UCL DEPARTMENT OF
COMPUTER SCIENCE

Research Note
RN/10/02

Cone of Silence:
Adaptively Nulling Interferers in Wireless Networks

30th January 2010

Georgios Nikolaidis

Astrit Zhushi

Kyle Jamieson

Brad Karp

Abstract

Dense 802.11 wireless networks present a pressing capacity challenge: users in proximity contend for limited unlicensed spectrum. Directional antennas promise increased capacity by improving the signal-to-interference-plus-noise ratio (SINR) at the receiver, potentially allowing successful decoding of packets at higher bit-rates. Many uses of directional antennas to date have directed high gain between two peers, thus maximizing the strength of the sender's signal reaching the receiver. But in an interference-rich environment, as in dense 802.11 deployments, directional antennas only truly come into their own when they explicitly null interference from competing concurrent senders. In this paper, we present Cone of Silence (CoS), a technique that leverages software-steerable directional antennas to improve the capacity of indoor 802.11 wireless networks by adaptively nulling interference. Using in situ signal strength measurements that account for the complex propagation environment, CoS derives custom antenna radiation patterns that maximize the strength of the signal arriving at an access point from a sender while nulling interference from one or more concurrent interferers. CoS leverages multiple antennas, but requires only a single commodity 802.11 radio, thus avoiding the significant processing requirements of decoding multiple concurrent packets. Experiments in an indoor 802.11 deployment demonstrate that CoS improves throughput under interference.

Cone of Silence: Adaptively Nulling Interferers in Wireless Networks

Georgios Nikolaidis Astrit Zhushi Kyle Jamieson Brad Karp
University College London
UCL CS Research Note RN/10/02
{g.nikolaidis, a.zhushi, k.jamieson, bkarp}@cs.ucl.ac.uk

ABSTRACT

Dense 802.11 wireless networks present a pressing capacity challenge: users in proximity contend for limited unlicensed spectrum. Directional antennas promise increased capacity by improving the signal-to-interference-plus-noise ratio (SINR) at the receiver, potentially allowing successful decoding of packets at higher bit-rates. Many uses of directional antennas to date have directed high gain between two peers, thus maximizing the strength of the sender's signal reaching the receiver. But in an interference-rich environment, as in dense 802.11 deployments, directional antennas only truly come into their own when they *explicitly null* interference from competing concurrent senders. In this paper, we present Cone of Silence (CoS), a technique that leverages software-steerable directional antennas to improve the capacity of indoor 802.11 wireless networks by adaptively nulling interference. Using *in situ* signal strength measurements that account for the complex propagation environment, CoS derives custom antenna radiation patterns that maximize the strength of the signal arriving at an access point from a sender while nulling inteference from one or more concurrent interferers. CoS leverages multiple antennas, but requires only a single commodity 802.11 radio, thus avoiding the significant processing requirements of decoding multiple concurrent packets. Experiments in an indoor 802.11 deployment demonstrate that CoS improves throughput under interference.

1. INTRODUCTION

Spurred by the availability of low-cost commodity radio hardware and freely usable unlicensed spectrum, users have enthusiastically embraced 802.11 wireless networking in home and office environments. As these networks proliferate rapidly, particularly in populous urban areas, their deployment density increases significantly. Measurements of 802.11 base station deployments in major US cities taken in 2005 already showed thousands of cases in which four or more 802.11 access points mutually interfered [1]. As only three non-overlapping channels are available in 802.11b/g, these increasingly dense deployments pose a wireless capacity challenge—physically proximal networks must share finite bandwidth.

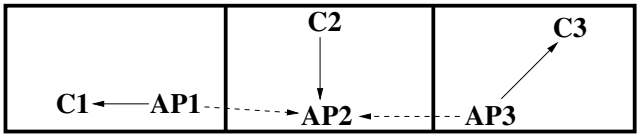

Figure 1: Typical dense 802.11 deployment in apartments or offices. AP1, AP2, and AP3 are access points; C1, C2, and C3 are clients. Solid arrows denote desired transmissions; dashed arrows denote unintended interference.

Consider a dense deployment of 802.11 networks on overlapping channels, typical in today's residential and commercial areas, as shown in Figure 1. The occupant of each of three apartments (or offices) operates his own access point (AP) with a standard omnidirectional antenna, and because of their close proximity, AP1 and AP3 interfere with reception by AP2. In particular, if client C2 and AP1 are hidden from one another, client C2 may transmit to its AP, AP2, and AP1 may simultaneously transmit to its client C1. Consider C2 as the sender, AP2 as the receiver, and AP1 as the interferer. Interference from AP1 will reduce the throughput C2 achieves at AP2. Indeed, more than one interferer may transmit concurrently; *e.g.,* AP3 might transmit to clients of its own, too, further interfering at AP2.

When multiple wireless networks operated by independent, non-cooperating individuals interfere, a receiver in one network derives no benefit from successfully decoding transmissions from an interferer in another network; data from another network are typically of no interest. Moreover, operators of these networks do not centrally coordinate transmit schedules, or share decoding information among nodes. Many recent advances in mitigation of interference have targeted environments where one enterprise operator closely coordinates multiple cooperating senders [3, 6], or where concurrent transmissions' contents are all of interest to a single receiver (*i.e.,* where interferers are part of the *same* network) [10]. In contrast, in this work, we specifically focus on mitigating interference in the ubiquitous "chaotic," non-cooperative deployments described above.

Directional antennas hold great promise for improving throughput on wireless links in such dense, interference-rich deployments. The throughput achieveable on a link depends on how well the receiver can discern a sender of interest's signal, while distinguishing it from compet-

ing background noise and interference from other concurrent senders—on the *signal-to-interference-plus-noise ratio* (SINR). The greater the bit-rate at which a packet is transmitted, the greater the SINR with which the receiver must receive the packet in order to decode it successfully. And at a given transmit bit-rate, as SINR increases, bit-error rate (BER) decreases, reducing costly link-layer retransmissions.

Extracting the greatest SINR from a receiver's directional antenna entails solving two distinct problems. First, how can one direct gain toward a sender of interest's signal, thus improving the strength with which it is received? And second, how can one *avoid* directing gain toward interfering signals from concurrent transmitters, and thus *null* interference? A system may independently address either or both of these problems. Solving either increases SINR, and can thus improve throughput.[1] The relative benefits to SINR of directing gain toward a sender's signal *vs.* nulling interference depend heavily on the deployment scenario. In an interference-rich environment, nulling is *vital* to achieving the full SINR and throughput improvements a directional antenna can offer. Lakshmanan *et al.* offer a technique for maximizing gain toward a sender of interest indoors [5], but this technique does not explicitly null interferers.

Multipath propagation, commonplace indoors, significantly complicates effective use of directional antennas by causing a sender's signal to arrive at a receiver in multiple components from unpredictable bearings, each with a different phase.[2] Figure 2 offers an idealized illustration of this phenomenon in a simple topology, where a receiver equipped with a directional antenna attempts to receive from sender S while an interferer I transmits concurrently. The solid arrows indicate multiple components from S while the dashed arrows indicate multiple components from I. We see that multiple components arise from reflections of each transmitter's signal (off walls and any of the many other reflective objects in the indoor setting) that depend not only on the locations of the sender and receiver, but on the time-varying minutiae of the physical surroundings. In order to maximize SINR successfully, and extract full benefit from directionality, the receiver must configure its antenna such that high-gain *lobes* are directed toward the incident bearings of S's signal, while low-gain *nulls* are directed toward the incident bearings of I's signal. The heavy line undulating about the receiver represents just such a *gain pattern* for its directional antenna. Radial distance from the center of the pattern to this line indicates the gain of the antenna in dB in that radial direction, and lobes are oriented toward S's components, while nulls are oriented toward I's. Software-steerable, *phased array* antennas, such as the one used in this work (described in Section 2.2), allow quick shaping of the gain pattern entirely electrically, under software control, *without* any mechanical motion. While phased array antennas have previously been used successfully indoors with fixed, essentially single-lobe beam shapes [6], such beam shapes cannot maximize SINR as effectively as ones flexibly and dynamically *customized* in accordance with the specific multiple arrival bearings of senders' and interferers' signals.

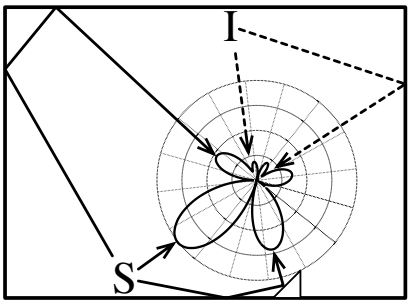

Figure 2: Example of multipath propagation between a sender S, inteferer I, and a receiving directional antenna. Solid arrows represent components from S; dashed from I. Boundaries are reflective walls; the triangle represents a reflective object.

