Practical Multi-antenna Spatial Reuse in WLANs

Sriram Lakshmanan[1], Karthik Sundaresan[2],
Mohammad Khojastepour[2], and Sampath Rangarajan[2]

[1] School of ECE, Georgia Institute of Technology, Atlanta, GA
[2] NEC Laboratories America Inc., Princeton, NJ

Abstract. Smart antennas can improve spatial reuse in a wireless network through interference suppression. However, interference suppression requires support from clients in the form of channel estimation, which existing clients do not support. In this work, we explore practical solutions to obtain spatial reuse with smart antennas without requiring hardware changes to clients. To this end, we design a novel solution for improving spatial reuse in indoor WLANs which uses 'approximate' channel estimates and still yields close to ideal performance. Our solution called Light-weight Multi-Antenna Spatial Reuse (*LSR*) consists of (i) a multi-link channel estimation scheme that can be realized with simple *Received Signal Strength (RSSI)* measurements that existing WLAN clients provide readily, (ii) a low-complexity scheduler to decide the subset of beamformed links that must be active concurrently. We demonstrate that the estimates obtained using this scheme when used with a multi-link beamforming technique such as Zero Forcing yields significant interference suppression benefits. We implement the channel estimation scheme on a testbed of software radio clients to demonstrate its practical feasibility. Further, we evaluate the performance of *LSR* using extensive signal strength traces from 802.11g Access Points equipped with eight element antenna arrays in an indoor office environment. The results indicate that *LSR* achieves close to the performance obtained with an optimal scheme that uses accurate channel estimates and also improves the median sum rate of indoor users by up to 2.7x over competing approaches.

Keywords: smart antennas, spatial reuse, interference suppression.

1 Introduction

The growing density of wireless networks and limited spectrum availability are posing severe challenges to achieving high performance in wireless networks. Consequently, significant attention is being devoted to improving the number of successful concurrent transmissions in a given network area (also called spatial reuse). Due to their ability to control signal transmissions spatially, smart antennas have emerged as promising candidates to achieve high link throughput and to improve the spatial reuse by interference suppression.

I. Tomkos et al. (Eds.): BROADNETS 2010, LNICST 66, pp. 158–177, 2012.

Although smart antennas have been shown to significantly improve the spatial reuse (by interference suppression) in theory [1], relatively few works have focused on whether the promised gains are achievable in practice or the design of practical solutions to realize those gains. Interference suppression algorithms can be classified into two main groups, namely coordinated and uncoordinated. Coordinated beamforming algorithms [2], typically compute the best set of weight vectors in a coordinated iterative manner. These algorithms yield the best performance but are also the most complex. On the other hand, with uncoordinated beamforming algorithms, the beamformers are computed at each AP in isolation. A popular instantiation of this class is interference nulling through Zero-Forcing (ZF).

Interference suppression between smart antenna links requires the estimation of the channel both from the desired and interfering transmitters (i.e. APs) at each of the receivers (i.e. clients). When the estimates are fed back to the APs, a joint processing helps determine how each individual AP should adapt its transmissions to maximize the signal to the desired client while minimizing the interference to other clients. Thus, accurate channel estimation at the clients is critical to realize interference suppression. However accurate channel estimates are not readily available in legacy omni-directional clients. Further, even with multiple antenna clients, they are available only at the physical layer of a single link (e.g. 802.11n), but not at the higher layers where the decisions about joint beamforming or multi-link scheduling must be taken. Additionally, practical networks must accommodate a variety of client platforms (smart phones, VOIP phones, laptops, etc) from several different manufacturers and all of them may not provide the required information for taking medium access decisions. Given the above challenges, the key questions to be answered are: (i) *Can channel estimation be realized without requiring hardware changes at the clients?* and (ii) *Can we enable interference suppression and achieve spatial reuse in practical networks which comprise such clients?*

We establish that the answer to both the above questions is affirmative but relies on appropriate algorithm design for channel estimation and link scheduling. We showed in our previous work [3] that single-link channel estimation and beamforming is possible even with conventional clients using just received power measurements and no hardware/software changes [1]. However, in this work we first show that a straight-forward extension of that approach has poor accuracy and scalability when applied in a multi-link scenario. We design several mechanisms to ensure that the channel estimation using just RSSI measurements is both *accurate and scalable* in a multi-link scenario. Given accurate channel estimates, the next step toward realizing network level gains of interference suppression is link scheduling. While the large potential set of links in a dense AP deployment makes this problem especially challenging, we propose a simple and light-weight scheduler that selects the right subset of links to be concurrently active so that the aggregate network throughput is maximized. We show that the

[1] Received Signal Strength Indicator (RSSI) is readily provided by all cards without additional software requirements.

proposed scheduler is scalable. i.e. it provides benefits close to an optimal solution involving brute force search over all combination of links without incurring the exponential complexity with increasing number of links. We call our solution that integrates the above algorithmic components Light-weight Multi-Antenna Spatial Reuse (LSR). LSR performs both *accurate channel estimation and link scheduling without requiring changes at the clients* and keeps the overheads low compared to competitive approaches.

To evaluate the practical benefits of LSR, we develop a testbed of six 802.11g APs equipped with eight element antenna arrays and six software radio clients in an indoor office environment. We implement the multi-link channel estimation on the testbed and identify that the beamforming vectors computed with RSSI measurements are very close to those obtained with ideal channel measurements (with 95% correlation coefficient for over 90% of the locations profiled). Although the channel estimation itself was successfully implemented, we identified that realizing the complete 'interference suppression' gains involves addressing new practical challenges such as the inability of the hardware to change patterns fast and the power leakage from the transmit antennas. Despite these limitations, We evaluate the real potential of interference suppression, by collecting a large set of signal strength traces from our testbed. Our evaluation indicates that, LSR improves the median sum rate by up to 2.7x compared to single user beamforming strategies and by up to 2.3x compared to concurrent interference-unaware beamforming. *Interestingly, we observe that the median number of concurrent links to be scheduled in indoor WLANs to maximize network throughput with eight antenna APs is around 4.* This observation has significant system design implications and keeps the run-time complexity of LSR small for a significant fraction of the topologies.

