COPA: Cooperative Power Allocation for Interfering Wireless Networks

Georgios Nikolaidis, Mark Handley, Kyle Jamieson[*] and Brad Karp
Dept. of Computer Science, UCL

ABSTRACT

As 802.11 wireless networks proliferate, interference becomes increasingly severe, particularly in dense, urban environments. These networks are usually operated by different users (*e.g.,* tenants in apartments). In this paper, we develop techniques for mitigating interference between such *loosely cooperating* 802.11 MIMO APs and clients, which do not share a high-speed wired backplane or central controller. We propose CoOperative Power Allocation (COPA), an approach to concurrent wireless medium access that combines fine-grained, per-subcarrier power allocation, nulling, and multi-stream transmission to claim capacity that status-quo approaches cannot. Jointly turning these knobs allows COPA to allocate subcarriers to senders partially, rather than all-or-nothing, and to embrace a measure of interference when doing so increases capacity.

CCS Concepts

•Networks → Wireless local area networks;

Keywords

MIMO, nulling, interference, power allocation, OFDM

1. INTRODUCTION

Deployments of Wi-Fi wireless LANs in homes and offices have proliferated so widely that it is now commonplace for several such networks to operate in close proximity. These dense, uncoordinated deployments often interfere significantly with one another. Of late, the research community has explored interference mitigation approaches that centrally control wireless senders' concurrent transmissions. One such approach centrally instructs separate access points (APs) run by the same organization to use their respective antennas to cancel the interference each AP might deliver to other APs' clients [6, 11, 13, 14]. Another such approach is distributed MIMO, which pools the antennas of APs run by the same organization into one large "virtual" AP that concurrently sends multiple streams of data to multiple clients [22]. These approaches show great promise in the enterprise setting where a single entity controls all APs and (in the case of distributed MIMO) connects them via a gigabit wired backplane.

Wireless LANs in separate homes and offices are typically owned and administered by different parties, where there is no common controller that can orchestrate their transmissions centrally, nor any gigabit-speed wired backplane interconnecting their APs. How can one mitigate the interference uncoordinated, densely deployed APs cause to one another's clients, and thus improve aggregate throughput? One well-known tool for interference mitigation is *nulling,* where an AP uses multiple antennas to cause multiple instances of its transmitted signal to cancel one another at another AP's client. In this paper, we explore *selfish cooperation,* where two APs run by different parties coordinate over the wireless medium to send concurrently while nulling toward one another's clients. The APs are selfish in that they may only decide to cooperate when neither suffers a reduction in throughput when they send concurrently. We demonstrate that while nulling can significantly reduce interference, the practical capacity improvement nulling alone achieves in realistic indoor environments is limited because nulling overlooks—and indeed, elevates the importance of—the complementary problem of power allocation.

We present *CoOperative Power Allocation* (COPA), a system in which two APs alert one another of the clients to which they intend to send. Each AP explicitly considers the relative strengths of both APs' transmissions at both clients, then nulls toward the other's client, cooperating with its counterpart to allocate power to each narrow sub-band within the Wi-Fi channel. The APs choose power allocations that yield a signal to interference plus noise ratio (*SINR*) at each client that is most conducive to high aggregate throughput. As we discuss in §2, subcarrier-granularity power allocation is important to throughput because frame reception fails when a Wi-Fi receiver cannot correctly decode the data on just a few of the dozens of subcarriers that comprise the Wi-Fi channel. Thus, even when receivers hear some subcarriers with very high SINR, and others with low SINR, the sender has no choice but to send with a (lower bit-rate) mod-

---

[*] Also affiliated with Princeton University.

CoNEXT '15 December 01 - 04, 2015, Heidelberg, Germany

ISBN 978-1-4503-3412-9/15/12…$15.00

DOI: http://dx.doi.org/10.1145/2716281.2836103

ulation and code that protect the data on the subcarriers with poor SINR, sacrificing available wireless capacity.

Unfortunately, as we illustrate experimentally, nulling exacerbates the variability in per-subcarrier SINR at receivers. COPA is thus an important complement to nulling: COPA APs cooperatively allocate transmit power so as to explicitly avoid low-SINR subcarriers at their receivers, and thus improve aggregate throughput. Furthermore, a cooperative system must sometimes offer an incentive for different, possibly selfish users to cooperate. COPA supports two modes of power allocation: one that aims to achieve the greatest aggregate throughput, possibly at the expense of one AP's throughput; and another that aims to improve aggregate throughput subject to the constraint that no AP suffer reduced throughput. In §4, we experimentally explore the price COPA pays in aggregate capacity to achieve fairness.