In this paper, we present Cone of Silence (CoS), a technique for improving throughput under interference in 802.11 wireless networks. CoS incorporates two main techniques: *SamplePhase,* a method for accurately, robustly, and efficiently deriving a custom pattern for an AP's antenna that maximizes the signal strength received from a specific sender or interferer; and *Silencer,* a method that, given signal-maximizing patterns for a sender and one or more interferers, produces a single pattern that simultaneously nulls the interferers while maximizing signal strength from the sender of interest, thus maximizing SINR and throughput.

An evaluation of a prototype of CoS on an indoor 802.11b/g testbed demonstrates that CoS can improve a sender's throughput under interference over that achieved by an omnidirectional receiver by between 1.6× and 17×, and that CoS improves receive throughput by nulling one or two concurrent interferers. CoS achieves these substantial performance improvements while offering the following key properties:

- An AP using CoS can null even *uncooperative* interferers from which it can receive packets—CoS does not require APs to schedule transmissions collaboratively, as do previous techniques for mitigating interference with directional antennas [6] or multiple antennas [3].
- Unlike schemes that receive many concurrently transmitted packets, but require processing-intensive full decoding of each one [10, 3], CoS can null multiple concurrent interferers using multiple antennas connected to only *a single commodity 802.11 radio.* To our knowledge, CoS's Silencer is the first such implementation of a decorrelator [11].

[1]Complementary arguments apply when a sender transmits with a directional antenna. In this paper, we focus on *reception* using directional antennas, though as we discuss in Section 5, we believe the techniques we describe will be useful at senders, too.

[2]For simplicity of exposition, we ignore phase in this discussion; we describe the role of phase in detail in Section 2.

2. DESIGN

We begin with a brief overview of the use scenario for CoS, followed by a primer on the phased array antenna hardware platform on which CoS is built. Thereafter, we next present SamplePhase, an algorithm for measuring the wireless channel between the CoS AP and other radios. We then present the design of Silencer, which builds on SamplePhase to simultaneously steer our AP towards associated clients and null one or more interferers.

2.1 Use Scenario

Consider again the topology in Figure 2. Recall that in order to improve SINR at a receiving AP equipped with a beam-steerable directional antenna, CoS must derive a pattern for the antenna that gathers path components from the sender of interest, S, while nulling path components from the interferer, I. CoS goes about that goal in two logical steps:

- First, CoS considers the sender of interest and each interferer individually. For each such remote transmitter, a CoS AP applies the SamplePhase algorithm, described in Section 2.3, to derive one receive pattern for each remote transmitter that maximizes received signal strength from that remote transmitter alone.
- Second, for each sender of interest, a CoS AP applies the Silencer algorithm, described in Section 2.4, whose input consists of the receive pattern that maximizes signal strength at the AP from the sender of interest, as well as one pattern for each interferer that does the same for that interferer. Silencer produces one pattern for each sender of interest that nulls all interferers while directing gain to maximize signal strength from the sender of interest.

CoS must determine the identities of interferers. Doing so for 802.11 transmitters who interfere strongly at the AP is not difficult; an AP may simply scan the channels that overlap its own periodically, and record the MAC addresses of any senders that occupy the channel heavily. It is precisely the strongest interferers that stand the greatest chance to reduce a sender of interest's throughput to the AP that will be most easily identified in this fashion. When receiving from *any* sender of interest, CoS *always* nulls toward *all* interferers that it has identified. CoS cannot be certain that an interferer is sending at any given time, and so may needlessly null that interferer. That choice would be problematic if nulling significantly reduced throughput from the sender of interest as a side effect. In Section 3.2, we present experimental evidence to argue that nulling does not do so—in effect, that nulling an interferer tends to be "safe" for the sender of interest.

Like any AP, CoS maintains a list of associated clients. CoS stores the Silencer-produced, throughput-maximizing pattern tailored to each associated client, and configures its directional antenna to the appropriate such pattern each time a client transmits to the AP. To do so, however, CoS must have foreknowledge of when each client will transmit.[3] As do others who have proposed interference mitigation techniques for 802.11 networks [3], we envision that an AP would use a TDMA schedule among its own clients to allow it to predict when each client transmits. Because TDMA scheduling among a base station and its clients is a well-understood area in the literature—the original 802.11 specification includes the PCF MAC, a TDMA approach [4]—we assume the availability of a TDMA implementation in this work, and focus on nulling interference given that a CoS-enabled AP can use TDMA to predict which of its clients will send.

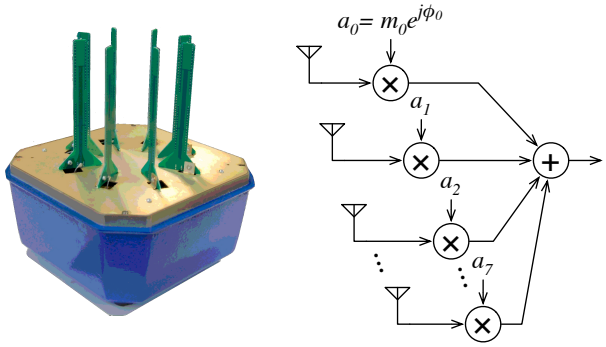

Figure 3: *Left:* Phocus phased array 802.11b/g antenna with eight elements. *Right:* simplified array model.

2.2 Hardware Platform

We have built the CoS prototype atop the Phocus phased array 802.11b/g antenna [2] manufactured by Fidelity Comtech Inc. and shown in Figure 3, left. The array consists of eight elements spaced equally on the circumference of a circle, each of which consists of four stacked dipoles. A signal processing module (Figure 3, right) mixes the signal from or to element k with a complex gain a_k allowing adjustment of its phase ϕ_k and magnitude m_k during reception and transmission, respectively. The eight resulting signals are summed and in turn connect to an antenna port of a standard Atheros AR5413 802.11 chipset. The gain and phase applied to each element may be independently controlled in software, the phase in single-degree increments between -180 and 180 degrees and the gain in 1% increments between 0% and 100%.

Taken together, the phase shifts and magnitudes configured for all elements define a complete pattern. Changing the array's pattern takes approximately 120 μs.

As supplied by the manufacturer, the antenna's software supports only 17 factory-configured "stock" patterns: one omnidirectional, and the remaining 16 "high-gain", each of which consists of a single high-gain lobe approximately 43 degrees in width, pointing toward one of 16 equally spaced directions about the antenna's center. The CoS software is

[3]It is important to note that the AP need *only* predict transmissions from *its own clients*—as explained above, CoS assumes that inteferers, who are uncooperative because they are part of other networks, transmit constantly.

not bound by this restriction; it can configure each element independently to any of the supported phase and gain values, yielding a vast variety of possible multi-lobed patterns.

2.3 Measuring the Channel: SamplePhase

The first challenge we face is measurement of the wireless channel from each client to all of the AP's antennas. Reflections off walls and other objects found in the typical indoor home or office mean that indoor wireless networks operate in the presence of strong multipath reflections, wherein transmissions arrive from multiple directions at the AP.

In order to produce the best end-to-end performance, we argue that a channel measurement algorithm should satisfy the following objectives:

- *Performance:* The algorithm should produce measurements that result in the best throughput.
- *Efficiency:* The overhead of the algorithm should be as low as possible without sacrificing performance.
- *Reliability:* The algorithm should meet the above two objectives consistently, with low performance variance, even in challenging wireless environments.

In prior work, Lakshmanan *et al.* [5] have proposed a channel measurement algorithm that we improve upon here. We experimentally compare the two algorithms in Section 3.1 and briefly touch upon efficiency differences between the two algorithms in Section 4.

The SamplePhase algorithm. To reduce complexity, our approach leverages received signal strength (RSS) readings measured at a client, from packets sent by the AP.[4] Consider an approximation of the wireless channel between the k^{th} AP element and the client as a single complex number of a certain phase θ_k and magnitude $\sqrt{P_k}$.[5] Then based on RSS readings, SamplePhase outputs a set of eight channel measurements: $\sqrt{P_1}e^{j\theta_{r1}}$, $\sqrt{P_2}e^{j\theta_{r2}}, \ldots,$ $\sqrt{P_8}e^{j\theta_{r8}}$, where r is the index of a reference element in the array, and $\theta_{kl} = \theta_l - \theta_k$ is the phase of element k relative to element l. SamplePhase only measures the relative differences between the channel phases θ_k, since these determine beam shape.