The rest of this paper is organized as follows. Section 2 provides a brief background on beamforming with interference suppression and also discusses related work in this area. Section 3 describes the solution components and also summarizes the key properties of the solution. Section 4 presents a detailed performance evaluation using real-world traces. Section 5 discusses issues and Section 6 concludes the paper.

2 Background and Related Work

2.1 A Beamforming Primer

In a scenario with M APs each equipped with k antennas and N clients each with a single (omni-directional) antenna, the baseband channel model at client n is given by,

$$y_n = \sum_{m=1}^{M} \mathbf{h}_{mn}{}^T \mathbf{x}_m + z_n \tag{1}$$

where each column vector $\mathbf{h}_{mn} = [h_{mn1}, h_{mn2} \dots h_{mnk}]^T$ denotes the channel between AP m and client n, $\mathbf{x}_m$ is the $k \times 1$ vector of transmitted signals from

AP m, y_n is the received signal and z_n is additive White Gaussian noise with zero mean and variance σ^2. The transmitted signal is related to the symbol (s_m) to be transmitted by a weight vector $\mathbf{w}_m$ as $\mathbf{x}_m = \mathbf{w}_m s_m$. By appropriately modifying $\mathbf{w}_m$, different beam patterns can be generated.

In directional beamforming *Dir*, the weights are set such that the beam patterns point a main lobe in a certain direction. Typically a fixed set of such beams (independent of the channel) are used to cover the the entire azimuth of 360 degrees. Such beams are known to be less effective indoors than outdoors due to multipath fading [1]. On the other hand, when the beam is adapted to leverage the multipath effects and maximize the SNR (Signal to Noise Ratio) at the receiver, the technique is called adaptive beamforming. When applied to a single link (which we refer to as Single User Adaptive Beamforming or SUA), the weights that optimize the received SNR are related to the channel gains $\mathbf{h}$ as $\mathbf{w} = \frac{\mathbf{h}*}{|\mathbf{h}|}$.

In a multi-AP setting, two types of beamforming are well known: (1) coordinated beamforming, where the APs share the channel estimates from each of their clients. (2) Uncoordinated beamforming where each AP only uses the channel estimates from itself to the other clients. While coordinated beamforming can provide higher gains than uncoordinated beamforming, it incurs more overheads due to joint computation of weight vectors. Consequently uncoordinated beamforming provides a good balance of performance and overheads making it desirable in practice.

Zero forcing (ZF) [1], is a popular uncoordinated beamforming method used for interference suppression. In ZF, the weight vectors are chosen by each AP such that the received signal power at the intended client is maximized subject to the constraint that the interference caused to the other clients is reduced to zero. The solution at each AP is related to the N channel vectors from that AP to the clients as described in [1]. For a 2 AP, 2 client case, the channel vectors from AP1 are given as $\mathbf{h_{11}}, \mathbf{h_{12}}$ and the vectors from AP2 as $\mathbf{h_{21}}, \mathbf{h_{22}}$. In this case, the weight vectors for AP1 is given as

$$w_1 = h_{11}^* - \frac{h_{12}^T h_{11}^*}{h_{12}^H h_{12}} h_{12}^* \tag{2}$$

where ^{H} denotes the Hermitian transpose and * denotes conjugate. The weight vector for AP2 is computed similarly. Compared to single user beamforming , a rate improvement up to 2x can be obtained in a two link setting using ZF.

2.2 Related Work

Theory: There has been an abundance of theory and protocol works on smart antennas [1,4] in the areas of channel estimation, beamforming and scheduling. Theoretical solutions to channel estimation [5,6,7] require accurate measurement of the baseband symbols at the receiver. Further, they also assume that antennas can be switched to create new beam patterns very fast (within a burst of symbols) and/or assume complicated receiver designs that require prior knowledge of the channel covariance matrix or rely on reciprocity which is difficult in the presence of interference.Recently in [8], the problem of user selection and beamforming has been considered, where a single AP beamforms to multiple clients simultaneously. Later works [2,9] have also considered the case where multiple APs cooperate to beamform to multiple clients. All these solutions assume that perfect channel state information is available at the AP.

Practical solutions in indoor wireless networks: There has been relatively fewer works on practical solutions to use beamforming in indoor wireless networks. In [10], the authors presented a system for downlink MIMO transmission, whereas the work in [11] proposes a channel estimation algorithm for random beamforming. Both these works use custom built narrow bandwidth hardware. DIRC [12] uses a fixed set of steerable 'directional' beams (independent of the channel) to increase spatial reuse whereas we use beams adapted to the channels that enable more spatial reuse compared to schemes that use directional beams.Interference alignment and cancellation [13] optimizes the performance by aligning the transmissions from multiple APs. However, it requires multiple elements at the clients and accurate channel state information.

In our prior work [14], we studied single user adaptive beamforming using software radios. Recently in [3], we considered the benefits of single-user beamforming using commercial Wifi devices. Both these works are intended to improve single link performance, do not consider interference and cannot be used to improve spatial reuse in a multi-link scenario. In contrast, in this work we design new estimation and scheduling algorithms to enable multiple concurrent AP-client links by interference suppression.

3 Light-Weight Multi-antenna Spatial Reuse

3.1 Network Model

We consider indoor WLANs where the APs posses multiple antennas arrays using which they can manipulate their beams. The APs are connected by a high-bandwidth ethernet connection. The clients posses a single (omni-directional) antenna and follow the Wifi association model wherein a single AP can only communicate to a single client at a time on a fixed channel. The objective is to maximize the weighted sum rate of all clients, where the weights are adjusted based on some fairness model.

3.2 Solution Overview

In this section we present our solution for improving spatial reuse in high density wireless networks which leverages the interference suppression capabilities of APs equipped with multiple antennas. The overarching goal is to design a solution which does not have sophisticated and unrealistic requirements and can be deployed in a WLAN without significant changes at the clients. We call our solution Light-weight Multi-antenna Spatial Reuse solution (LSR).

