of smart antenna technologies, depending on the sophistication of signal processing at the transmitter (Tx) and receiver (Rx), namely Multiple Input Multiple Output (MIMO), Directional antennas and Beamforming antennas. MIMO is a popular strategy that uses multiple antenna elements at the Tx and Rx and leverages the rich scattering nature of the environment and the knowledge of the channel at the Rx, to increase the link capacity [1]. Directional antennas are used predominantly in outdoor scenarios and their transmission pattern is pre-set to point a main lobe (providing high gain) in the direction of the Rx [1]. Such a radiation pattern is obtained by applying a fixed set of phases to the signal transmitted from each antenna at the Tx. Most of the current works on smart antennas [12, 14, 15] fall under this category. Since the transmit pattern is oblivious to channel state information, directional beams are affected by multipath propagation and have been shown to be less effective in indoor environments [11, 16].

In contrast, *Beamforming* is a closed-loop technique that uses the channel information at the Tx to modify the transmitted signal such that the signals received through the multiple paths are reinforced at the Rx thereby improving the link Signal to Noise Ratio (SNR). Consequently, the resulting beam pattern may not have the single main lobe structure (pointing in the direction of the Rx) of a directional antenna but results in better performance in multipath rich environments. Our recent measurements in a real indoor office environment have shown that beamforming can yield an average of 6.5 dB (and maximum of 12 dB) of SNR improvement when compared to directional antennas [16]. These signal quality improvements can be used to significantly improve the datarate of the link. For instance, consider a link whose receive SNR is 8 dB when the Tx forms a directional beam to the Rx using an array of eight antenna elements. For this SNR, an 802.11g link can support a datarate of 9 Mbps [17]. On the other hand, when us-

IMC'09, November 4–6, 2009, Chicago, Illinois, USA.

ing beamforming, the additional SNR gain of 6 dB [16] results in a datarate of 18 Mbps (yielding a 2X rate improvement compared to directional). Thus, the importance of beamforming over directional beam is clear. Unlike other MIMO techniques, beamforming does not require multiple elements at the clients and hence conventional off-the-shelf single element antenna clients[1] could be used to leverage this technique. Further, beamforming is complementary to other MIMO techniques and could also be used with multi-antenna clients.

Background: Given the importance of beamforming, there has been significant research in theory which suggests that with accurate channel estimation at the Rx and its feedback to the Tx, one could exploit the multipath channel optimally to improve link throughput. When a transmitter with K antenna elements communicates with a receiver which has a single antenna element, the wireless channel so formed is called a Multiple Input Single Output (MISO) channel. The baseband channel model for a MISO channel can be represented as

$$y = \mathbf{h}^T\mathbf{x} + z \tag{1}$$

where the column vector $\mathbf{h} = [h_1 h_2 \ldots h_K]^T$ is the vector of channel gains between each Tx antenna and the Rx antenna, $\mathbf{x}$ is the $K \times 1$ vector of the transmitted signals, $\mathbf{y}$ is the received signal and z is the additive White Gaussian noise. A beamformer is defined as a vector of complex numbers $\mathbf{w} = [w_1\, w_2 ...\, w_K]$ which translates each transmit symbol s to the signal vector $\mathbf{x} = \mathbf{w}s$ to be transmitted from the K antennas. By estimating the channel at the Rx and conveying it to the Tx, the beamformer can be adjusted to be the complex conjugate of the channel, so that the spatial channels from each Tx antenna combine coherently and reenforce each other at the Rx. Thus, the optimal beamformer is given by $\mathbf{w} = \mathbf{h}^*$, where $*$ denotes the complex conjugate. We note that given the channel vector $\mathbf{h}$, the weights can be computed in a straightforward manner. Applying a complex weight to each antenna element adjusts the magnitude and phase of the signal transmitted from that element appropriately.

Channel Estimation: Beamforming requires channel estimation between each Tx and Rx antenna, each of which is a complex value which is composed of a magnitude and a phase component. A conventional channel estimation procedure involves activating each antenna in isolation and transmitting a fixed sequence of bits (preamble) on each antenna. This procedure is repeated for each of the antennas [1]. The received complex baseband symbols (i.e., the amplitude and phase) are then used to determine the complex channel gain h_i using the structure of the preamble. The accuracy of the received symbol value for all the preamble bits determines the accuracy of the channel estimates. While the magnitudes of the channel gains can be obtained easily, estimating the phase accurately is more involved as described next.