Our main contributions in this work are to elucidate the real-world performance of nulling and materially improve it. Our salient findings include:

- In 83% of topologies in an office environment in which two 4-antenna Wi-Fi APs send to two 2-antenna clients, nulling *underperforms* CSMA. On these topologies, in which the variability of SINR across subcarriers introduced by naïve nulling forces the APs to transmit at lower bit-rates, COPA improves nulling's throughput by a mean of 64%, such that the throughput of COPA's approach to nulling *exceeds* CSMA's in 76% of the same topologies.
- In the remaining 17% of the same topologies in which naïve nulling outperforms CSMA in throughput, it does so by a median of 12%. On these topologies, in which cross-interference between each AP and the other AP's client is relatively weak, but naïve nulling still introduces throughput-limiting variability of SINR across subcarriers, COPA improves nulling's throughput improvement over CSMA to a median of 45%.

2. PROBLEM

Consider the Wi-Fi deployment shown in Figure 1, where two independently operated Wi-Fi LANs are in proximity (*e.g.,* in adjacent offices or apartments). Because of the proximity of the two LANs, if AP1 transmits to C1 while AP2 transmits to C2 concurrently, each will likely interfere with the other's transmission.

The throughput of AP1's transmission to its intended receiver C1 is determined by the *signal-to-interference-plus-noise ratio* (SINR), where *signal* denotes the received power of AP1's signal at C1, *interference* denotes the received power of AP2's signal at C1, and *noise* denotes the *noise floor,* or the power level of background RF noise in the environment. Thus, from C1's perspective, stronger received interference from AP2 or a weaker received signal from AP1 reduces SINR at C1, thus reducing throughput.

In Wi-Fi, there are two main techniques to mitigate this throughput reduction: *carrier sense* and *nulling.* Carrier sense (CS) attempts to eliminate interference by avoiding concurrent transmission: a sender defers if it hears another sender transmitting. CS results in sequential transmissions, *e.g.,* in two greedy senders' taking turns transmitting on average. Given collisions and medium acquisition overheads, the aggregate throughput of CS is bounded above by that of a perfectly scheduled Time Division Multiple Access (TDMA) scheme. Nulling, on the other hand, entails concurrent transmissions by the two senders, and attempts to cancel out interference at the unintended receiver. A sender with multiple antennas can cancel its own transmission by sending one instance of its signal from one antenna and a phase-shifted instance of the same signal from a second antenna. If it chooses the phase shift so that the two signals arrive perfectly out-of-phase at the antenna of the unintended receiver, they will sum to zero.

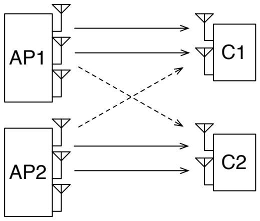

Figure 1: AP1 and AP2 transmit concurrently to clients C1 and C2, respectively; each AP interferes with the other AP's transmission to that other AP's client.

But CS and nulling share a common deficiency: Wi-Fi senders using these techniques today allocate power equally across the 20 MHz or 40 MHz frequency band of a transmission,[1] and they do so in a fashion oblivious to the detailed interference effects each sender causes to the other's receiver. To see why this matters, let us examine in greater detail how a 20 MHz Wi-Fi channel behaves in practice.

2.1 OFDM and Narrow-Band Fading

CS and TDMA both assume that concurrent transmissions by two APs will decrease SINR sufficiently that neither client can decode correctly. However, modern Wi-Fi senders first redundantly code outbound frames, then break the wideband wireless channel into some number of *subcarriers*, sending distinct "slices" of the redundantly-coded frame's bits on each subcarrier—a scheme known as Orthogonal Frequency-Division Multiplexing (OFDM). The SINR of each individual OFDM subcarrier determines the bit-error rate with which a receiver can decode information sent on that particular subcarrier. To limit hardware complexity, the Wi-Fi standard constrains a sender to use the same OFDM modulation on all subcarriers.

Indoors, it is expected that many OFDM subcarriers will suffer from *narrow-band fading* due to multipath propagation effects. Some subcarrier frequencies are affected more than others by reflection, scattering, and shadowing. In addition, the fading pattern often drastically differs at different locations separated by a distance of just 12.5 cm (one radio

---

[1] Platform limitations confine our results to 20 MHz.

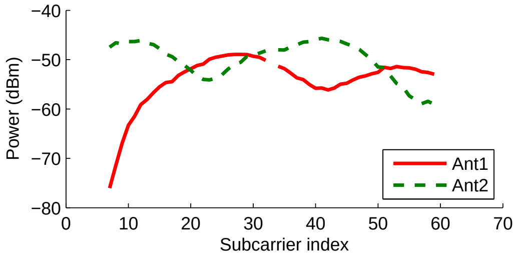

Figure 2: Received power from a single send antenna at two different receive antennas, by OFDM subcarrier index.

wavelength), as Figure 2 from our testbed shows. The sender here allocates equal power to each subcarrier resulting in a received power shown in the figure.