SamplePhase measures the individual channel magnitudes from each individual element $\sqrt{P_k}$ $(k = 1 \ldots 8)$ directly. Our algorithm transmits 25 contiguous bursts of three *probe* packets each from each individual element of the AP to the remote node to which the channel is being measured. The bursts are interleaved across elements, so that a period of interference impacts just a small number of packets from any particular element.

SamplePhase's phase measurements are based on the following observation about the wireless channel. Suppose the phase of the channel from element k to the remote node is some value θ_k, and the phase of the channel from element l to the remote node is another value θ_l. Further, suppose that the AP transmits data with a phase difference δ between elements k and l that we choose and program into the AP. Then, by the principle of superposition (Lakshmanan *et al.* provide a detailed derivation), the power of the elements' combined transmissions at the client is

$$P_{kl}(\delta) = P_k + P_l + 2\sqrt{P_k P_l}\cos(\theta_l - \theta_k + \delta). \quad (1)$$

Rearranging the above, we find the following:

$$\cos(\theta_{kl} + \delta) = \frac{P_{kl}(\delta) - (P_k + P_l)}{2\sqrt{P_k P_l}} \quad (2)$$

The above suggests the following way of estimating θ_{kl}: using empirically measured values of P_k, P_l, and $P_{kl}(\delta)$, sample the expression on the right-hand side of Equation 2 at one or more values of δ. Then, the best estimator of θ_{kl} will minimize the sum of squared errors between the empirically measured values from the right-hand side of Equation 2 and computed values from the left-hand side of the same equation. Sampling multiple evenly spaced values of δ removes any phase ambiguity with high likelihood, as explained in the Appendix A. For simplicity and to overcome practical limitations on the number of antenna patterns that the Fidelity Comtech array can store at once, SamplePhase uses four evenly spaced values of δ for its phase measurements.

SamplePhase microbenchmarks. Figure 4 shows example measurements of $\sqrt{P_{kl}(\delta)}$ for three representative element pairs on an AP-client link in our testbed (described in Section 3), as we vary δ between -180 and 180 degrees. We see the expected sinusoidal relationship, and note that estimating θ_{kl} by using measurements of P_{kl} near the peak or trough of the sinusoid may decrease accuracy, because of the quantization of the sample data and the decreased slope of the sinusoid near those points. On the same plots, the vertical lines represent the peaks of the sinusoid as predicted by the estimators of θ_{kl} produced by the SamplePhase and the Lakshmanan *et al.* methods. We see that by using multiple sample points, SamplePhase finds the peaks of the sinusoidal data better than the prior method for these representative element pairs.

Mean absolute error	Variance of the absolute error	
SamplePhase	10	122.5
L. et al	26	688

Table 1: Mean and variance of absolute error for SamplePhase and L. *et al.* methods.

Table 1 lists the mean and the variance of the absolute error for both methods. To derive these results, we fit the empirical data to the sine wave whose phase minimizes the

[4] In our prototype, the CoS AP sends measurement probes to all remote nodes, which record RSS measurements and return them out-of-band to the AP. In a production deployment, we envision that the CoS AP would send 802.11 null data frames to its clients to elicit ACKs, and measure RSS on these returning ACKs (in fact, the Phocus array software already implements this functionality as shipped). For interferers, CoS could simply measure RSS on received frames.

[5] The 802.11b/g wireless channel is 20 MHz wide and so in fact cannot be completely characterized by a single complex value; we touch on this point in Section 4.

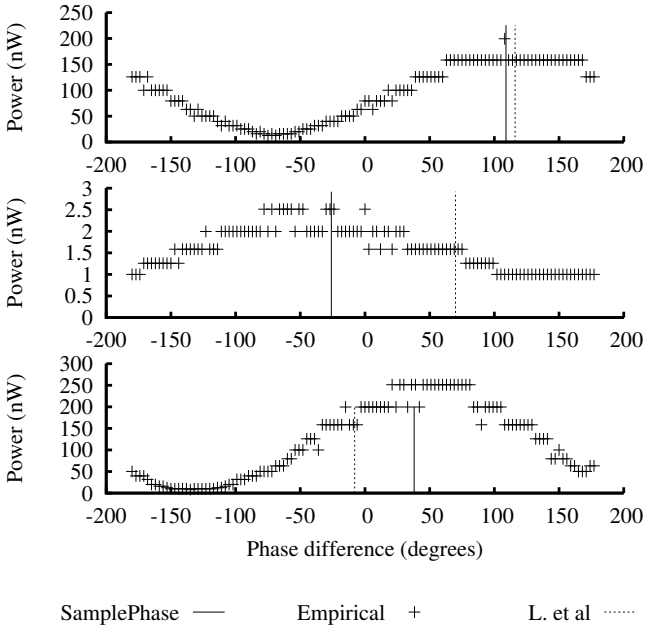

Figure 4: Empirical power measurements (P_{kl}) of RSS for three representative element pairs for an AP-remote node link when pairs of AP antenna elements transmit simultaneously with varying phase difference (δ) to a remote node ("Empirical" points). The SamplePhase and Lakshmanan *et al.* estimates of the peak of the sinusoid appear as vertical lines.

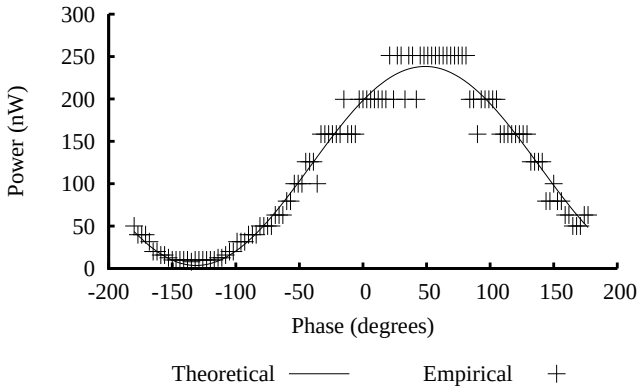

Figure 5: Empirical data with the theoretical fit superimposed.

sum of squared errors (Figure 5). Note that this fitting to a sine wave uses many more samples then both SamplePhase and Lakshmanan *et al.*'s method. We measured the absolute error of each of the two methods by computing the absolute difference between each method's estimate and the peak of the best-fit sine wave. Then looking across all links in our testbed (shown in Figure 8 on p. 7), and at all array element pairs on every testbed link, we computed the average absolute error and variance of the absolute error for both SamplePhase and the Lakshmanan *et al.* method. The results are shown in Table 1. In this table, we see that SamplePhase offers both a lower absolute mean error and a lower variance of absolute error. We conclude that SamplePhase more accurately finds the antenna phase that maximizes the signal strength when two elements send together.

In Section 3, we show that SamplePhase furthermore finds better overall beamforming and interference-cancelling patterns than the prior method, as determined by the end-to-end metric of throughput.

Beamforming toward a client. Once the AP has measured the channel between itself and a client, it can beamform its transmissions to or receptions from that client by weighting the kth element's input by the channel measurement to the kth element, $\sqrt{P_k}e^{j\theta_{kr}}$, where r is the reference element chosen during SamplePhase's measurement. This results in co-phasing the signals from all antennas so that they align and constructively interfere. The combination of co-phasing and weighting proportional to $\sqrt{P_k}$ maximizes signal-to-noise ratio (SNR) at the receiver and is known as maximal ratio combining (MRC) in the literature. MRC does not maximize signal to noise plus interference ratio (SINR), however, and so interfering transmissions will impact a beamforming AP's throughput, as we show in Section 3. We therefore seek a way to null interfering clients and maximize SINR.

2.4 Nulling Interferers: Silencer

Silencer is an implementation of a decorrelator [11], a computational structure that allows distinct signals to be received concurrently. What distinguishes Silencer from other decorrelator implementations is that Silencer can recover a signal from a sender of interest while nulling other concurrently received signals *without* decoding these other signals.

Using channel measurements from the methods in Section 2.3, we can represent the channels to clients as vectors in an eight-dimensional space (since our AP has eight elements):

$$\mathbf{h}_c = \begin{bmatrix} \sqrt{P_1}e^{j\theta_1} \\ \sqrt{P_2}e^{j\theta_2} \\ \vdots \\ \sqrt{P_8}e^{j\theta_8} \end{bmatrix} \quad (3)$$

where the measurements for $\mathbf{h}_c$ are taken at client c. To null an interferer i (either another AP or an interfering client), the AP measures the channel $\mathbf{h}_i$ between itself and the interferer, and using the Gram-Schmidt algorithm, computes a basis in $\mathbb{C}^8$ for the vector space orthogonal to h_i (indicated by V_i in Figure 6). Then, Silencer projects the received signal $\mathbf{y}$ onto V_i (indicated by $\text{Proj}_{V_i}(\mathbf{y})$ in Figure 6).