LSR consists of two main components, namely a light-weight channel estimation scheme and a low-complexity beamforming and scheduling scheme. The first component of LSR is a technique to perform channel estimation in a multiple AP multiple client Wifi network using just RSSI measurements at the client devices. i.e. *LSR does not require any phase or baseband symbol measurement at the receiver*. Conventionally channel estimation involves using special training symbols at the transmitter and enhanced receivers that provide both the magnitude and phase of received symbols [1]. While this has been accomplished in the open-loop context in 802.11n receivers, it is still non-trivial to feed it to the APs and perform closed-loop beamforming. Additionally, legacy clients do not have sophisticated hardware to provide channel estimates. Combined with the large variety of Wifi clients (e.g. PDAs, laptops, smartphones, etc), current techniques leave a significant fraction of Wifi clients unable to leverage beamforming. In contrast to conventional estimation procedures, LSR uses an intelligent antenna excitation scheme that relates the RSSI to the channel gain magnitudes and the differential phases. While we introduced the idea of estimating beamforming coefficients for single AP beamforming in [3], it is unclear whether they can be applied to multi-link beamforming. We first show that a straight-forward extension of the single AP procedure would suffer from poor accuracy and scalability when applied to a multi-AP setting. Then, we introduce optimizations to accommodate multi-link channel estimation. The second component of LSR involves selecting the right set of links to perform joint beamforming. Here LSR incorporates a Zero Forcing beamformer and a novel scheduler that identifies the best subset of links to operate on, while keeping the overall complexity low.

3.3 LSR Channel Estimation

We first revisit the basic procedure for single link estimation [3] and describe its enhancement for multiple links.

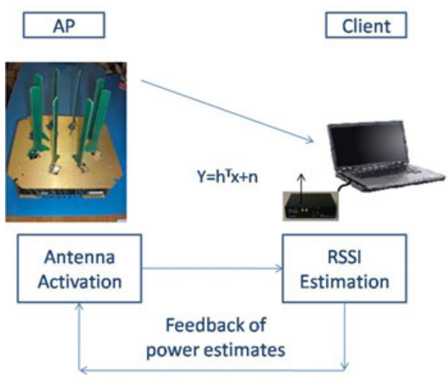

Fig. 1. LSR channel estimation

Single-Link channel estimation. We briefly summarize the idea of received power based channel estimation for single link beamforming that we presented in [3]. The key idea is to estimate 'differential' channel phases (instead of absolute channel phases) by employing *tandem activation of more than one antenna* and using received power estimates. Assuming a channel model given by 1, 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 that the transmit 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 signal power can be computed as $P_{ij} = |h_i + h_j|^2$. Thus, for tandem activation, the received power P_{ij} is given as

$$P_{ij} = P_i + P_j + 2\sqrt{P_i.P_j}.cos(\theta_{ij}) \tag{3}$$

where θ_{ij} is the *channel phase difference* between h_i and h_j. The resulting power at the client depends on the relative channel phase θ_{ij}. When $\theta_{ij} = 0$, the signals combine constructively causing the powers of the individual elements to add up. 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 3 as

$$\theta_{ij} = \cos^{-1} \frac{P_{ij} - P_i - P_j}{2\sqrt{P_i.P_j}} \tag{4}$$

By repeating this idea for pairs of antenna elements, the relative phases and the channel magnitudes can be obtained as described in [3].

Multi-link vs multiple single link estimation. In a M-AP N-client network, a straight-forward application of this procedure would be to perform the basic algorithm for each AP-client pair (All Pairs Estimation-APE). APE has the following issues. (1) The channel estimation for each AP-client pair would be separated significantly in time. Since interference suppression requires the channels to the different clients to be estimated as close in time as possible, the accuracy of the estimates could be affected. (2) The excitation overhead for APE increases proportional to MN. In contrast, by interleaving the estimation process across clients intelligently, the overhead in LSR is reduced significantly to be proportional to M. (3) In APE, the feedback overhead from each client (in packets) scales linearly with M. But in LSR, by using aggregated feedback the overhead is restricted to a few packets for practical AP densities. Thus LSR carefully considers these issues which are not addressed by APE.

Multi-Link Estimation and beamforming. A straight-forward extension of the above procedure to a multi-link scenario affects the accuracy of the estimates. Further it increases the time required for estimation and the feedback overheads. Hence, we develop additional mechanisms that must be designed into each step of the single-link solution [3] for efficient multi-link estimation.

i. Simultaneous multi-client measurement
As with single-link estimation, each one of the k elements at each AP is activated in isolation first. Following this, one antenna is designated as the reference and is activated with each of the remaining $k-1$ antennas in tandem pairs. This step is performed by each AP one after another by transmitting beacon packets. All clients estimate the channels simultaneously from each AP, enabling the measurements to be consistent across links. This also reduces the overheads compared to exciting each AP-client pair at a time.

ii. Aggregated Feedback from clients to APs
Single link estimation involves feedback of $2k-1$ received power values for the activations of step i. described previously. Instead of transmitting the feedback for each AP-client pair separately, the estimates from different APs are aggregated at each client into one or more packets and transmitted to the AP to which it is associated. An AP which receives the packet successfully broadcasts it on the ethernet backbone to the other APs so that they can extract the channel estimates required.

iii. Multi-link Ambiguity resolution
For the channel estimates between each AP-client pair, the magnitudes of the channel gains are obtained correctly. But the channel phases ϕ_{ij} (in radians) have an ambiguity due to the use of the cos^{-1} function in Equation 4. i.e. the correct θ_{ij} can be either of $\phi_{ij}, -\phi_{ij}, \pi - \phi_{ij}, -(\pi - \phi_{ij})$. The ambiguity is resolved by tandem activations with the estimates of Steps i and ii as described in [3] across each AP-client pair. As an additional optimization, the ambiguity resolution can also be performed by *applying the interference suppression weights* (instead of single link beamforming weights) across pairs of antennas so that the overhead of ambiguity resolution is reduced compared to that in [3].