There is a power budget for the entire 20 MHz channel that state-of-the-art Wi-Fi OFDM senders divide equally between subcarriers, but other allocations are possible. We observe that in the presence of narrow-band fading, the ability to unequally allocate subcarrier power creates the opportunity for two APs to send concurrently to their respective clients. For example, suppose AP1's transmission on some subcarrier propagates strongly to its client C1, but much more weakly to AP2's client C2. If AP1 and AP2 transmit concurrently, AP1's transmission may cause an insignificant SINR drop on that subcarrier at C2. Even if AP2's transmission on that subcarrier causes a throughput-reducing SINR drop at C1, AP1 could potentially alleviate this by increasing the power it sends on that subcarrier.

To take advantage of such opportunities, the two APs would need to make power allocation decisions cooperatively: if both APs knew the per-subcarrier SINR of each AP at both clients, they could maximize aggregate throughput by choosing how much power to transmit on each subcarrier. Today's Wi-Fi networks and OFDMA system design proposals [20, 24] ignore other senders when they allocate power; instead we propose *cooperative power allocation.* This approach is possible even for single-antenna APs.

Prior approaches to power allocation do not fit the scenario of cooperating Wi-Fi APs. The technique of *water-filling* maximizes achieved link capacity [26] for idealized radios that transmit Gaussian signals, but performs poorly for practical radios like those used in Wi-Fi, which transmit discrete constellations. The related technique of *mercury and water filling* [15] optimally distributes power among subcarriers for discrete constellations. But it does not consider eschewing weakly received subcarriers outright. Nor does it address power allocation under dynamically changing interference, such as results when two concurrently sending APs perform power allocation toward their own clients while nulling toward each anothers' clients.

2.2 Nulling in Practice: Residual Interference

In principle, when one AP nulls toward another's client, the nulling should eliminate *all* interference from the first AP at the client's antenna. In practice however, there is residual interference left over after nulling. Furthermore, senders null on a per-subcarrier basis, and so efficacy may vary significantly from subcarrier to subcarrier, even though averaged across subcarriers, nulling reduces interference well.

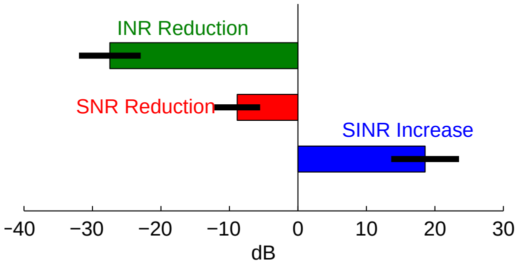

Figure 3: End-to-end effect of nulling on SINR, SNR and INR; 30 indoor office topologies; two four-antenna APs sending to two 2-antenna clients. Error bars denote one standard deviation.

Nulling viewed on average. To evaluate the efficacy of nulling we transmit concurrently from two four-antenna APs to two two-antenna clients in an indoor office environment. Each AP sends two MIMO streams to its own client and nulls toward each of the two antennas of the other client.[2]

How much does nulling improve SINR at C1 and C2? In our testbed, we take measurements at C1 and C2 in 30 different indoor office topologies, each consisting of different placements of two four-antenna APs and two clients. Our results in Figure 3 show a reduction in interference ("INR reduction" in the figure) averaging 27 dB. Although nulling reduces interference significantly, the reduction does not generally exceed -30 dB. Furthermore, nulling may also reduce the signal an AP delivers to its own client. We term this "collateral damage," as two signals intended to cancel at the other AP's client may also partially cancel at the intended receiver. Experimentally, we see that the cost of nulling the signal of interest ("SNR reduction" in the figure) averages -8 dB, offsetting the reduction in interference for a net 18 dB SINR improvement on average, with SINR improvement generally no better than 23 dB.

Nulling viewed per subcarrier. We now examine the effect of nulling on individual subcarriers at client C1 in one of our testbed topologies. The "SNR BF" curve in Figure 4 shows the SNR of each of AP1's subcarriers at C1 when only AP1 sends—this is the baseline, when AP1 has complete freedom to adjust the phase of its transmissions, thus beamforming towards C1. The "SNR Null" curve shows the result of AP1 nulling towards C2 while sending to C1. Not only has mean power decreased, but the SNR is more variable. The primary reason is that in order to null, AP1 can no longer fully align the phase of its transmissions as received at C1. The effect resembles that of narrowband

---

[2] Our APs and clients are Rice WARP v2s communicating over a 20 MHz channel in the 2.4 GHz band, using OFDM and Wi-Fi's 802.11n high-throughput bit-rates; full details of our experimental setup are in Section 4.