After the interference cancellation step, Silencer directs gain in the direction of the intended client's channel $\mathbf{h}_c$ (in the V_i vector subspace). If we represent projection onto V_i with the 8×8 complex matrix $\mathbf{Q}_i$, the overall operation on the received signal $\mathbf{y}$ is therefore $(\mathbf{Q}_i\mathbf{h}_c)^*\,\mathbf{Q}_i\mathbf{y}$. We program the AP with the eight-element, complex-valued vector $\mathbf{Q}_i^*\mathbf{Q}_i\mathbf{h}_c$ to implement this operation.

Generalization to multiple interferers. Silencer easily generalizes to multiple interferers i_1 to i_7, each with a different channel estimate $h_{i_1} \ldots h_{i_7}$, by using the Gram-Schmidt process to construct a vector subspace orthogonal to the span of all the inteference vectors. In Section 3.4, we present ex-

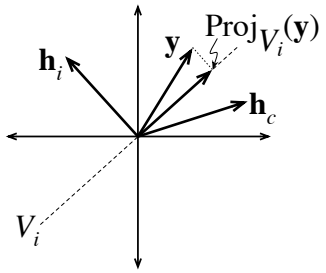

Figure 6: By projecting the received signal $\mathbf{y}$ onto the vector subspace orthogonal to an interferer's channel $\mathbf{h}_i$, Silencer (shown here in $\mathbb{R}^2$ for ease of exposition) nulls the signal from the interfering client.

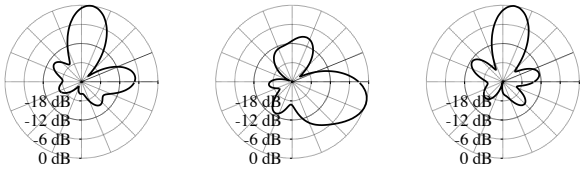

Figure 7: Impact of Silencer on antenna gain pattern: note that these empirical figures show gain *vs.* direction, but do not show how the antenna manipulates the phase of the received signals. *Left:* an MRC gain pattern maximizing SNR for a sender. *Center:* an MRC gain pattern maximizing SNR for an interferer. *Right:* the resulting Silencer gain cancelling the interferer and beamforming toward the sender.

perimental results nulling up to two out of three simultaneously transmitting senders.

Practical limitations in nulling interferers. Although CoS can theoretically entirely remove interference from up to seven simultaneous interferers, several practical design issues limit real-world system performance:

- *Hearing the interferer.* In order to compute $\mathbf{h}_i$, CoS needs to receive a sufficient number of packets from interferer i. This precludes nulling the most distant interferers, since our commodity hardware detects packets only down to -94 dBm. Nonetheless, we show in Section 3.2 that CoS can even null interferers with PRRs to the AP as low as 9%.
- *Estimating the channel to the interferer.* The accuracy of the channel estimation algorithm will impact the degree to which CoS can beamform towards clients and null interferers. For OFDM modulations, since CoS measures and beamforms across all OFDM subcarriers simultaneously, it does not capture inter-subcarrier differences in the 20 MHz WLAN channel. This is a practical tradeoff: measuring inter-subcarrier differences would require a software-defined radio, or a PHY interface that returns per-subcarrier RSS readings. CoS sacrifices some performance for the simplicity of running on a commodity hardware platform and the speed of only requiring three $\mathbb{C}^8$ matrix multiplications and a Gram-Schmidt iteration in order to compute a beamforming pattern that nulls a new interferer.
- *Adapting when an interferer ceases sending.* When CoS nulls an interferer that has since ceased transmission, it sacrifices some amount signal power that would have improved the overall SINR had it not nulled that interferer. In Section 3.2 we quantify the throughput impact of nulling towards an interferer despite that interferer's not transmitting.
- *The degree of similarity between the client's channel and the sender's channel.* The more orthogonal the sender of interest's channel is to each of the interferers' channels, the less of the sender's signal Silencer will null along with those of the interferers. Fortunately, since CoS uses eight antennas and complex-valued channel vectors, there are many more degrees of freedom than the two shown in Figure 6. We examine how well CoS can null interferers end-to-end in Section 3.
- *The time between channel estimation and interference cancellation.* One important concern is how much time the beam-steering and inteference-cancelling patterns we derive last, because their longevity, together with the time needed to measure the channel, determine CoS's overhead. In Section 3.3 we measure pattern lifetime.

3. EVALUATION

There are several key performance questions surrounding indoor interference nulling. First, by how much does SamplePhase increase received signal power and throughput on a single link? Next, how well does Silencer null interference and allow that same link to function in the presence of an interferer? For both of the preceding questions, how long do the patterns derived last? Finally, how many simultaneous interferers can CoS null? In this section, we answer these questions using experiments in a typical indoor office environment, on the 13-node testbed shown in Figure 8.

Experimental setup. Our testbed consists of three Phocus phased-array antenna nodes on which we run CoS and 10 Soekris nodes each equipped with a single omnidirectional antenna. All nodes use Atheros 5413 WiFi cards and the madwifi driver under Linux. Soekris nodes use madwifi v0.9.4, whereas the phased arrays use madwifi v0.9.2.1, including patches from the OpenWRT project (for back-ported bug fixes) and Fidelity Comtech (for antenna phase and gain control functionality). To explore many different topologies, we use Soekris nodes to transmit as either senders or interferers. Notionally, these may be thought of as omnidirectional APs.

Methodology. Our experiments run in channels one (2.412 MHz) and six (2.437 MHz) of the 2.4 GHz ISM band. Using a WiSpy[6] dBx spectrum analyzer to monitor the entire 2.4 GHz spectrum, we measured the noise floor of the network at -94 dBm throughout our experiments. We also verified

[6]http://metageek.net

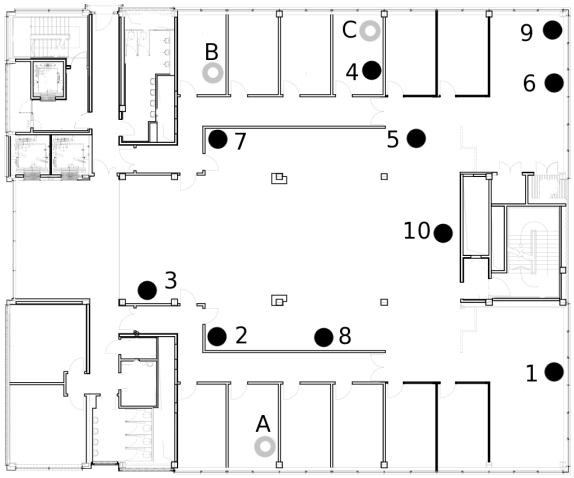

Figure 8: The indoor office environment and wireless network topology for the experiments in this paper. Filled dots represent nodes with a single omnidirectional antenna and hollow dots represent nodes with a phased-array antenna.

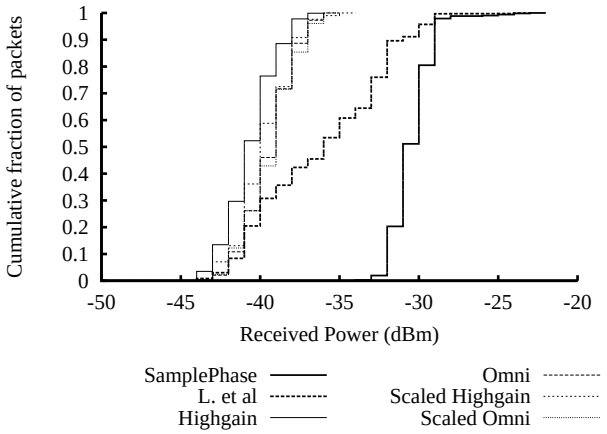

Figure 10: RSS distributions of packets drawn from fixed bit-rate (54 Mbit/s) experiments on testbed link 2A.

the presence of light background traffic from one other network being received at an average -92 dBm (measured at the middle of the testbed in Figure 8) on channel six and the occasional presence of background traffic on channel one.

Senders in our throughput experiments send 1500-byte UDP packets. For the auto bit-rate experiments, we enable bit-rate selection at senders using the madwifi implementation of the SampleRate algorithm. Experiments proceed by measuring the throughput of each of the patterns evaluated over 30 seconds.