Motivation: Practical channel estimation has the following problems, **a)** it requires specialized clients with the capability to measure the amplitude and phase of the received signals, **b)** even with such clients, hardware synchronization impairments (frequency and phase offsets) tend to corrupt the estimated channel coefficients [16], and **c)** the feedback overhead and processing delays are very high making it hard to deploy. In [16], we presented an intelligent channel estimation scheme that performs measurements in *space* rather than time to estimate the channel while overcoming hardware synchronization impairments. Even though beamforming using such an accurate channel estimation procedure is shown to provide much gain over Omni and directional beams, it requires specialized clients which are capable of symbol level channel estimation (such as the software radio used in [16]). However, such capability is not currently available on commercial WiFi clients. Thus accurate channel estimation is a burden to achieving beamforming benefits. Given the above practical impediments in procuring channel state information at the Tx, the key question to ask is: *Can we develop a procedure that enables one to realize the benefits of beamforming using off-the-shelf clients?*

Our contribution: We answer the above question in the affirmative. We provide a new beamforming solution that performs approximate channel estimation using signal power measurements at the Rx in conjunction with an intelligent antenna activation algorithm at the Tx. This solution still uses the concept of differential phase estimation as in [16] but now computes the complex beamformer weights using only the received signal power estimates (as opposed to per symbol amplitude estimates). Consequently, (i) As an approximation to received power, RSSI measurement can be used to compute beamformer weights, (ii) hardware oscillator impairments are overcome by the use of differential phases, (iii) significant benefits are obtained without hardware modifications to clients, and (iv) the feedback overhead is low. Overall, the solution provides a better balance of the performance-complexity tradeoff when compared to both directional and beamforming based on accurate channel estimates.

We validate the practical feasibility of the algorithm by implementing it on off-the-shelf WiFi clients. Using the developed system, we conduct measurements in an office environment to evaluate the approximate beamforming solution using an eight-element phased array AP and off-the-shelf single antenna omni-directional clients. A summary of the results is as follows.

- The proposed beamforming solution reduces the multipath fading induced packet loss rate from 6.5% to less than 0.3% in a real indoor environment and improves throughput by up to 7.1X and 1.64X compared to Omni and directional.

- Beamforming provides a median SNR gain of 10.5dB over Omni and 3dB over directional antennas which is close to the anticipated median gains of 11.5dB and 6dB when using perfect channel estimates.

- Compared to Omni, beamforming also improves the connectivity by reducing coverage holes and enhances the link stability significantly.

The rest of the paper is organized as follows. In Section 2, we present our beamforming solution that uses RSSI measurements. Section 3 describes the results from an evaluation study of the proposed scheme. Conclusions and future directions are presented in Section 5.

2. BEAMFORMING USING RSSI MEASUREMENTS

The beamforming procedure consist of two stages: (1) The *transmit activation* stage and (2) the *estimation and feedback* stage as illustrated in Figure 1. Conventional beamforming uses a simple activation stage where the AP transmits a known sequence on each antenna followed by a symbol level estimation and feedback. Since commercial WiFi clients cannot provide symbol level estimates, we design a more intelligent activation stage where multiple antennas

[1]It should be noted that single element antennas can only produce omni-directional radiation patterns

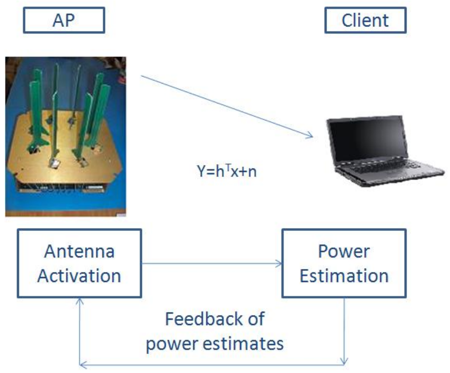

Figure 1: **Beamforming AP and client. The Beamforming procedure is highlighted.**

are activated in tandem followed by a simpler channel estimation stage where signal power values are fed-back from the clients. Note that RSSI is a good approximation of the received power [2], thereby enabling the procedure to work with Off-the-Shelf clients.