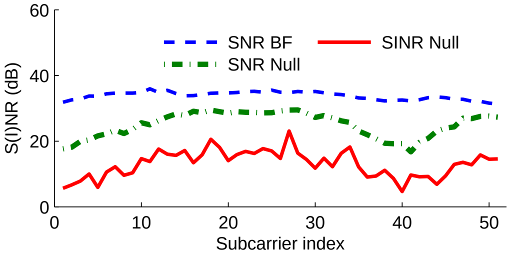

Figure 4: Per-subcarrier effects of nulling; two four-antenna APs sending to two two-antenna clients.

fading: we see an increased SNR variance across subcarriers. Lastly, the "SINR Null" curve shows the SINR C1 experiences when AP1 and AP2 send concurrently, but null towards each other's clients. In addition to the effect of imperfectly aligned phase for "SNR Null", incomplete nulling has further reduced the mean SINR and introduced further variance across subcarriers. Several noise sources conspire to cause imperfect nulling, including receiver noise when measuring the channel state in order to calculate the nulling phase and transmitter imperfections and noise when sending the nulled signal.

The consequence of an increase in SINR variance across subcarriers is often reduced throughput. Wi-Fi employs only a single modulation and convolutional code across all subcarriers, and uses all subcarriers in every packet transmission. When the SINR varies widely across subcarriers, a few low-SINR subcarriers can cause the decoding of an *entire packet* to fail, and thus cause the sender to reduce bit-rate. As Wi-Fi does not cooperatively allocate power to subcarriers by taking SINR at all receivers into account, it cannot reallocate power from good subcarriers to weak ones to "save" them from having catastrophically low received SINR.

3. DESIGN

In this section, we describe the design of COPA, starting with a description of how COPA APs learn about each other's transmissions (§3.1), followed by a description of how each COPA AP chooses the amount of power to allocate to each OFDM subcarrier (§3.2). A description of the system's overall design, including how both the above design elements integrate with transmit nulling concludes the section (§3.3).

3.1 Coordination Protocol

Since COPA APs don't necessarily belong to the same administrative domain, nor are even connected by a high speed LAN, they must coordinate over the air in order to:

1. Inform each other about opportunities for concurrency,
2. Disseminate information about the channels (*channel state information*, or *CSI*) between APs and clients, and
3. Exchange information about the choices each AP makes.

COPA strives to accomplish these goals with a minimum of protocol overhead. The basic mechanism we choose is the use of *control* messages transmitted using an omnidirectional spatial profile. We term our control message the *Intention-to-Send* or *ITS*.

Learning CSI When a COPA AP overhears frames from nearby clients or other APs, it measures the channel from those senders, and caches the resulting CSI in a table indexed by sender address. Since the wireless channel is reciprocal between sender and receiver,[3] COPA APs learn CSI information to nearby clients by overhearing their recent transmissions as shown in Step ① of Figure 5.

How recent must clients' transmissions be in order to ensure that COPA APs use accurate CSI information? CSI does not need to be refreshed at the start of every 4 ms 802.11 *transmit opportunity* (the time granularity of medium acquisition), but instead once every *coherence time*, the amount of time for which the wireless channel remains mostly constant, given the speed of nearby mobile clients, objects, and people. The coherence time is given by $t_c = \frac{m \cdot \lambda}{v}$, where λ is the wavelength of the carrier frequency, v is the speed of the host, and m is a parameter that characterizes the physical environment. A conservative value for m is 0.25 [26, 10], which results in coherence times of 28 ms for a speed of 4 km/h and 112 ms for a speed of 1 km/h. In our experimental evaluation (§4), we measure CSI once every 30 ms—sufficiently often for COPA to work in an environment with a coherence time of that short a duration.

COPA compresses CSI information and precoding matrices using adaptive delta modulation across subcarriers' amplitude and phase (separately), and compressing the result using a lossless variant Lempel-Ziv data compression algorithm. This yields a compression ratio of two on average for the channels in our testbed (§4).

Finally, to avoid interference between different OFDM carriers, concurrent transmissions need to be synchronized in time within a cyclic prefix (800 ns). This is possible in today's Wi-Fi medium access control protocol by senders timing their transmissions off the end of ongoing transmissions, and COPA leverages the same mechanism in its medium access control protocol, which we now describe.

The ITS Exchange When traffic from the wired backhaul arrives at the COPA AP for downstream transmission to a client, the AP first checks whether there are any ongoing COPA transmissions, and obeys the status quo Wi-Fi carrier sense deference rule of waiting until ongoing transmissions complete and then competing for the medium by means of a bounded exponential backoff. Once the medium becomes clear, all APs ready to send traffic to clients then contend to send an ITS INIT control frame as shown in Step ② of Figure 5. ITS INIT expresses an AP's intent to send to a specified receiver, and the AP that wins the contention is considered elected as *Leader AP*. The ITS INIT contains both the Leader AP's identity and the identity of the client to which the leader AP is about to send (Client 1 in our example).

After the ITS INIT frame, any other APs that have traffic to send to their clients then cancel the transmission of their

---

[3] Except for radio front-end differences, which modern Wi-Fi radios calibrate away before operation [1].