Unless otherwise stated below, when we compare different antenna patterns, we normalize total antenna gain, running SamplePhase and using its total radiated power as a reference power level, and scaling the Silencer, directional, and omnidirectional patterns to emit the same total power. We label the latter two "Scaled Highgain" and "Scaled Omni," respectively. To put our performance into perspective, we also compare against omnidirectional ("Omni") and stock high gain ("Highgain") patterns with the highest antenna gain configurable by the user: 2.15 dBi for omnidirectional[7], and 15 dBi with a 43° beam width for high gain. Unless otherwise stated, senders and interferers in the experiments below transmit at full power (18 dBm).

Table 2 gives a roadmap for the key experiments we present in this section, and the performance gains they acheive.

3.1 Beamforming with SamplePhase

We first examine how well SamplePhase improves throughput at the receiver compared to the Lakshmanan *et al.* method and simple omnidirectional patterns. We also determine whether SamplePhase's measurements of the channel yield any throughput improvement as compared with the highest-throughput pattern among the manufacturer's fixed high-gain patterns.

We determine which high-gain pattern among the 16 possible such patterns offers the greatest throughput by empirical measurement. To do so, we interleave two iterations through the 16 directional patterns, each spaced equally around the 360° circle at 22.5° increments. We test each pattern for 30 seconds, completing the experiment in 16 minutes. We pick the stock high gain pattern that yields the best overall average throughput, and compare SamplePhase against this pattern, both at full array power (27 dBm) and scaled on a per-link basis to use the same total power as the SamplePhase pattern.

In a topology with the AP A receiving, and a single sender S sending, we compare receive throughputs at the AP using a SamplePhase-derived pattern steered toward S, the best high-gain pattern for S, the scaled best high-gain pattern for S, an omni pattern, and a scaled omni pattern.

For many different two-node (S,A) topologies, we measure receive throughput at A for the above patterns. Each measurement we report in Figure 9 is the mean of 10 one-minute measurements with error bars representing 95% confidence intervals. We re-run the optimization at the start of each one-minute measurement interval.

Figure 9 presents the main SamplePhase throughput result, in which we compare patterns generated by the SamplePhase and the Lakshmanan *et al.* measurement methods, described in Section 2.3, with the high gain and omnidirectional antenna patterns described above. From the figure, we see that in the absence of interference, SamplePhase offers the greatest throughput.

Two factors explain this throughput improvement. The first is that SamplePhase derives patterns that maximize RSS better than other methods. In Section 2.3, we present microbenchmarks that show that SamplePhase derives patterns that maximize element-pairwise RSS better than competing methods. In order to show that SamplePhase's patterns im-

[7]The vendor does not provide a figure for gain relative to an isotropic antenna. The figure above is based on the assumption that the antenna in omnidirectional mode acts as a half-wavelength dipole.

Experiment	Conclusion or performance improvement	Discussed in
SamplePhase throughput	Over many links sending one at a time in our testbed, SamplePhase improves throughput over an omnidirectional pattern by 1.5–89%, improves throughput over the best directional pattern by 6.4–219%, and improves throughput over the Lakshmanan *et al.* method by 4.3–124%.	§3.1
Silencer throughput	Over many testbed links with an interferer placed near the receiver of each, CoS improves throughput over an omnidirectional pattern by 40–1013% and improves throughput over the best stock directional pattern by 31–222%. Silencer can also null traffic on adjacent WiFi channels.	§3.2
Longevity of patterns	In a busy indoor office environment, CoS patterns work effectively for on the order of 10 hours during quiet times and two hours during busy times.	§3.3
Nulling many interferers	Over many testbed links, Silencer can null two simultaneous interferers (a total of three concurrent transmissions), acheiving throughput gains of 3.1× over the best omnidirectional pattern with the same interference. Nulling distinct interferers yields additive throughput gains.	§3.4

Table 2: Summary of experimental results for the proposed techniques, and the corresponding conclusion or performance improvement.

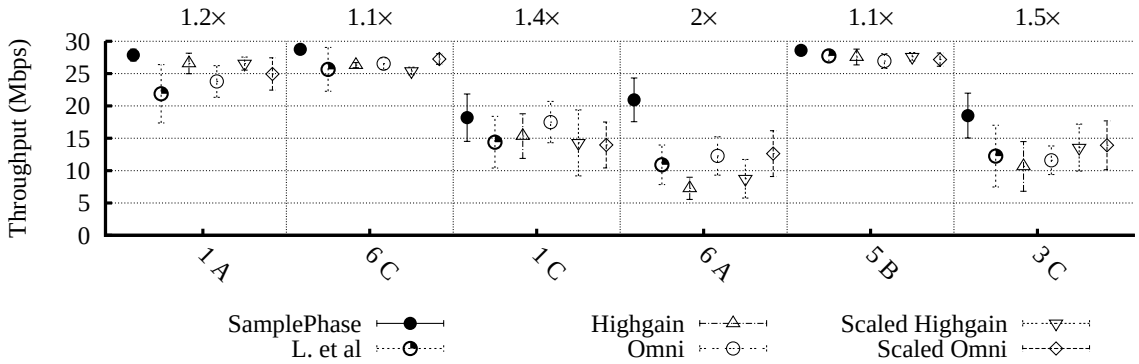

Figure 9: Empirically measured throughput for six different antenna gain patterns across six different testbed links (labeled by sender identifier and access point identifier: *cf.* Figure 8 on p. 7). In the absence of interference, SamplePhase offers the greatest throughput because of the accuracy of its channel measurement method.

prove RSS over other antenna patterns, we ran a fixed bit-rate microbenchmark over a link with high packet delivery rate at that bit-rate. This microbenchmark eliminates a sampling bias favoring strongly received packets that we would introduce if we examined RSS readings from the auto bit-rate experiments in this section. In Figure 10 we see that these pairwise gains in RSS translate into improved RSS for the pattern as a whole, which increases the SNR of received packets, making successful decoding more likely.

The second phenomenon that explains the throughput improvement in Figure 9 is that as a result of incurring fewer bit errors, SampleRate, the bit-rate adaptation scheme, chooses to use higher rates when using SamplePhase-derived patterns. For a representative link in our testbed, we examined SampleRate's data structures over a 30-second representative period in the middle of our throughput experiment. Figure 11 shows the fraction of packets that SampleRate chooses to send at each bit rate over this link. From Figure 9, we see that with the SamplePhase pattern, SampleRate chooses the top bitrate (54 Mbit/s) for slightly more than 42% of all packets it sends and either of the top two highest (54 and 48 Mbit/s) for a total 80% of all packets sent. None of the other measurement methods chooses the top two bitrates for more than 40% of all packets sent. Omnidirectional patterns use 24 Mbit/s and lower bit-rates for half of all packets.

3.2 Nulling Interferers with Silencer

802.11 networks operate at relatively high SNRs, so the cause of poor performance is often interference. In this experiment, we test how well Silencer nulls a single interferer, the most common case in a light to moderately loaded network. In a topology with an AP A receiving, a sender S sending, and an interferer I interfering, we measure the improvement in receive throughput at the AP using a Silencer pattern compared to those of a SamplePhase-derived pattern steered toward S, the best high-gain pattern for S, the scaled best high-gain pattern for S, an omni pattern, and a scaled omni pattern.