Properties. The previously described channel estimation approaches uses assumptions about the channel coefficients and channel noise. While we experimentally showed that using power measurements provides single link beamforming benefits [3], here we provide the analytical reasoning for the effectiveness of the proposed solution for both single and multi-link beamforming. We briefly state the following properties and refer the interested reader to [15] for the details.
P1. Channel magnitudes and 'Differential phases' computed by LSR achieve the same beamforming benefits as using magnitudes and 'absolute' channel phases.

Conventionally beamforming requires accurate estimation of the channel vector $\mathbf{h}$ both in magnitude and phase. However, LSR uses differential phases measured using RSSI. We show analytically that the interference suppression benefits of LSR are close-to-ideal interference suppression benefits. We show how the weight vectors using absolute and differential phases differ only by a complex number whose magnitude is unity. Thus both the signal and interference powers are the same in both cases.

P2: For practical operating SNRs, thermal noise has minimal impact on the accuracy of LSR.

We explore how thermal noise affects RSSI based channel estimation. We show that for any AP-client pair, the differential phase estimate between two elements i and j when using RSSI measurements is related to the accurate differential phase and noise variance σ^2 as

$$\tilde{\theta}_{ij} = \cos^{-1} \frac{\cos\theta_{ij}.|h_i|.|h_j|}{\sqrt{(|h_i|^2 + \sigma^2))(|h_j|^2 + \sigma^2)}} \tag{5}$$

For practical Signal to Noise ratios (SNRs), the noise term is negligible compared to the signal term. Hence, we obtain $\tilde{\theta}_{ij} \tilde{=} \theta_{ij}$.

The magnitude of the channel gains can also be obtained from the RSSI as $|\tilde{h}_i| = \sqrt{P_i}$ or $|\tilde{h}_i| = \sqrt{|h_i|^2 + \sigma^2}$. Again for practical SNRs, this can be simplified to yield $|\tilde{h}_i| \tilde{=} |h_i|$. Thus, the effect of thermal noise on the accuracy of LSR is minimal as elaborated in [15].

3.4 Link Scheduling

Given a set of links, the above procedure estimates the channels using RSSI measurements at the desired and interfered clients that in turn helps achieve interference suppression. To achieve low complexity, we chose the ZF based beamforming solution, where each AP forces the interference power it causes to other clients to be reduced to zero[2]. Thus, given the channel vectors $\mathbf{h}_{mn}$ from AP m to client n, ZF solution maximizes $|\mathbf{h}_{mm}^T w|^2$ to the desired client m while minimizing the interference to client n, namely $|h_{mn}^T w|^2 = 0, \forall n \neq m$.

However, the design of a MAC solution would also require the determination of the set of links (scheduling) that must operate concurrently using ZF to maximize a desired system objective. To incorporate both throughput and fairness, we consider the popular weighted sum rate as the system objective to be maximized. For simplicity of discussions, consider a network with N co-channel APs and one client associated with each AP. The proposed solution will also apply to the case with multiple clients per AP. To determine which subset of links must operate concurrently a brute force search could be employed. However, it would need to evaluate the system objective for all combinations of n links, where $n \in [1, N]$. This results in a total of $\binom{N}{1} + \binom{N}{2} ... + \binom{N}{N} = 2^N$ combinations, making it exponential in N. To keep the scheduling solution scalable, we propose a greedy algorithm to identify the set of concurrent links that runs in $O(N^2)$ time, while yielding a performance comparable to that of the exhaustive approach. Note

[2] We note here that other closed loop MU-MIMO strategies can also be used at this stage given that the channel estimates computed in LSR with RSSI are close to the actual channel estimates as we show later.

Algorithm 1. Link Scheduler for LSR

1: Run LSR to obtain the channel estimates for desired and interference links using RSSI measurements.
2: $\mathcal{L} \leftarrow$ Set of active links; $C = \emptyset$
3: **while** $C \neq \mathcal{L}$ **do**
4: $\quad \ell^* = \arg\max_{\ell \in \mathcal{L} \setminus C} \sum_{k \in C' : C \cup \ell} T(\rho_{k,C'}) - \sum_{m \in C} T(\rho_{m,C})$
5: $\quad \hat{S} = \sum_{k \in C'} T(\rho_{k,C'})$
6: $\quad$ **if** $\hat{S} > S$ **then**
7: $\quad\quad C \leftarrow C \cup \{\ell^*\}$
8: $\quad\quad S = \hat{S}$
9: $\quad$ **else**
10: $\quad\quad$ Exit
11: $\quad$ **end if**
12: **end while**
13: Run ZF(C) to obtain weight vectors $\mathbf{w}_\ell$, $\forall \ell \in C$
14: Execute $\mathbf{w}_\ell$, $\forall \ell \in C$

that, one could also optimize the greedy algorithms further to improve system performance. Our goal here is to provide a low complexity greedy scheduler that will help us evaluate the performance of our RSSI based interference suppression scheme (LSR) in a network-wide setting.

The algorithm (Algorithm 1) works as follows: First the channel estimates are obtained using RSSI measurements for all the desired and interfering links (step 1). The scheduler initializes the set of selected links C to an empty set initially (step 2). At each iteration, it grows the selected set C by adding the link from the remaining set $\mathcal{L} \setminus C$ that provides the highest increase in the system utility using the rate table $T(\rho, C)$ for a given set of SINRs ρ for links in C (steps 4-8). The iterations stop when the system utility cannot be increased any more by the addition of a link (steps 9-11). The scheduler then runs the zero-forcing based beamforming solution on the selected set of links C to obtain the weight vectors and executes them (steps 13-14). The while loop and step 4 contribute to the time complexity and is bounded by $\frac{N(N+1)}{2}$ in the worst case, which is $O(N^2)$.