Core idea: The algorithm is based on the idea of estimating differential channel phases by employing *tandem activation of more than one antenna* and using received power estimates. Thus, the estimation process is distributed across space (elements) instead of time. In conventional channel estimation, when a single antenna is activated at a time, the received power is dependent only on the channel magnitude and is given by $P_i = |h_i|^2$ (assuming the Tx power is unity). Hence the information about the channel phase $arg(h_i)$ is lost when the power is computed. In contrast, by the tandem activation of more than one antenna element, the effects of the channel phases are also reflected in the received power in a manner that depends on the relative channel phases. i.e. when two elements i and j are activated simultaneously with equal weights (such that the transmitted power still adds up to unity), the received power can be computed as $P_{ij} = |h_i + h_j|^2$. Thus, for tandem activation, the received power P_{ij} is given as

$$P_{ij} = P_i + P_j + 2\sqrt{P_i P_j}cos(\theta_{ij}) \tag{2}$$

where θ_{ij} is the *channel phase difference* between h_i and h_j. Depending on the relative channel phase θ_{ij}, the two signals combine together to change the signal power at the Rx. When $\theta_{ij} = 0$, the signals combine constructively causing the powers of the individual elements to add up at the Rx. However, when $\theta_{ij} = 180$ the signals combine destructively causing the received power to be the difference of the powers transmitted from the individual antennas. Hence, the change in the received power across a strategic set of activations can be used to identify the relative channel phase between the channel gains by rewriting Equation 2 as

$$\theta_{ij} = \cos^{-1} \frac{P_{ij} - P_i - P_j}{2\sqrt{P_i P_j}} \tag{3}$$

By repeating this idea for pairs of antenna elements, the relative phases can be obtained. Since all the channel phases must be measured with respect to the same reference for estimates to be meaningful, we designate element 1 as the reference element. The channel gain magnitudes can be obtained directly from the power measurements by activating each antenna element individually as $|h_i| = \sqrt{P_i}$. When used along with the relative phases, the beamformer weights can be determined as $w_i = \sqrt{P_i}e^{j\theta_{i1}}$ for $i > 1$ and $j = \sqrt{-1}$ with $\theta_{11} = 0$. We also note that, irrespective of the number of antennas used, we ensure that the total transmitted power remains constant by normalizing the weights. In the rest of this section, we describe the algorithm to identify optimal beamforming weights.

[2]In the rest of the paper we use power synonymously with RSSI for ease of explanation

2.1 Algorithm steps

The algorithm consists of the following steps and is performed at the Tx and Rx consecutively.

1. Single and tandem activation with equal weights:

In the single antenna activation stage, each one of the K elements at the Tx is activated in isolation, i.e. one at a time using S consecutive packets for each antenna element. S is a parameter that can be increased for more accurate estimates but is chosen to be small to keep the overhead of the estimation process low (we use $S = 5$ in our experiments). S becomes specially important to perform right ambiguity resolution. This is followed by activating two antenna elements at a time. One of the two antennas in each activation is the reference antenna element and the other is chosen successively from second to the K^{th} antenna.

2.RSSI measurement and computation of channel magnitudes and phases:

The $K-1$ received signal power values for each of the tandem activations is noted at the Rx along with the K average signal powers for the single activations. These $2K-1$ values are then used to compute the magnitudes $|h_i| = \sqrt{P_i}, 1 \leq i \leq K$ and the relative phases $\phi_{i1}, 1 < i \leq K$ from Equation 3. The K magnitudes $|h_i|$ and the $K-1$ relative phases are then conveyed to the AP in a single packet.

3. Ambiguity resolution through tandem activation with unequal weights:

While the magnitudes are obtained correctly, the phases ϕ_{ij} (in radians) have an ambiguity due to the use of the cos^{-1} function in Equation 3. i.e. the correct θ_{ij} can be either of $\phi_{ij}, -\phi_{ij}, \pi - \phi_{ij}, -(\pi - \phi_{ij})$. To resolve the ambiguity, the $k-1$ pairs activated in tandem in Stage 1 are again activated but with modified amplitude and phase weights[3]. i.e., element 1 is activated using the magnitude $\sqrt{\frac{P_1}{P_1+P_i}}$ and phase '0', whereas element $i, i > 1$ is activated with a magnitude $\sqrt{\frac{P_i}{P_1+P_i}}$ and each of the phases $\phi_{ij}, -\phi_{ij}, \pi - \phi_{ij}, -(\pi - \phi_{ij})$. Hence, for each of the $K-1$ pairs, there are four activations corresponding to these four phases, which we call the quadruple.

4. Accurate beam weight determination:

Of the four choices in each quadruple, the receiver identifies the choice which yields the largest signal strength at the receiver and notes this as the unambiguous relative phase for each of the non-reference antenna elements i.e. element 2 to element K. The final beamforming weights for each antenna element i is given by the magnitude $|w_i| = \sqrt{\frac{P_i}{\sum_{l=1}^{K} P_l}}$ and the phase θ_{i1}.