For several three-node (S, I, A) topologies, we measure receive throughput at A for the above patterns. Each measurement we report in Figure 12 is the mean of 10 one-minute measurements with error bars representing 95% confidence intervals. We re-run the optimization for each one-minute

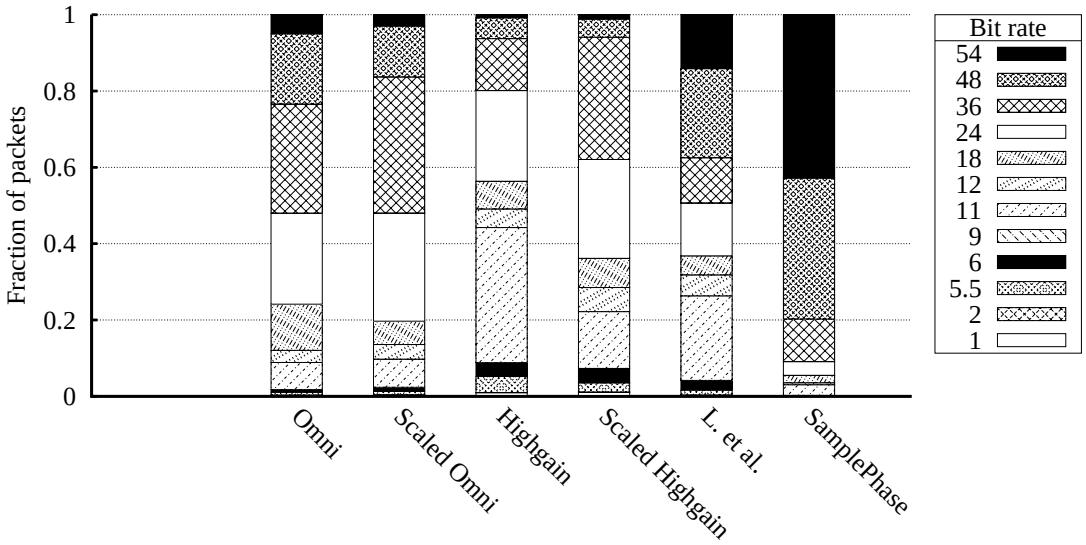

Figure 11: Bit-rates chosen by SampleRate during a representative 30-second interval at link 6A (*cf.* Figure 9). SamplePhase chooses significantly higher bit rates than other methods, choosing the top two bit-rates (48 and 54 Mbit/s) 80% of the time.

Topology	Sender (dBm)	Interferer (dBm)
3 A 6	18	18
7 C 8	18	18
4 B 1	12	18
5 B 2	12	18
7 C 1	12	18
5 B 10	12	6
7 A 8	18	6
3 C 6	18	18
3 A 6*	12	18

Table 3: Single-interferer experiments: power levels used at sender and interferer.

measurement interval. The limited physical extent of our testbed makes it difficult to place sender-interferer pairs such that neither senses the other's carrier. Therefore, we use the `TX_STOMP` register of the Atheros chipset to run experiments with carrier sense turned off at both *S* and *I*. Doing so yields more flexibility to try different power levels (shown in Table 3) at the sender and interferer, to more broadly explore how Silencer performs at the AP. Because the path between sender and interferer is distinct and independent of the sender-AP and interferer-AP paths, it is reasonable to turn off carrier sense in order to emulate topologies in which the sender and interferer are mutually hidden terminals.

In all experiments, the interferer sends broadcast packets at 54 Mbit/s.

On the final link (labeled "3A6*") in Figure 12, we test the ability of Silencer to null interference on an adjacent channel. For these data points, we tune the interferer's radio to WiFi channel 7, leaving the sender on channel 6. We allow SamplePhase to send measurement probes on channel 7 to the interferer, and on channel 6 to the sender. We then offer the resulting patterns as inputs to Silencer in the usual way. Once Silencer has generated a pattern that nulls the interferer, we tune the AP to channel 6. We note that Silencer still effectively nulls interference and increases the throughput of the link, despite combining patterns generated by SamplePhase on adjacent channels.

Penalty associated with nulling interferers. Because CoS *always* nulls inteferers of which it is aware, but interferers do not send during every packet time, there may be an opportunity cost associated with "needless nulling" (of an interferer not sending). That is, nulling an interferer may collaterally also reduce the signal strength from the sender of interest. To evaluate whether such an effect noticeably reduces throughput, we performed an experiment in which we compared the throughput achieved by a pattern derived with SamplePhase with that of a pattern derived using Silencer. In the latter case, we used the same sender as in the former, with the addition of an interferer in order to calculate the Silencer-derived pattern. We then measured the throughput achieved by the same sender in the absence of interference when receiving using these two patterns. The result was a 0.1% reduction in throughput for Silencer, suggesting that nulling a non-active interferer may not reduce a sender's throughput significantly.

3.3 Longevity of Interference-Nulling Patterns

In this experiment, we ask the following question: How long can we reasonably expect to be able to use a throughput-maximizing pattern before changes in the channel cause the pattern's performance to degrade significantly? In other words, how do received power and throughput evolve over time when a CoS-enabled AP receives after deriving beam patterns using SamplePhase and Silencer?

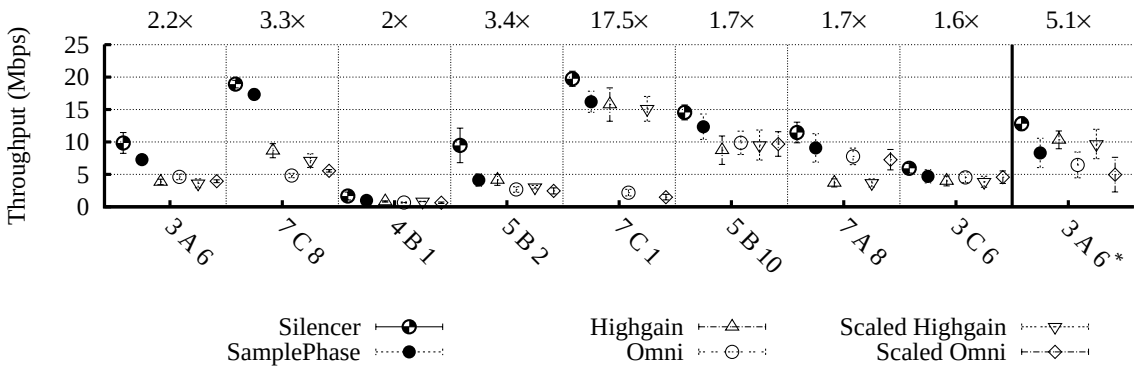

Figure 12: Empirically measured throughput in the presence of one interferer, for six different antenna gain patterns across six different testbed links (labeled by sender identifier, access point identifier and interferer identifier: *cf.* Figure 8 on p. 7). By nulling, Silencer yields the best end-to-end throughput in the presence of interference.

For *one link* on which Silencer provides a throughput gain over omni and high gain patterns, we receive using both the Silencer-derived and high gain directional patterns. In Figure 13 we plot a time series of the thirty-second moving average of throughput. We interleave throughput measurements of Silencer, a SamplePhase pattern that does not attempt to null, the high gain directional pattern that achieves the highest throughput on the link, the omnidirectional pattern and power-normalized versions of the last two. We can see that the pattern derived by Silencer outperforms all the others consistently for about 45 minutes, after which its performance degrades roughly to theirs.

3.4 Nulling Multiple Interferers

Since multiple simultaneous interferers are common in densely deployed wireless networks, Silencer's performance when more than one interferer sends concurrently is paramount. We now answer this question experimentally.

For five links in our testbed, we set up two interferers and a sender, for a total of three senders, each of which transmits packets simultaneously, as fast as possible, with carrier sense disabled as in the single-interferer experiments in Section 3.2. Over all links evaluated in this experiment, the sender transmits at 18 dBm power and the interferers transmit at 12 dBm. We compare the throughputs of a CoS-enabled AP in omnidirectional mode, Silencer nulling one of the interferers, and Silencer nulling both interferers. The results of this experiment appear in Figure 14. We see that Silencer nulling one interferer while not taking into account the other offers substantial throughput gains over an omnidirectional pattern, doubling throughput on some links. Furthermore, when Silencer measures the channel to both interferers and nulls both, it achieves additional substantial throughput gains over Silencer nulling just one interferer.

4. RELATED WORK

Beamforming—shaping the transmit or receive patterns of a multi-antenna array to maximize signal strength between a sender and receiver—is a well-known communications technique for multipath communication channels, but has only recently been investigated in local area wireless networks. In contrast, base stations in mobile telephone networks perform transmit beamforming on the downlink and receive beamforming on the uplink in order to multiplex transmissions. This is made possible in part by well-planned cellular topologies where interference is carefully managed.

In local-area wireless and mesh networks, recent MIMO 802.11n chipsets perform receive beamforming. As these ASIC-based implementations have direct access to physical-layer information, they can estimate the channel for each OFDM subcarrier indepedently, and beamform independently for each subcarrier. That approach allows more effective maximization of received signal strength than either SamplePhase or the earlier method of Lakshmanan *et al.* can achieve, as these latter two techniques only observe channel measurements from a commodity 802.11 card's per-packet RSSI measurements, which are averaged across all OFDM subcarriers.

Ruckus Wireless, a startup company, manufactures ZoneFlex APs that perform transmit beamforming to maximize receive signal strength at clients. While the algorithmic details of the techniques used by these products are proprietary, an examination of the marketing literature on the company's web site [8] strongly suggests they do no explicit nulling of any kind, neither in the transmit nor receive direction.