Note that, the above algorithm does not require the division of the network topology into multiple contention domains to determine the set of concurrent links in each contention domain. Instead, by virtue of being greedy, it automatically leverages the natural spatial reuse available in the network, to determine links both within and across contention domains that can operate concurrently.

4 Performance Evaluation

4.1 Measurement Setup

We perform experiments in a testbed that consists of six APs and six clients deployed in an indoor office environment whose layout is shown in Figure 2. Each of the APs is a commercial 802.11 b/g AP (Phocus Array [16]), which has

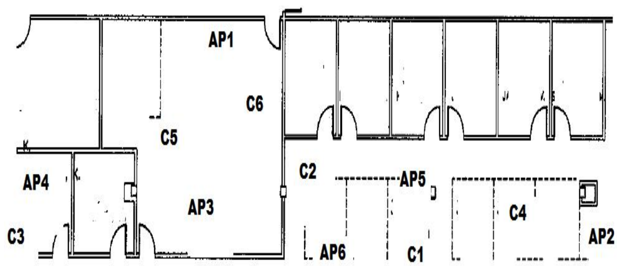

Fig. 2. Experimental Testbed

an eight element circular antenna array. Each client is a software radio which employs the Universal Software Radio Peripheral (USRP [17]) for the hardware and the GNURadio software [18] package along with the 802.11b code from [19]. The AP provides a command line interface to set specific beam patterns by writing weight vectors. We use this to implement the channel estimation component of LSR. The excitation steps of LSR are implemented as user level code on the Phocus Array.For every experiment, we also implement the complete channel estimation proposed in our previous work [14], which we refer to as 'perfect' channel estimates. In both LSR and our previous work, the first stage of channel estimation is performed concurrently. The ambiguity resolution is performed with the perfect channel measurement procedure first, immediately followed by LSR. We collect the channel estimate traces for both the schemes across every pair of AP and client, across physical locations and across time (multiple runs over several days).

We use the weight vectors for the 16 directional beams provided by the manufacturer [16]. For the competitive strategies ZF and SUA, we use the estimated channels to compute weight vectors. We compute the resulting Signal to Interference plus Noise ratio (SINR) for each of these strategies and the data rate table for 802.11g systems for each link. This table maps an input SINR to a data rate from 6 Mbps to 54 Mbps and is well known in the rate control literature [20]. We compute the sum rate of the concurrent links as the final metric (i.e. we use equal weights in our weighted sum rate objective function).

We organize the evaluation into three classes, namely channel estimation in isolation, joint channel estimation and beamforming, integrated operation (which includes channel estimation, beamforming and the scheduling components).

4.2 LSR Channel Estimation in Isolation

RSSI being a quantized estimate of the received power, results in a quantization error that may affect LSR. To quantify the performance difference between using full precision power estimates and RSSI values, we run the channel estimation algorithm on a two AP-two client scenario using full (16 bit) resolution for power and then with RSSI expressed as an integer in dB. Since we are interested in the similarity in the direction between two vectors, we define the similarity Index (SI) between two vectors as

$$SI(a,b) = \frac{|a^H b|}{|a||b|} \tag{6}$$

SI is the cosine of the angle between the vectors a and b. An SI value of 1 indicates that the two vectors are similar (highly correlated) whereas a value close to zero indicates that they are not similar.

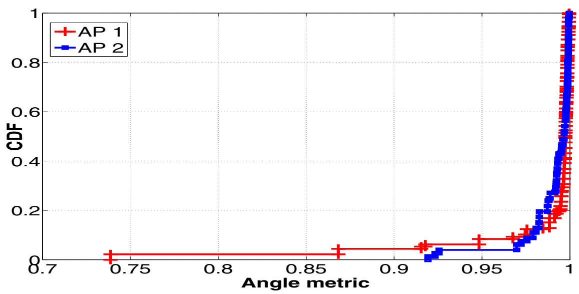

Fig. 3. Two Link Similarity Index

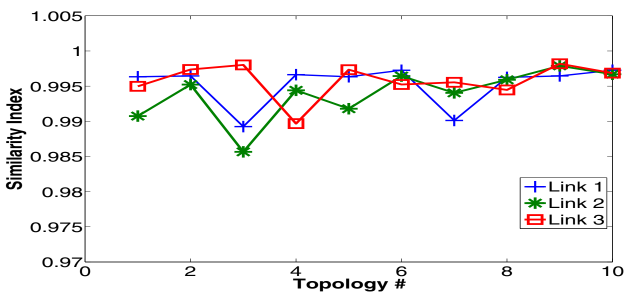

Fig. 4. Three Link Similarity Index

We first focus on a single pair of links in our testbed. We compute the weight vectors for zero forcing and determine the similarity index between the weight vectors computed using the ideal and RSSI measurements. The CDF of the similarity index for several runs of two link scenarios in our testbed is plotted in Figure 3. The figure reveals that the SI is very close to 1 for a majority of the cases and at most goes down to 75% still indicating a good correlation between the weight vectors using power and with RSSI.

We then explore if the weight vectors are also similar or more different for three link zero forcing. We select subsets of three links (with fixed AP-client association), and compute the three AP zero forcing vectors using both the perfect estimates and the estimates from LSR. For each link, we plot the similarity index (SI) in Figure 4. We observe from the figure that for all the 10 topologies, the similarity between the weight vectors is very high approaching 1. Thus both perfect estimates and using RSSI measurements yield zero forcing vectors which are highly correlated. Thus, LSR which requires RSSI measurements only, enables close to ideal weight computation for practical operating conditions.