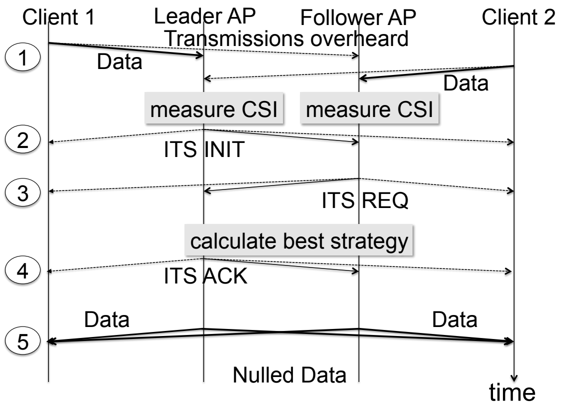

Figure 5: A timeline of COPA MAC operation. ITS REQ frames include CSI from Follower AP to both clients, while ITS ACK frames include the precoding matrix for the Follower AP.

ITS INIT frames and contend to send a frame that indicates that the AP (which we term a *Follower*) requests to participate in the next transmission opportunity with the Leader. This frame is called the ITS *request* (ITS REQ) frame and is shown in Step ③ of our example. The ITS REQ frame contains the identities of the Leader, Follower, and both clients, and the CSI from Follower to Client 1 and Client 2.

Once the Leader receives an ITS REQ frame from a Follower, it estimates the best joint strategy for both APs. This computation is described in the next two sections (§3.2, §3.3). Once it has made its choices, the Leader AP sends an ITS ACK (Step ④) containing all four parties' identities, a field indicating either that the Leader has decided that the two APs would be best off taking turns in time or transmitting concurrently. In the first case, if one of the two APs wins the initial contention, the other does not send an ITS REQ back for the rest of the coherence time. This means that the first AP has the opportunity to either engage in an ITS exchange with some third AP, or alternatively transmit on its own. On the other hand, if the two APs decided to transmit concurrently, the ITS ACK also includes the precoding matrix that the Follower should use for a concurrent transmission. Finally both APs transmit: concurrently if the calculation shows that to be the best strategy, or sequentially if no good concurrent solution is available.

Discussion: Other MAC Considerations Since COPA's ITS exchange relies on probabilistic contention between senders, collisions are possible, which result in garbled ITS frames. In these cases, senders follow the standard bounded exponential backoff and retry their transmissions.

In the experimental evaluation we conduct in this paper, we limit ourselves to cases with two sending APs. When there are more than two senders present, the fairness of COPA's ITS mechanism merits consideration. When two COPA senders elect to send sequentially after an ITS exchange, they implicitly win two consecutive contention rounds. To avoid unfairness to other senders who may be present, after two COPA senders send sequentially, they should defer to other senders in the immediately following contention round by using a modified contention window of $[\mathrm{aCWmin}+1, 2\cdot\mathrm{aCWmin}+1]$, rather than the default of $[0, \mathrm{aCWmin}]$. We expect this deference to improve fairness; we leave an evaluation of this modification to future work.[4]

Coherence time (ms)	COPA Conc (%)	COPA Seq (%)	CSMA CTS (%)	CSMA RTS/CTS (%)
4	9.3	7.7	2.7	3.7
30	5.1	3.5	2.7	3.7
1000	4.5	2.8	2.7	3.7

Table 1: Throughput costs incurred by MAC overhead.

All ITS control packets contain an *airtime* field indicating the duration on the wireless medium of the data each AP will send to its client. To nearby radios not participating in the coordinated transmission, ITS control packets thus function in the same way as 802.11's RTS/CTS exchange: other radios defer for the duration of the coordinated transmission, even if overhearing only one side of the ITS exchange.

The ITS exchange adds overhead beyond that of CSMA. The magnitude of this overhead depends on how often COPA must disseminate CSI (which in turn depends on the coherence time of the environment), and on whether COPA decides on sequential or concurrent transmission. Table 1 analytically compares the throughput costs incurred by COPA's ITS exchange in the sequential and concurrent transmission cases, *vs.* by CSMA's CTS-to-self and CSMA's RTS/CTS, for different coherence times.

3.2 Per-Subcarrier Power Allocation

Given precoding matrices and CSI, COPA calculates the expected SINR at both clients, on every OFDM subcarrier: this depends on how much power each AP sends on each subcarrier. COPA's goal, however, is not necessarily to maximize average SINR, but instead to maximize throughput. Current hardware constrains us to using a single decoder at the receiver, so subcarriers with a very poor SINR may cause a high bit error rate (BER) at the receiver, even if most of the subcarriers have good SINR.