2.2 Performance and robustness:

Computing the beamforming weights using the differential phases (as described above) yields the same SNR improvement as using absolute channel estimates as described below.

With ideal beamforming, the channel vector $\mathbf{h}$ is measured accurately in both magnitude and phase. Let the magnitude be given as $a_i = |h_i|$ and the phase as $b_i = arg(h_i)$. Hence, the beamforming

[3]The relative magnitudes are chosen such that it is the same as what would eventually be used by the beamformer; only ambiguity in phase is being resolved at this point.

weights are given by $w_i = \frac{a_i e^{-jb_i}}{\sqrt{\sum_{i=1}^{K} a_i^2}}$. When used in Equation 1, the received signal (ignoring noise) is then given by

$$y_{ideal} = \sum_{i=1}^{K} a_i e^{jb_i} \frac{a_i e^{-jb_i}}{\sqrt{\sum_{i=1}^{K} a_i^2}} \tag{4}$$

Consequently, the received power is

$$P_{ideal} = \sum_{i=1}^{K} a_i^2 \tag{5}$$

Using the algorithm proposed in this section, the magnitudes are given by $\sqrt{P_i} = a_i$ and are the same as in the conventional beamforming case. On the other hand, the (relative) phases are given by $c_1 = 0$ and $c_i = \theta_{i1}, i > 1$, which can be simplified as $c_i = b_i - b_1$. hence, the beamforming weights in this case are, $w_i = \frac{a_i e^{-jc_i}}{\sqrt{\sum_{i=1}^{K} a_i^2}}$. When applied to Equation 1, the received signal is given by

$$y_{rssi} = \sum_{i=1}^{K} a_i e^{jb_i} \frac{a_i e^{-j(b_i - b_1)}}{\sqrt{\sum_{i=1}^{K} a_i^2}} \tag{6}$$

which can be simplified as

$$y_{rssi} = e^{jb_1} \sqrt{\sum_{i=1}^{K} a_i^2} \tag{7}$$

Consequently, the received power is given by,

$$P_{rssi} = \sum_{i=1}^{K} a_i^2 \tag{8}$$

From Equations 5 and 8, it is easy to see that the proposed algorithm which uses relative phases provides the same beamforming gain as the ideal beamforming algorithm which computes the exact phases. The underlying intuition is that, a common phase error across the estimates of $\mathbf{h}$ does not affect the relative aligning of the spatial channels. The channels combine coherently, as long as the phases relative to a common reference are accurate.

We also note that the proposed algorithm is robust to the oscillator phase and frequency offsets identified in [16] since it uses signal power measurements which are unaffected by the phase and frequency offsets (unlike symbol level measurements).

2.3 Overheads

In this section, we quantify the overheads incurred by the above procedure in terms of the number of excitations at the transmitter and the number of feedback bits. With K antennas at the Tx, the total number of excitations for the estimation in the first stage is $2K - 1$. The ambiguity resolution stage requires $4K - 4$ excitations, making the total number of excitations $6K - 5$. With S samples per excitations, the total number of excitations is $(6K - 5)S$. However, the number of receive observations that must be fed back is $3K - 2$. With B bits to represent each symbol [4], the total number of bits to be fedback is $(3K - 2)B$.

In practice, each excitation requires only a few samples(symbols) to obtain reliable estimates (for e.g. 10). Hence, with $K = 8$ antennas, the total number of symbols for the first stage is $15 * 10 = 150$, whereas it is $28 * 10 = 280$ for the ambiguity resolution stage. Each of these excitations can be accomplished using one Wifi control packets of size < 40 bytes. The total feedback overhead from the client to the AP is limited to 44 bytes (i.e $22 * 16$ bits) which is less than a normal 802.11 packet size. When an averaging value $S = 5$ is used the feedback packet size required is $44 * 5 = 220$ bytes, which is within the Maximum Transmission Unit (MTU) of Wifi links. However, our current hardware (Phocus Array v2.1 [18]) does not allow fast changing of patterns within a packet. Hence, for the implementation, each excitation takes one packet duration. But with a more sophisticated implementation, we believe that the excitation process can be accomplished much faster.

[4]For instance, the number of bits used to represent each baseband sample magnitude in the USRP is 16.