DIRC [6] increases indoor network capacity by having APs transmit directionally, but *always receives using only an omnidirectional antenna pattern.* DIRC further is intended for use in an enterprise setting where all wireless infrastructure is controlled by one authority, as it centrally schedules all APs' transmissions to avoid causing interference. By contrast, CoS is intended to mitigate interference in dense "chaotic" deployments, where APs run by non-cooperating users interfere. While based on the same Phocus array antenna platform as CoS, DIRC's APs perform no beamforming of any kind when they transmit—neither to maximize signal strength nor to null. Instead, DIRC's APs transmit using 16 fixed, manufacturer-supplied "high-gain"

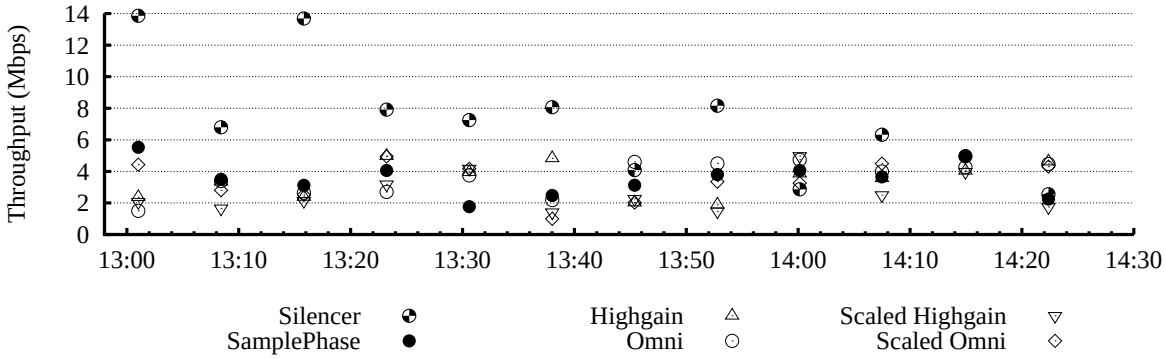

Figure 13: Longevity of a Silencer interference-nulling pattern. We run the SamplePhase measurement method only once, at the beginning of this time series. Then, as the wireless channel changes, we measure throughput in the presence of a continuous interferer. We see that the Silencer pattern's throughput does not degrade significantly for approximately 45 minutes.

patterns, each with one main lobe and several side lobes. These patterns are identical to the "high-gain" patterns used in the evaluation of CoS. CoS and DIRC are complementary: because DIRC identifies APs whose transmissions interfere at receivers and prevents those APs from transmitting concurrently, we believe that DIRC could potentially experience improved transmit concurrency by nulling toward unintended *receivers* during transmission, using patterns produced by CoS's Silencer.

The Phocus array antenna has also found use for multicast [9] and vehicular network access [7]. Like DIRC, both these systems use only the manufacturer's 16 fixed, high-gain patterns; neither does any beamforming or nulling.

Space-Division Multiple Access (SDMA) [11], which allows a receiver equipped with multiple antennas and multiple radios to receive multiple concurrent packets successfully, has been a topic of investigation in the communications theory community since the late 1990s, and has been applied to mobile telephone base stations in the past decade. More recently, Tan *et al.* [10] have combined SDMA with successive interference cancellation (SIC) [11] for use with packetized data, and successfully decoded two concurrent transmissions in a dual-antenna MIMO 802.11 receiver implemented atop a software radio platform. (They also describe techniques intended to decode more than two concurrent transmissions.) SDMA is a good fit in settings where a base station must receive data from many clients. In the dense, "chaotic" environments CoS targets, where APs run by different users interfere, however, data from an interferer on one network is of no interest to users on another, and the computational cost of decoding many packets is unattractive. CoS instead opts for an approach that nulls multiple interferers with eight antennas and a single radio, decoding only the packet from the sender of interest.

Gollakota *et al.* [3] describe a MIMO interference mitigation system in which APs align their concurrent transmissions in time, and apply SIC by exchanging information over a wired Ethernet. These techniques fit the case in which all interfering APs are controlled by a single authority, but are ill-suited to today's commmon dense 802.11 deployments, where many independent administrators operate APs that intefere. Again, CoS nulls interfering APs without requiring coordination among them, and thus is well suited to dense deployments of uncoordinated 802.11 networks.

Finally, Lakshmanan *et al.* [5] implement transmit beamforming using MRC for the Phocus phased array antenna, the identical experimental platform we use for CoS. Their technique inspired SamplePhase, but it does not explicitly null interferers; *i.e.,* their technique includes no functionality analogous to CoS's Silencer. SamplePhase uses a randomly chosen base element in its measurements to avoid consistently using any one base element that cannot successfully transmit to a client. The method of Lakshmanan *et al.* always uses the same base element, and thus is vulnerable to the aforementioned pathology. One subtle but significant practical difference between SamplePhase and Lakshmanan *et al.*'s method concerns the complexity of the two algorithms' measurements. Essentially, SamplePhase can statically compute all antenna patterns used during its measurements for deriving an MRC pattern, while the method of Lakshmanan *et al.* cannot. This difference arises because the latter method cannot analytically determine the correct sign of the phase difference between two elements in an MRC pattern; it must instead experimentally compare the received power levels achieved with phase differences of opposite sign. The Lakshmanan *et al.* method therefore requires two sequentially dependent measurement steps, the second of which uses antenna patterns known only after the first step. Because the Phocus array platform incurs significant delay when being configured with new patterns (on the order of 250 ms for each pattern configured), the extra measurement step of the Lakshmanan *et al.* method is a significant cost.[8] As a consequence, SamplePhase will in many cases be able to produce an MRC receive pattern using mea-

[8]Note that this delay is only for *configuring* a newly defined pattern into the array; switching among previously configured patterns takes only 120 μs.

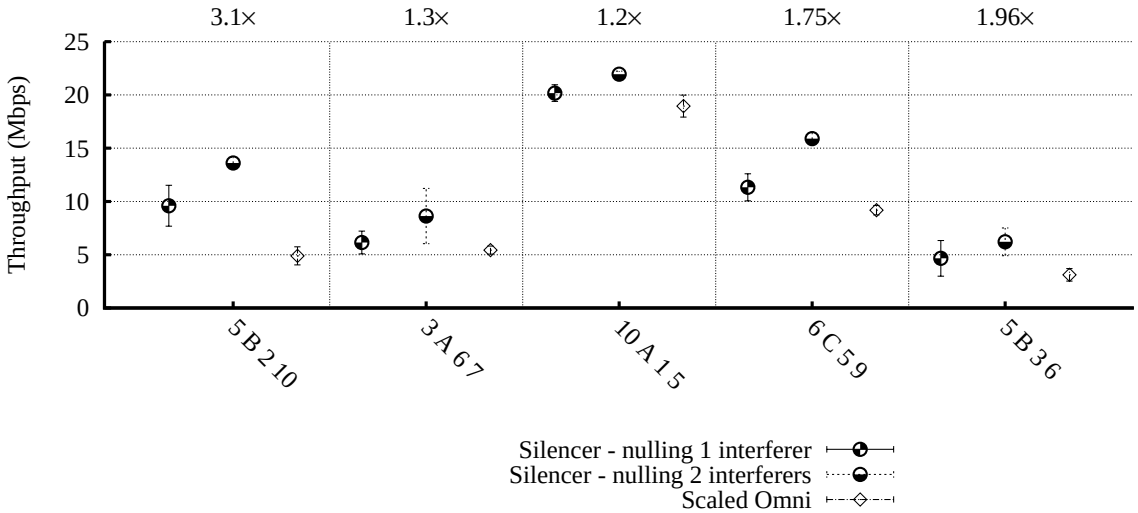

Figure 14: Empirically measured throughput in the presence of two concurrent interferers (for a total of three concurrent transmissions), over multiple testbed links (labeled by sender identifier, access point identifier and both interferers' identifiers: *cf.* Figure 8 on p. 7). Adding channel measurements for each additional interferer and incorporating that interferer into Silencer's nulling algorithm yields cumulative additional benefit.

surements over a shorter period that is more likely to fall within the channel's coherence time.

5. CONCLUSION

The broad adoption of 802.11 wireless networks forces an ever-increasing population of co-located users to share finite unlicensed spectrum. Much of this sharing occurs among distinct networks deployed in uncoordinated fashion in close proximity. Capacity-enhancing proposals for enterprise wireless networks, in which a large organization wishes to increase capacity in a network of many centrally controlled APs, have profited greatly from multi-antenna techniques in recent years, but assume cooperation among APs.