To prevent bad subcarriers from causing bit errors, COPA simply drops them. It does so by indicating to the receiver in the A-MPDU's preamble which subcarriers to attempt to decode. Dropping subcarriers mitigates decoding errors, but also frees power to be added to other subcarriers (improving the bitrate achievable on them), and allows a concurrent sender to use the wireless capacity interference-free.

Although it is not possible for Wi-Fi hardware to radiate zero power on a subcarrier—the typical carrier leakage from adjacent subcarriers given by the Maxim 2829 datasheet is -27 dB [2], similar to what we experimentally observed—

---

[4] We expect this change either to negligibly reduce throughput or possibly even improve it, owing to a possible decrease in collisions. Note that this modification only takes place after a full ITS exchange, since for the rest of the coherence time the two hosts do not engage in further ITS exchanges.

dropping a subcarrier leads to a drop in interference, leaving that subcarrier free for other hosts to use. This in combination with nulling by the new user of the subcarrier so that it doesn't cause cause problems at the AGC of the original receiver allows for better use of available frequencies.

For a single AP transmitting to a single client without interference from a concurrent sender, the procedure is shown in Algorithm 1. We call this algorithm *Equi-SNR* as it equalizes the received SNR on all subcarriers not discarded.

Algorithm 1: Power allocation for one MIMO stream.

```
Sort the subcarriers into order of increasing SNR
for i in range(0, NUM_SUBCARRIERS) do
    Allocate no power to the first i subcarriers
    Allocate power to remaining subcarriers so as to
    equalize their SNR
    Calculate max achievable 802.11 modulation
    Calculate throughput given modulation and number
    of subcarriers used
end
Use number of subcarriers that maximizes throughput
```

3.2.1 Concurrent subcarrier power allocation

When two APs transmit concurrently to two clients, whether performing nulling or not, the task becomes much more complicated. For each AP we can calculate the subcarriers used and power allocation for those subcarriers, as above. However, we must then take into account the interference caused by this choice of power allocation. To illustrate the point consider the following scenario:

> We've decided AP2 won't use a subcarrier because of interference from AP1. This reduces interference at Client 1, so AP1 can reduce the power it uses on that subcarrier to improve others. Since AP1 reduces the power used, the SINR at Client 2 improves, so we now decide AP2 can use that subcarrier after all. To get acceptable SINR at Client 1, AP1 must now increase the power on that subcarrier, reducing power on others to compensate.

Since total power is limited, any change in power allocation to one subcarrier requires adjusting all others, changing interference on all subcarriers at the other AP. It should be clear that an optimal solution requires considering all possible power allocations and combinations of used subcarriers between the two APs, which becomes impractical.

To achieve an acceptable heuristic solution, we use the processing shown in Figure 6. Based on the precoding matrix and CSI, we first calculate a power allocation solution *independently* for each MIMO stream between AP1 and Client 1 and between AP2 and Client 2. This initial power allocation assumes a per-subcarrier interference based on the other sender allocating its transmit power equally across all subcarriers. After this initial allocation, we compute the revised interference each stream causes to all the other streams. We feed this interference matrix back into the allocation algorithm and recompute which subcarriers should be used and how to allocate power between them. The process iterates until it converges or an iteration limit is reached. The iteration is not guaranteed to reach a global maximum, and because of independent allocations to each stream, it may occasionally regress from the best solution, in which case we choose the best solution previously found.

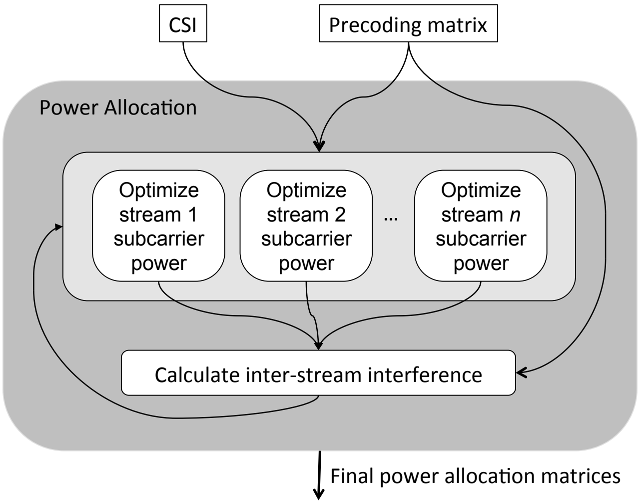

Figure 6: COPA's *Equi-SINR* iterative power allocation.

This algorithm resembles Algorithm 1 (Equi-SNR), except it equalizes SINR rather than SNR to take account of this inference, so we term it *Equi-SINR*.