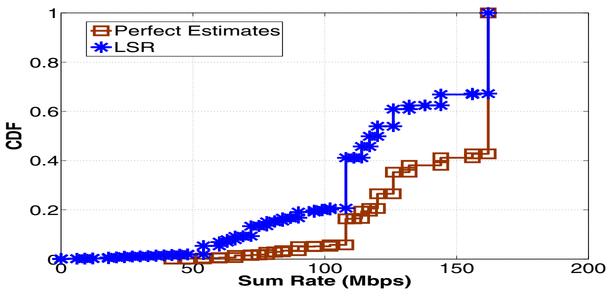

Fig. 5. CDF of sumrates

4.3 Joint LSR Channel Estimation and Beamforming

While our implementation of channel estimation operated successfully as described in Section 4.2, we encountered two practical challenges when implementing zero-forcing, (i) the inability of the hardware to update patterns quickly and (ii) power leakage from the transmit antennas at the APs. We believe that these can be overcome by better hardware and software designs in the future. Nevertheless, we evaluate the practical benefits of the interference suppression with LSR using realistic signal strength traces from our testbed. We consider the aggregate rate of concurrent links employing LSR, and compare it with other competing strategies.

Aggregate rate. We are interested in studying how the aggregate rate varies across different two link and three link topologies. For each of the two link runs, we choose two out of six APs and two out of six clients (yielding $\binom{6}{2} * \binom{6}{2} = 225$ topologies). We plot the CDF of the sum rates in Figure 6. The CDFs indicate that the distribution of rates is similar when both perfect and LSR estimates are used. Specifically, while more than 70% of the link configurations can sustain a data rate of 108 Mbps with perfect estimates, with LSR almost 60 % can sustain the same rate. The median rate is 108 Mbps, which is double the single link rate. Similarly, for the case of three links we study the performance over 400 topologies and plot the result in Figure 5. The resulting median rates are 162Mbps and 120 Mbps when using perfect estimates and LSR estimates. These results suggest that LSR yields significant benefits for a majority of indoor locations and is also close to the performance of zero forcing with ideal estimates.

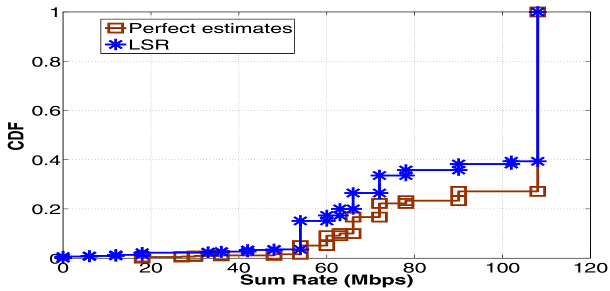

Fig. 6. Joint LSR channel estimation and beamforming: Two Link rates

Comparison of strategies. We are interested in analyzing the competitive advantage of LSR over other approaches. We consider the 225 link pairs (i.e two link topologies) for this evaluation. We compare with directional beamforming (Dir) [12] and adaptive beamforming (SUA) from an AP to its clients. Each of these strategies can be employed either in (1) a time division manner or (2) simultaneously across links. Activating multiple single user beamformed links simultaneously can yield a high sum rate if the channels of the users are well separated in vector space. Hence, for comparison, we consider the following four strategies Dir-TDMA, Dir-concurrent, SUA-TDMA, SUA-concurrent. We implement a search over all combination of beams with Dir to compare with an ideal implementation of [12].

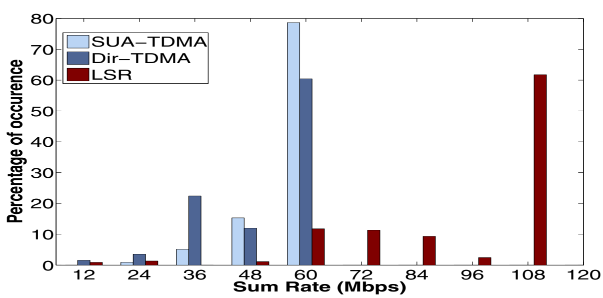

Fig. 7. Joint LSR channel estimation and beamforming: Strategies - TDMA

The histogram of the sum rates obtained with time division versions of these strategies is presented in Figure 7 and for the concurrent versions in Figure 8. From the figure, the following inferences can be made. (1) LSR outperforms both the TDMA approaches by a factor of 2 in the median sum rate. (2) For more than 80 % of the link pairs, LSR outperforms Dir-concurrent and SUA-concurrent. (3) The relative benefits of LSR are higher over SUA-concurrent than over Dir-concurrent since SUA-concurrent has lesser configuration flexibility than Dir-concurrent which can choose out of $16 * 16 = 256$ beam combinations.

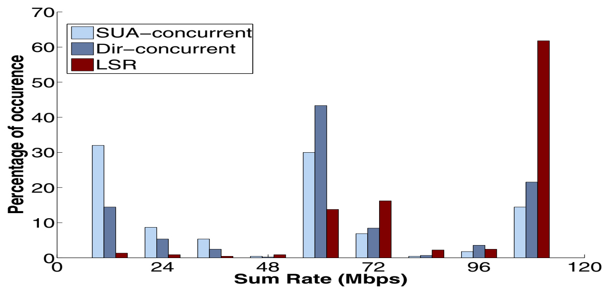

Fig. 8. Joint LSR channel estimation and beamforming: Strategies - Concurrent

4.4 LSR Integrated Operation

We consider an integrated operation of all components of LSR, namely the channel estimation, beamforming and scheduling.

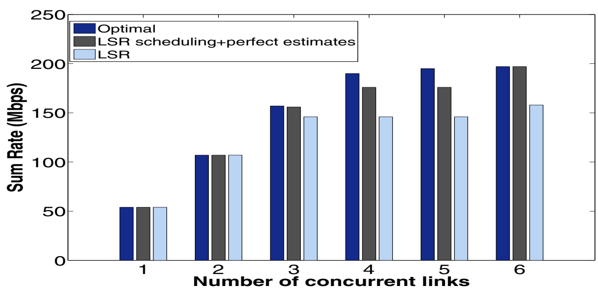

Fig. 9. LSR Integrated operations: Scalability of LSR

Scalability with number of concurrent links. We investigate how the performance benefits provided by LSR scale with the number of concurrent links as compared to ZF with perfect channel estimates. We feed in topologies with 2 link, 3 links, 4 links and so on upto 6 links.