In this paper, we have presented Cone of Silence (CoS), a system that leverages multiple antennas to increase receive throughput at an AP in the presence of interference from uncooperative nearby senders. CoS adopts long-known techniques for multi-antenna systems, maximum ratio combining (MRC) and the decorrelator, but applies them in the novel context of an 802.11 receiver with a single, commodity radio. CoS's decorrelator, Silencer, nulls multiple interferers without decoding their signals, while maximizing signal strength from a sender of interest. Our experimental prototype CoS 802.11b/g access point equipped with an 8-element phased array antenna demonstrates a throughput improvement under interference of between 1.6× and 17.5× over that achieveable with an omnidirectional antenna, and achieves consistently higher throughput than beamforming toward the sender alone. SamplePhase and Silencer achieve throughput gains by allowing automatic bit-rate adaptation on clients to choose higher bit-rates for more packets; they increase SINR at the AP, improving the reliability of higher bit-rates.

We believe the CoS design holds promise for implementation in commodity 802.11 interfaces, which today do not null inteferers. CoS also suggests useful future enhancements to the host-wireless PHY and MAC interfaces. Making per-OFDM subcarrier signal strength measurements available to the host would allow the host to perform higher-fidelity CoS (both SamplePhase and Silencer) in software. An open question is how quickly per-subcarrier signal strengths change *vs.* host CPU processing speeds. Even with a hardware implementation of CoS in the 802.11n ASIC, one could enhance the host-wireless MAC interface to allow the host CPU to specify per-packet which remote MAC addresses to null toward. Such an interface might provide a powerful primitive for building a nulling-enabled MAC protocol.

6. REFERENCES

[1] AKELLA, A., JUDD, G., SESHAN, S., AND STEENKISTE, P. Self-management in chaotic wireless deployments. In *MobiCom* (Aug. 2005).

[2] FIDELITY COMTECH, INC. Phocus Array 3100X v2.1 data sheet. `http://www.fidelity-comtech.com/PDFs/DS%20Phocus%203100x%20v2.1%20v4%20Final.pdf`.

[3] GOLLAKOTA, S., PERLI, S., AND KATABI, D. Interference alignment and cancellation. In *SIGCOMM* (Aug. 2009).

[4] IEEE Standard 802.11-2007: Wireless LAN MAC and PHY Specifications, June 2007.

[5] LAKSHMANAN, S., SUNDARESAN, K., RANGARAJAN, S., AND SIVAKUMAR, R. Practical

$$\hat{\theta}_{kl} = \arctan\left\{\frac{\sum_{i=1}^{n}\left[\left(\zeta\left(\frac{2\pi i}{n}+\phi\right)-\zeta\left(\frac{2\pi i}{n}+\pi+\phi\right)\right)\cos\left(\frac{2\pi i}{n}+\phi\right)-\left(\zeta\left(\frac{2\pi i}{n}+\phi\right)-\zeta\left(\frac{2\pi i}{n}+\pi+\phi\right)\right)\sin\left(\frac{2\pi i}{n}+\phi\right)\right]}{\sum_{i=1}^{n}\left[\left(\zeta\left(\frac{2\pi i}{n}+\phi\right)-\zeta\left(\frac{2\pi i}{n}+\pi+\phi\right)\right)\cos\left(\frac{2\pi i}{n}+\phi\right)+\left(\zeta\left(\frac{2\pi i}{n}-\frac{\pi}{2}+\phi\right)-\zeta\left(\frac{2\pi i}{n}+\pi+\phi\right)\right)\sin\left(\frac{2\pi i}{n}+\phi\right)\right]}\right\} \quad (4)$$

Figure 15: The generalized SamplePhase estimator, sampling n equally-spaced element phase differences with an initial offset of ϕ.

beamforming based on RSSI measurements using off-the-shelf wireless clients. In *IMC* (Nov. 2009).

[6] LIU, X., SHETH, A., KAMINSKY, M., PAPAGIANNAKI, K., SESHAN, S., AND STEENKISTE, P. DIRC: Increasing indoor wireless capacity using directional antennas. In *SIGCOMM* (Aug. 2009).

[7] NAVDA, V., SUBRAMANIAN, A., DHANASEKARAN, K., TIMM-GIEL, A., AND DAS, S. Mobisteer: Using steerable beam directional antenna for vehicular network access. In *MobiSys* (June 2007).

[8] RUCKUS WIRELESS, INC. ZoneFlex Mid Range product description. `http://www.ruckuswireless.com/products/zoneflex-mid-range`.

[9] SEN, S., XIONG, J., GHOSH, R., AND CHOUDHURY, R. Link layer multicasting with smart antennas: No client left behind. In *ICNP* (Oct. 2008).

[10] TAN, K., LIU, H., FANG, J., WANG, W., ZHANG, J., AND VOELKER, G. SAM: Enabling practical spatial multiple access in wireless LAN. In *Proc. of the ACM MobiCom Conf.* (Beijing, China, Sept. 2009), pp. 49–60.

[11] TSE, D., AND VISWANATH, P. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

APPENDIX

A. THE SAMPLEPHASE ESTIMATOR

Starting from Equation 2, define

$$\zeta(\delta) = \frac{P_{kl}(\delta) - (P_k + P_l)}{2\sqrt{P_k P_l}}. \quad (5)$$

For a given base element, we program the AP to transmit with two elements active for $\delta = \left\{-\frac{\pi}{2}, 0, \frac{\pi}{2}, \pi\right\}$, and compute values for $\zeta\left(-\frac{\pi}{2}\right)$, $\zeta(0)$, $\zeta\left(\frac{\pi}{2}\right)$, $\zeta(\pi)$ directly from multiple empirical RSS measurements.

We now explain how to compute the SamplePhase estimator $\hat{\theta}_{kl}$ for the true channel phase difference between elements k and l, θ_{kl}. Define the measured error at angle δ, $\varepsilon(\delta)$, as follows:

$$\varepsilon(\delta) = \zeta(\delta) - \cos(\theta_{kl} + \delta). \quad (6)$$

Using the least-squares method, we fit the measured values $\zeta(\cdot)$ to a sinusoidal function. Let $S\left(\hat{\theta}_{kl}\right)$ be the sum of the squared errors $\varepsilon(\cdot)$, as a function of the SamplePhase estimator:

$$S\left(\hat{\theta}_{kl}\right) = \varepsilon^2\left(-\frac{\pi}{2}\right) + \varepsilon^2(0) + \varepsilon^2\left(\frac{\pi}{2}\right) + \varepsilon^2(\pi). \quad (7)$$

Using Equation 6 to expand Equation 7 and applying trigonometric identities, we find

$$\begin{aligned} S\left(\hat{\theta}_{kl}\right) = {} & \zeta^2\left(-\frac{\pi}{2}\right) + \zeta^2(0) + \zeta^2\left(\frac{\pi}{2}\right) + \zeta^2(\pi) \\ & + 2 - 2\left[\zeta\left(-\frac{\pi}{2}\right)\sin\left(\hat{\theta}_{kl}\right) + \zeta(0)\cos(\hat{\theta}_{kl})\right. \\ & \left. - \zeta\left(\frac{\pi}{2}\right)\sin(\hat{\theta}_{kl}) - \zeta(\pi)\cos(\theta_{kl})\right]. \end{aligned} \quad (8)$$

In order to minimize S, we seek the zeros of S':

$$\begin{aligned} S'\left(\hat{\theta}_{kl}\right) = {} & -2\zeta\left(-\frac{\pi}{2}\right)\cos\left(\hat{\theta}_{kl}\right) + 2\zeta(0)\sin\left(\hat{\theta}_{kl}\right) \\ & + 2\zeta\left(\frac{\pi}{2}\right)\cos\left(\hat{\theta}_{kl}\right) - 2\zeta(\pi)\sin\left(\hat{\theta}_{kl}\right) \\ = {} & 0 \\ \therefore \hat{\theta}_{kl} = {} & \arctan\left(\frac{\zeta\left(-\frac{\pi}{2}\right) - \zeta\left(\frac{\pi}{2}\right)}{\zeta(0) - \zeta(\pi)}\right) \end{aligned} \quad (9)$$

where arctan is generalized to have exactly one root in $[0, 2\pi)$ given the numerator and denominator in Equation 9.

From the above derivation, we can see that the it does not matter which angles are used evaluate ζ, so long as these angles are spaced evenly, with a $\frac{\pi}{2}$ difference between them.

Now, suppose we have n quadruplets of evenly spaced measurements starting at an arbitrary phase ϕ. Using further trigonometry, Equation 9 fully generalizes to Equation 4 (Figure 15).