3.2.2 Equi-SINR power allocation: Example

To get some better intuition about the behavior of COPA, we can take a look what happens on the subcarrier level for a single stream. Figure 7 shows the BER per subcarrier for a single stream for a pair of four-antenna AP, two-antenna client networks, with concurrent transmissions *i.e.,* when we use a nulling precoding matrix, and the same bitrate both for COPA and for the no-power allocation case. The throughput value for no power allocation ("NoPA") is the one achieved when using its optimal bitrate (13.5 Mbps). Although NoPA has better BER than COPA in several subcarriers, it exhibits a great variation. On the contrary, COPA has a lower overall BER variation and drops particularly bad subcarriers. As a result, although COPA drops 8 of the subcarriers, it achieves higher throughput because it manages to use a higher bitrate (39 *vs.* 13.5 Mbps) in the remaining ones, leading to a significant throughput increase.

3.3 Predicting the Best Strategy

The overall architecture for each AP's implementation of COPA is shown in Figure 8. After receiving the CSI, the leader AP calculates four precoding matrices: two "transmit beamforming" ones (one for itself and one for the follower AP) that maximize power at the intended receiver, and are calculated using the *Singular Value Decomposition* (SVD) of

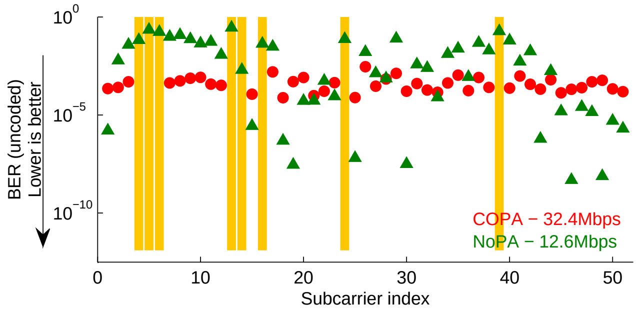

Figure 7: BER per subcarrier without coding when using the same nulling precoding matrix and bitrate (39 Mbps) with COPA and without ("NoPA"). Vertical bars denote subcarriers that COPA drops.

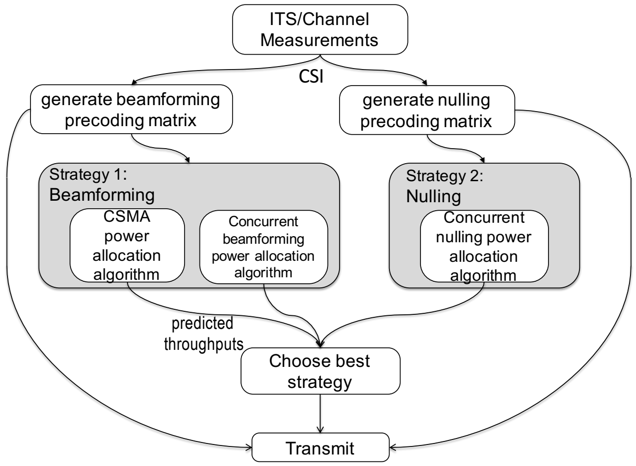

Figure 8: Overview of COPA

the appropriate channel, and two "nulling" ones (again one for itself and one for the follower) that use a combination of nullspace projection and the SVD to null interference at the unintended receiver while maximizing power at each AP's own client.

Returning to Figure 8, we see that COPA applies the Equi-SINR power allocation and subcarrier selection algorithm and calculates the possible throughput for multiple medium access strategies. First of these is *COPA-SEQ*, in which transmitters use the transmit beamforming precoding matrices and transmit sequentially. In our results, COPA-SEQ always beats stock 802.11n without power allocation, which is expected since the latter serves as its starting point. We can also use the same transmit beamforming precoding matrices for a concurrent strategy. This non-nulled concurrent solution only performs Equi-SINR power allocation and subcarrier selection to mitigate interference. For single antenna cases, where nulling is not possible, this is the only concurrent strategy we consider. But even when nulling is possible, if the cross-interference is very weak, nulling may unnecessarily waste transmit power without yielding any useful reduction in interference. Finally, another concurrent strategy is to use Equi-SINR with the nulling (or possibly alignment) precoding matrices. This subsumes traditional nulling, which serves as the starting point for the first iteration of Equi-SINR.

We note here again that additional power allocation schemes are possible. In our experiments we also apply mercury/water filling (see §2.1) instead of equalizing SINR in Step 4 of Algorithm 1.

Once all these power allocation matrices have been generated, we calculate the effective BER and hence the aggregate throughput the two APs would achieve if they chose each scheme. Again referring to Figure 8, COPA then transmits with the throughput-maximizing strategy.

Ideally, we would like to be able to tell in advance which strategy would win without trying different solutions and comparing them. In practice, this is not so easy: if cross-interference is relatively weak, nulling with power allocation always beats CSMA. If interference is very weak indeed, concurrent sending without nulling can even beat nulling as it sends more power to the intended recipient.

3.4 Overconstrained Nulling

Nulling requires that a transmitter have enough antennas to phase-align its transmissions so that they cancel each other at an unintended receiver. For example, if both APs have four antennas and both clients have two antennas, each AP can generate two MIMO streams to its own client and still null its transmission at both the other client's antennas.