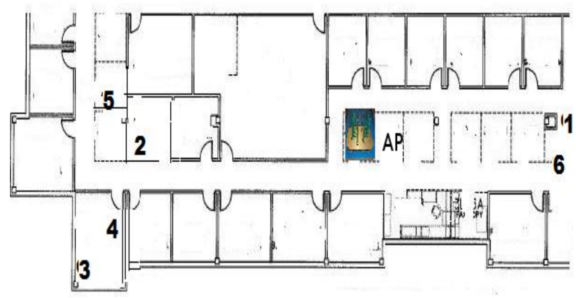

Figure 2: Experimental Testbed. The numbers indicate client locations.

The effect of the above overhead depends on the frequency of adaptation, which depends on the frequency of channel variation. Additionally, as explained in [16], for static clients *a beam pattern can be estimated and retained for several tens of seconds (thousands of packets) without incurring appreciable degradation in performance.* When considering mobile clients, the frequency of adaptation is higher. Nevertheless, pedestrian mobility is typically low (< 1 m/s), and the overhead of adaptation is just two packets. Hence we expect that the overheads are small compared to the benefits even for mobile clients. However, exploring beamforming solutions optimized for mobility is an interesting avenue for future work.

3. EVALUATION RESULTS

Experimental setup: To implement beamforming, we use a testbed consisting of an 802.11b/g access point with a phased array antenna from Fidelity-Comtech [18] and a laptop running Ubuntu 8.10 equipped with a D-Link 802.11b/g card in an indoor office environment as shown in Figure 1. The AP runs the open source Madwifi WLAN driver. It also consists of a set of 16 pre-computed directional antenna patterns that cover the entire 360 degrees and a command interface to set and write new beam patterns. In the AP, the pattern that is selected is used for both transmission and reception. The D-Link card uses the Atheros chipset and the Madwifi driver. We use Iperf as the traffic generating application and the *athstats* Madwifi utility on the laptop to obtain fine grained statistics from the card. The AP is placed at a fixed location on top of a cubicle and the laptop is placed at six different locations throughout the building as shown in Figure 2. The transmit power of the AP is fixed at 10dBm for all experiments. We use two main metrics: (1) the SNR (computed from the RSSI and the Noise Floor reported by the card) and (2) the throughput reported by the Iperf application on the receiver. We mainly highlight the benefits of the beamforming solution (Bf) in comparison to *Omni*(where the transmitter uses a single antenna element) and *Dir* where a directional pattern that points to the Rx is used. All experiments are conducted with 1500 byte packets by default and RTS/CTS is not used.

Solution Implementation: The proposed solution is implemented as scripts on the AP and client. The AP code consists of shell scripts which invoke the appropriate commands for writing patterns into the AP's hardware and setting the correct pattern numbers for each experiment. At the client side, the *athstats* Madwifi utility is used to collect the statistics required for beamforming. The *athstats* utility allows low level WiFi card statistics such as the number of Cyclic Redundancy Check (CRC) errors, the number of retransmissions, etc. to be collected. Both the client and the AP implement their repsective default rate control algorithms. The *athstats* utility is modified to take an input parameter T (in seconds, set to 5 by default), that determines the duration for which it is run. The output of *athstats* is written to a file and parsed to obtain the average of the individual received RSSI for each experiment. These values are processed using Octave scripts to identify the beamforming weights. Once identified, a new pattern file is generated and transferred to the AP. The AP provides a command line interface to write a new pattern file into the hardware. Similarly the 'pasctl' utility is used to set the right pattern. The AP starts Step 1 of the algorithm described in Section 2 followed by the client providing feedback. Then the AP again goes through the activation sequence described in Step 3 and the client returns the feedback in Step 4. The signal power is used for computing the weights in the algorithm and is obtained from the RSSI reported by the card in dBm. All experiments are carried out on Channel 6 and before each experiment it is ensured that there are no other co-channel transmitters in the test environment. The feedback from the client to the AP is accomplished through the wireless interface through an 'SCP' file transfer.

Power measurement: Current 802.11 cards only present the received power in units of dBm and quantized to take integer values of dBm as opposed to absolute powers in milliwatts. The use of quantized dBm values results in some inaccuracy in computing the beam weights. Specifically, a 1 dB error in dBm units can affect the absolute value by up to 25%. Hence, we compute the average of multiple RSSI samples after converting them from dB scale to absolute scale to improve the accuracy. We show later that although this inaccuracy creates a loss in beamforming performance compared to beamforming based on accurate channel estimation, still significant beamforming benefits are obtained.