The average sum rate for each case is plotted for optimal scheduling, greedy scheduling with perfect estimates and LSR in Figure 9. From the figure it is clear that the average rate for all three strategies is same for a small number of links. However, as the number of links increases, there is a degradation in rate due to the greedy nature of the algorithm employed in LSR. Further, when the number of links is large, the weight vector computation involves more channel vectors and the errors in the estimation could also affect the sum rate. However, the evaluation reveals the following surprising observation. The maximum degradation between optimal solution and LSR solution occurs around 5 links where the throughput is reduced from 195 Mbps to 146 Mbps. Although there is a degradation, we note that the performance of optimal Zero Forcing also begins to saturate around 4 links and the degradation does not increase further. Hence, LSR suffers minimal degradation while incurring much lesser complexity compared to optimal zero forcing.

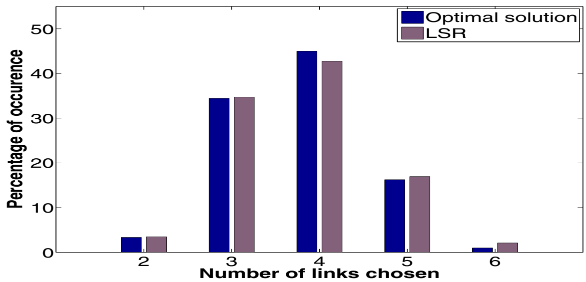

Fig. 10. LSR Integrated operations:# of links chosen

Comparison with optimal search and complexity. We investigate the run-time complexity of LSR under practical conditions by studying the impact of number of concurrent links and the sum rates achievable. The run-time complexity depends on the number of links at which the greedy algorithm terminates. If the optimal number of concurrent links is N, then the worst-case time complexity of $O(N^2)$ is always incurred by LSR. However, if the sum-rate optimal number of links is smaller than N, LSR incurs lesser complexity in practice. To study this, we generate $6! = 720$ topologies with 6 APs 6 clients from our traces by varying the AP-client association.

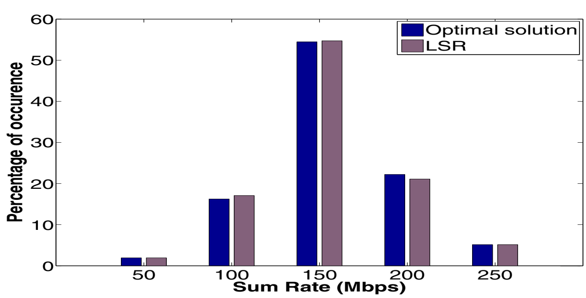

Fig. 11. LSR Integrated operations:Sum Rates obtained

For each topology, we determine the best sum-rate obtained by an exhaustive search of 2 links, 3 links and so on upto N links and the number of links that yields the best sum-rate. We then allow LSR to operate and determine the sum rate and the number of operating links identified by LSR. The results are plotted as histograms in Figure 10 and 11 for the optimal number of links and the sum rate. The main observations are as follows: (1) *With eight antennas at each AP, for more than 45 % of topologies, 4 concurrent links is the best* and for more than 35 % of the topologies activating 3 links yield the best rate. Thus the optimal number of links is much less than the number of antennas at each AP. (2) *In practical indoor deployments, the run-time complexity of LSR is likely to be lower than that expected in theory and brute force scheduling incurs a large complexity that is not justified.* (3) Around 55 % of topologies yield a sum rate around 150 Mbps. A very small number of topologies yields close to 250 Mbps and a smaller set yields close to 50 Mbps. Thus, the association in multi-AP smart antenna deployments is more important than their single antenna counterparts and a range of benefits can be obtained based on the topology.

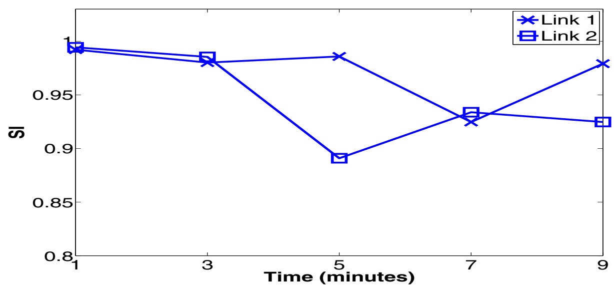

Fig. 12. Similarity index of weight vectors

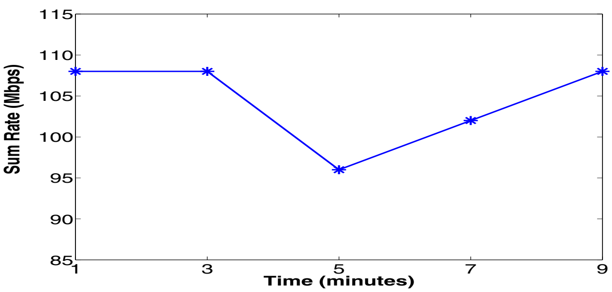

Fig. 13. Sum Rates vs. time

4.5 Impact of Channel Time Variation

We study the channel variation over time to understand how frequently beam patterns must be adapted. To study this effect, we first estimate the channel and compute weight vectors using the estimates. We then estimate the channel at a subsequent time duration to see whether the previously computed weights can be retainedin a two link scenario. We plot the Similarity Index between weight vectors computed at consecutive two minute intervals in Figure 12. It can be observed that the correlation is quite high for most of the time instants. The corresponding sum rates are plotted in Figure 13. It can be observed that as long as the correlation between vectors is greater than 95 %, the sum rate is retained at the maximum value of 108 Mbps. When transient fluctuations exist, the rate is reduced momentarily to 96 Mbps but is still quite close to the rate achieved with current estimates. We also note that we are constrained by the hardware in terms of how fast we can measure channels and write new patterns. With a kernel space implementation of LSR on the Phocus Array, we believe that the channel adaptation would be faster and the rate achieved would be closer to the rate obtained using accurate and timely channel estimates.