3.1 How well does beamforming work?

We first evaluate whether the proposed solution using RSSI works well compared to directional beams and if so what the magnitude of improvement is. We vary the position of the client among the positions in Figure 2 and profile the performance of Omni, directional and beamformed links. We use Iperf to transmit UDP traffic to the client for one minute in the following scenarios. We first create an omni pattern at the AP using a single antenna (Antenna 1). Second, we change the antenna pattern to a directional antenna pointed towards the client. Third, we run the beamforming algorithm described in Section 2, compute the beamforming weights and transmit antenna patterns using these weights. We then measure the throughput and SNR at the Rx in each of these cases. At the end of this run, we once again measure the performance for Omni, directional and beamforming to ensure that the channel has not changed during the measurement. In most of the cases, there was no change in performance at the beginning and end of the measurement run which is in agreement with our channel measurement results in [16]. The results are plotted in Figures 3 and 3 for each of the six different locations. It is seen that beamforming performs better than directional and Omni at each of these locations.

The average improvement of beamforming over Omni is 10.5 dB in SNR and 7.1X in throughput for the locations profiled. Similarly, when compared to directional antennas, the average gain is 3 dB in SNR and 1.64x in throughput with a best case gain in SNR of up to 4 dB. Hence, we confirm that the proposed solution that uses RSSI feedback can be used to achieve significant beamforming gains. Note that the average SNR gains reported in this work are lower than that achieved with ideal adaptive beamforming performed using symbol level channel estimates (which is around 6.5 dB) [16]. We identify the following reasons for this gain reduction: (1) the approximation error introduced by using RSSI values which are quantized to integral values in dB instead of absolute signal power, and (2) the experiments in [16] do not use an Automatic Gain control(AGC) circuitry which is implemented in commercial systems; the AGC circuitry tends to decrease the beamforming gains. In spite of this, we note that the proposed solution gives significant beamforming benefits while reducing the complexity of beamforming. We leave the fine-tuning of the algorithm to future work.

Table 1: Connectivity when using Omni and beamforming. PING connectivity is improved using beamforming

Runs	Omni		beamformed	
	Loss(%)	Latency(ms)	Loss(%)	Latency(ms)
1	100	> 1000	0	2
2	96	20.4	0	1.6
3	76	12.1	0	1.55
4	76	6.3	0	1.3
5	80	15	0	1.9

3.2 Link Stability benefits

We focus on the benefit that beamforming brings in overcoming the losses that occur due to multipath fading. Figure 3 plots the error performance on the y-axis for a fifteen second snapshot of a UDP traffic experiment in a single location (Location 1). The WiFi card increments a counter whenever a packet header is successfully decoded but the CRC for the MAC layer payload failed. The *athstats* utility prints out the number of packets received and the number of packets with failed CRCs for every one second. The ratio of the number of CRC errors to the total number of packets whose header was decoded represents the error rate on the channel due to random (multi-path fading) related wireless loss (recall that there are no other co-channel transmitters in the test area thereby ruling out hidden terminals and the SNR from the AP to the clients is also above the association threshold).

From the figure it is clear that while the (link level) packet loss rate seen with a single-element antenna transmission is around an average value of 6.5%, the use of beamforming reduces the losses to less than 0.4% for the same transmitted power. We highlight that the impact is higher when considering higher layer protocols such as TCP which are sensitive to losses and delay. Thus, we conclude that the beamforming solution effectively handles multi-path fading related channel impairments. Consequently, the (UDP) throughput of the link remains stable without fluctuations to channel fading as can be observed in Figure 3.1.

3.3 Connectivity and coverage improvements

We investigate whether the beamforming solution improves the connectivity of clients and reduces coverage holes. To do this, we pick specific spots (near locations 3 and 5 in Figure 2), that are not within connectivity range when using an Omni antenna at the Tx. In these locations, the client cannot associate with the AP since the received power is below the association threshold. We then run

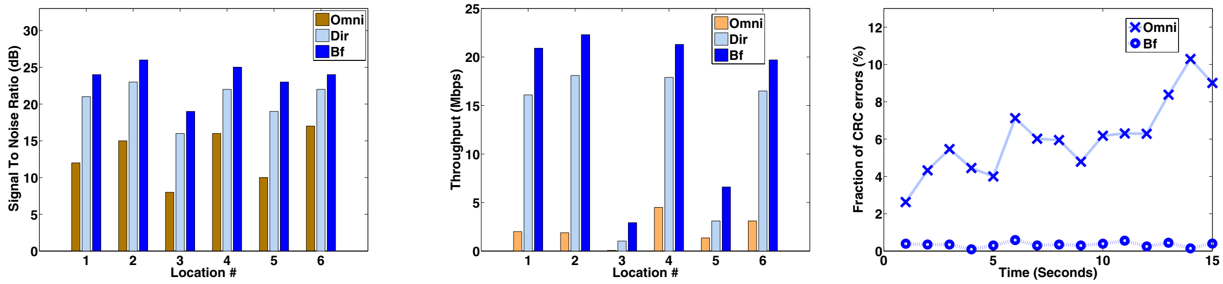

(a) RSSI gain. (b) Throughput (c) CRC Errors

Figure 3: Beamforming improves link quality

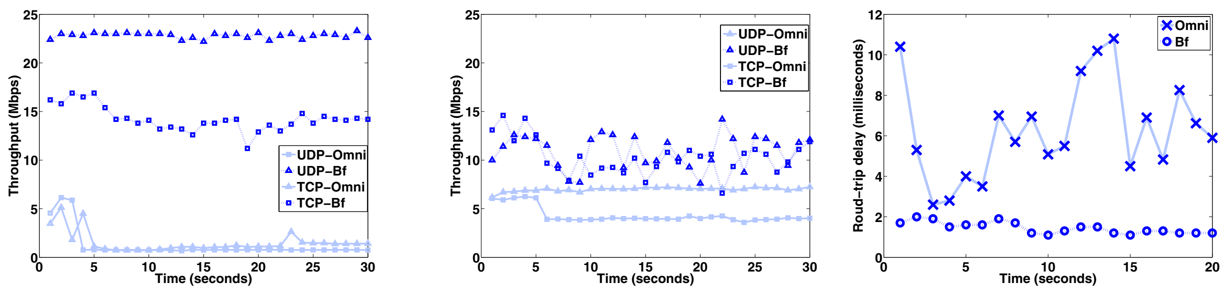

(a) Downlink throughput (b) Uplink throughput (c) Delay performance

Figure 4: Transport layer performance

the beamforming algorithm at the Tx. We observe that while the client cannot associate when using a single antenna element, the use of beamforming improves the signal strength and consequently the connectivity. We send a sequence of ICMP PING packets from the AP to the client, to determine the connectivity. We show the latency of successful ping responses and loss rate in Table 1. As can be observed, for several runs spaced apart by 10 minutes, both the average loss rate and latency of the ping response are very high when using Omni. We tabulate the delay for those packets that are successful. In some runs (Run 1, 2), the delays are very high and sporadic, because there is no connectivity at those locations. However, using the beamforming solution causes successful responses with an average latency of 1.67ms and zero loss.

3.4 Transport layer performance

In this part, we are interested in evaluating how the benefits of beamforming translate to application layer performance. We consider the following two dimensions: Uplink vs. downlink and TCP vs. UDP, which are the representative transport characteristics for several wireless LAN applications. For this experiment, we use a fixed location (location 2) and plot the throughput obtained. For the downlink case, we execute a 30-second run using Iperf from the AP to the client for the Omni and beamforming cases. The resulting throughput is plotted in Figure 3.1.

Downlink: From the figure, we observe that the downlink UDP throughput increases from an average of 0.9 Mbps to 22.1 Mbps, whereas the TCP throughput increases from 1.06 Mbps to 14.2 Mbps when beamforming is applied. Thus there is significant improvement for both transport protocols, but *the gain is higher with UDP, due to the adverse reaction of TCP to even minimal losses*.

Uplink: On the other hand, the benefits are much smaller for the Uplink case as reported in Figure 3.1. We recall that the dominant factor as far as the throughput is concerned is the rate at which the MAC layer DATA packet is sent, since the ACK packet is always sent at the basic rate of the 802.11 standard (1 Mbps for 802.11b and 6 Mbps for 802.11g). While *the downlink can directly improve the data rate of the 802.11 DATA transmissions, the beamforming gains for the uplink case can only reduce the error in the packet reception* but cannot adjust the datarate (as the client uses a single element). Since the MAC layer ACK is always sent at the base rate of 6Mbps it cannot be used to improve the datarate of the ACK transmissions, thereby leaving the beamforming gains unused. Thus, it is clear that the magnitude of the beamforming benefits depends on the traffic characteristics.