5 Discussion

– *Other beamforming and scheduling strategies:* While the proposed solution uses Zero Forcing and a greedy scheduling algorithm to obtain a good balance between complexity and performance, other alternatives to these could also be explored using the channel estimation component of LSR as is.
– *MIMO clients:* With the growing popularity of 802.11n, clients are likely to have more antenna elements. We believe that LSR can be extended to such clients and also combined with open-loop MIMO techniques for performance improvement.
– *Mobility:* We have considered primarily static clients in an indoor environment. While we believe that estimation and weight adaptation can be performed fast enough to account for indoor mobility, the design of the adaptation algorithm and its implementation are interesting directions to be explored.

6 Conclusions

In this work, we have developed LSR, a light-weight solution to obtain spatial reuse benefits in indoor wireless LANs equipped with smart antennas. The core components of LSR are: a scheme for multi-AP channel estimation which uses simple RSSI measurements, a Zero forcing beamformer which uses these estimates and a greedy link scheduler. We evaluate the solution using extensive traces collected from an indoor office environment. Our evaluation reveals that significant spatial reuse benefits can be obtained using LSR, and the median aggregate rate of the network improves up to 2.7x compared to related approaches.

References

1. Paulraj, A., Nabar, R., Gore, D.: Introduction to space-time wireless communications. Cambridge University Press (May 2003)
2. Skjevling, H., Gesbert, D., Hjorungnes, A.: Low-complexity distributed multibase transmission and scheduling. Eurasip Journal on Advances in Signal Processing (2008)

3. Lakshmanan, S., Sundaresan, K., Rangarajan, S., Sivakumar, R.: Practical beamforming using rssi measurements on off the shelf wireless clients. In: ACM Internet Measurement Conference (November 2009)
4. Ramanathan, R.: On the Performance of Ad Hoc Networks with Beamforming Antennas. In: ACM MOBIHOC (2001)
5. Ieee 802.11n working group draft, http://www.ieee802.org/11
6. Li, Y., Winters, J.H., Sollenberger, N.R.: MIMO-OFDM for Wireless Communications: Signal Detection with Enhanced Channel Estimation. IEEE Transactions on Communication 50(9), 1471–1477 (2002)
7. Minn, H., Al-dhahir, N.: Optimal training signals for mimo ofdm channel estimation in the presence of frequency offset and phase noise. IEEE Transactions on Communications 54(10) (October 2006)
8. Yoo, T., Jindal, N., Goldsmith, A.: Multi-antenna broadcast channels with limited feedback and user selection. IEEE JSAC 25(7) (September 2007)
9. Chae, C.-B., Inoue, T., Mazzarese, D., Heath, R.W.: Non-iterative multiuser mimo coordinated beamforming with limited feedforward. In: IEEE ICASSP (April 2008)
10. Samardzija, D., Huang, H., Valenzuela, R.A., Sizer, T.: An experimental downlink multiuser mimo system with distributed and coherently-coordinated transmit antennas. In: IEEE ICC (2007)
11. Kecicioglu, B., Ozdemir, O., Torlak, M.: Opportunistic multiple antenna systems: Channel estimation and experimental results. In: IEEE ICC (2006)
12. Liu, X., Sheth, A., Kaminsky, M., Papagiannaki, K., Seshan, S., Steenkiste, P.: Dirc: Increasing indoor wireless capacity using directional antennas. In: ACM SIGCOMM (August 2009)
13. Gollakota, S., Perli, S.D., Katabi, D.: Interference alignment and cancellation. In: ACM SIGCOMM (August 2009)
14. Lakshmanan, S., Sundaresan, K., Khojastepour, M., Rangarajan, S.: Towards adaptive beamforming in indoor wireless networks: An experimental approach. In: IEEE Infocom (Miniconference) (April 2009)
15. Lakshmanan, S., Sundaresan, K., Khojastepour, M., Rangarajan, S.: Spatial reuse with smart antennas in indoor wlans (July 2009), http://www.nec-labs.com/~karthiks/papers/tr-lsr.pdf
16. Fidelity-comtech inc., http://www.fidelity-comtech.com
17. Ettus inc., http://www.ettus.com
18. Gnuradio project, http://www.gnuradio.org/trac
19. Adroit project, https://acert.ir.bbn.com/projects/adroitgrdevel
20. Ramachandran, K., Kokku, R., Zhang, H., Gruteser, M.: Symphony: Synchronous two-phase rate and power control in 802.11 wlans. In: ACM MOBISYS (2008)

A Complexity and Overheads

In this section, we quantify the overheads incurred in LSR in terms of the number of excitations at the transmitter and the number of feedback bits. For an M AP-N client network, beamforming with interference suppression requires the estimates of MN channel vectors. The first stage involves $2k-1$ excitations since this is completed in parallel at all the clients together. The ambiguity resolution step involves $4k-4$ excitations per Ap-client pair. Hence, the total number of excitations can be shown to be $6k-2$ per AP-client pair. The the number of receive observations that must be fed back is $3k-2$. With B bits to represent each symbol [3], the total number of bits to be fed back from each client to each AP is $O_s = (3k-2) * B$. Since the observations are aggregated, a single client sends $O_s = (3k-2) * B * M$ bits.

In practice, with 8 antenna elements and 16 bit precision, the total number of excitations at each AP for each client is 43. Even with 20 bytes per excitation (the size of a control packet in 802.11), the number of bytes to be transmitted from the AP is 860 bytes. Since the excitation is in two stages, each AP would need about two packets per client for excitation. In a 6 AP 6 client environment, 12 packets per AP leads to 72 packets for the entire estimation procedure. The process for the entire network can be completed within one or two seconds. The total feedback overhead from a client to each AP is limited to 44 bytes (i.e $22*16$ bits). Since multiple feedback packets can be aggregated, even with 6 APs, the feedback overhead per client is just 264 bytes, which is less than a normal 802.11 data packet. Thus, the overhead of LSR channel estimation is very small and can be easily accommodated in any deployment.

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