Delay: In addition, we also study how the link delay changes with beamforming. We plot the packet round trip delay in Figure 3.1 for location 2, when Omni and beamforming are used. It is clear that even for locations which have connectivity, the average delay is reduced from 6.5ms to 1.3ms when beamforming is employed. In addition, the standard deviation in delay (which is important to video applications) is reduced from 2.49 to 0.28 ms.

4. DISCUSSION

MIMO Clients: In this work we have considered clients with omni-directional antennas since the majority of current deployed clients have a single (omni-directional) antenna. However, recently multiple antennas at the clients(e.g. the IEEE 802.11n standard) are also becoming popular. With such clients, adaptive beamforming can be used along with other approaches such as Spatial Multiplexing when the latter approaches suffer in performance. More importantly, when performing multi-link MIMO, the channels between each transmit antenna and each receive antenna must be measured in order to perform interference suppression.Thus, even with other strategies that are possible with MIMO, channel estimation developed in this work is important. A simple extension of the current procedure to such cases would involve performing this procedure for each receive antenna, one at a time. Exploring more optimized variants of the current procedure is an interesting direction for future work.

5. CONCLUSIONS

In this work, we develop a new beamforming algorithm that approximates the channel estimation procedure by using power measurements. It is shown that RSSI can be used as an approximation for the received power and used to compute the channel coefficients, provided that an intelligent estimation technique is used at the Tx. Using an experimental deployment, beamforming using such channel estimation is shown to achieve significant benefits without incurring excessive complexity; thus it becomes practical to implement and achieve beamforming gains using off-the-shelf wireless clients.The ability to exploit the multipath channel through beamforming with simple RSSI measurements at the client has opened up several interesting directions for exploration, such as: **a)** investigating if the beamforming solution can be extended for simultaneously serving multiple users from a single AP, and **b)** exploring if such a beamforming solution can be used for improved spatial reuse and better interference management across multiple APs.

6. REFERENCES

[1] A. Paulraj, R. Nabar, and D. Gore, “Introduction to space-time wireless communications,” *Cambridge University Press*, May 2003.

[2] ITU-IPTV focus group. [Online]. Available: http://www.itu.int/ITU-T/IPTV/

[3] Ruckus wireless inc. [Online]. Available: www.ruckuswireless.com/solutions/triple-play-carrier-services

[4] J. Proakis and M. Salehi, *Digital Communications*. McGraw-Hill Science/Engineering/Math, 2007.

[5] A. Miu, H. Balakrishnan, and C. E. Koksal, “Improving loss resilience with multi-radio diversity in wireless networks,” in *ACM MOBICOM*, 2005.

[6] K. Jamieson and H. Balakrishnan, “PPR: Partial packet recovery for wireless networks,” in *ACM SIGCOMM*, 2007.

[7] Wimax forum, http://www.wimaxforum.org.

[8] The long term evolution project, http://www.3gpp.org.

[9] Wlan standard for enhancements for higher throughput, http://www.ieee802.org/11/reports/tgn_update.htm.

[10] WirelessHD,http://www.wirelesshd.org.

[11] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan, “On the effectiveness of switched beam antennas in indoor environments,” in *Passive and Active Measurements Conference (PAM)*, Apr 2008.

[12] K. Ramachandran, R. Kokku, K. Sundaresan, M. Gruteser, and S. Rangarajan, “R2D2: Regulating beam shape and rate as directionality meets diversity,” in *ACM MOBISYS*, 2009.

[13] V. Shrivastava, S. Rayanchu, J. Yoon, and S. Banerjee, “802.11n under the microscope,” in *ACM/USENIX Internet Measurement Conference (IMC)*, 2008.

[14] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. T. Giel, and S. R. Das, “Mobisteer: using steerable beam directional antenna for vehicular network access,” in *ACM MOBISYS*, 2007.

[15] A. P. Subramanian, H. Lundgren, and T. Salonidis, “Experimental Characterization of sectorized antennas in dense 802.11 wireless mesh networks,” in *in ACM Mobihoc*, 2009.

[16] S. Lakshmanan, K. Sundaresan, R.Kokku, A.Khojestapour and S.Rangarajan, “Towards adaptive beamforming in indoor wireless networks: An experimental approach,” in *IEEE Infocom (Miniconference)*, Apr 2009.

[17] K. Ramachandran, R. Kokku, H. Zhang, and M. Gruteser, “Symphony: Synchronous two-phase rate and power control in 802.11 wlans,” in *ACM MOBISYS*, 2008.

[18] Fidelity-comtech inc, http://www.fidelity-comtech.com.