Sometimes though, the problem is overconstrained. If, for example, two APs have three antennas each, and their clients each have two antennas, then the APs are faced with a stark choice: either send two MIMO streams each but don't null, resulting in significant interference, or send only one MIMO stream each and null that stream at the other AP's client. The problem arises from the fact that if a receiver receives more concurrent streams than the number of its antennas, it cannot use MMSE to disentangle them, which can make the intended streams undecodable.

We have investigated whether in this scenario it is possible to send two streams while *partially* nulling by optimizing the precoding matrix for aggregate throughput. The short answer is that in many cases it is, but the optimization is extremely computationally expensive, requiring tens of seconds of compute time for calculating the precoding matrix for a single subcarrier. Instead, because our APs are cooperating, we can adopt a more effective and much cheaper solution: simply shut down a receive antenna.

When the second AP responds to the first AP's ITS, it already knows whether the problem is overconstrained. In its responding ITS, it indicates that it wishes to participate in a concurrent transmission (if that is what COPA decides is the best strategy), but with reduced rank so that the problem is no longer overconstrained. It chooses whichever of its client's antennas has the best expected SINR, and indicates that the other antenna should be shut down for this transmission. By doing so, the receiver does not have to deal with the potentially high levels of interference that will inevitably end up in the shut antenna and with which it cannot deal with using MMSE since it doesn't have enough inputs.

In the case of two three-antenna APs sending to two two-antenna clients, this would result in AP1 sending two MIMO streams to client 1 while nulling to client 2's remaining antenna, whereas AP2 sends one MIMO stream to client 2 and nulls to both of client 1's antennas. In principle this may give up to a 50% throughput improvement over CSMA.

Although this solution is asymmetric - one client gets more throughput than the other - randomness in the DCF results in each AP sending the first ITS about half the time, so on average the asymmetry cancels out. This simple solution works well, as we will show in the evaluation.

3.5 Incentive Compatibility

Our goal so far has been to maximize aggregate throughput. This clears any transmission backlog fastest, but it may result in one receiver getting lower throughput than it would if its AP had not cooperated. In general, we would like a solution to be *incentive-compatible*, in the sense that no client ever loses out if its AP cooperates using COPA. In this way, APs always have an incentive to cooperate.

A simple tweak to COPA makes it incentive-compatible: simply revert to sequential transmission with power allocation and subcarrier selection if concurrent transmission would reduce either client's throughput. This is done by including this as an additional criterion when deciding on the "best strategy" in Figure 8, and we evaluate this next in §4.

4. EVALUATION

Summarizing our findings thus far, COPA is motivated by experiments that show nulling does not perform terribly well in our environment. But as Figure 3 shows, nulling does significantly reduce interference, and as a result increases SINR. However, it introduces an increase in the SINR variability across subcarriers (Figure 4), resulting in bit errors on weak subcarriers, thus seriously limiting throughput.

In this section we will experimentally demonstrate that in many scenarios, leveraging cooperative subcarrier selection and power allocation rescues performance and allows APs to transmit concurrently. Despite this, we will see that our nulling/power-allocation solution does not always work well, and so a system must consider multiple possible algorithms including sequential transmissions.

4.1 Experimental Methodology

We ran experiments throughout our lab, which includes open-plan floor space as well as offices and corridors.

We use the WARP version 2 platform for both senders and receivers running in the 2.4 GHz band using 20MHz channels and 15 dBm of maximum transmit power. All four nodes are connected to a PC that runs a modified version of the WARPLab framework that is used to calculate the precoding matrices and power allocations for our experiments. Our WARPs are mounted on trolleys so we can move them around the building to test many different topologies.

We choose topologies that include both short and long links, and weighted the scenarios so that usually the signal of interest was more powerful than the interfering signal. The rationale for this is that hosts are normally (but not always) closer to their own AP than to an interfering AP. In a few topologies we deliberately positioned the receiver so its direct line of sight was blocked by a metal filing cabinet. Figure 9 shows each receiver in each topology as a point. SNR at a client from its own AP is plotted against the SNR of the signal from the interfering AP (effectively the INR). Clients below the $x = y$ line have stronger intended signal than interfering signal. It can be seen from this plot that our topologies include a fairly wide range of signal and interference strengths, though there are few really bad channels because of the nature of the building. We will separately examine channels with weaker interference in §4.4.

In each topology we ran the following scenarios:

- Two single-antenna APs transmitting to two single-antenna clients.
- Two four-antenna APs transmitting to two two-antenna clients. We refer to this as the "constrained case", where four MIMO streams and full nulling should be possible.
- Two three-antenna APs transmitting to two two-antenna clients. We refer to this as the "overconstrained case", where there are not enough transmit antennas to both send four MIMO streams and fully null.

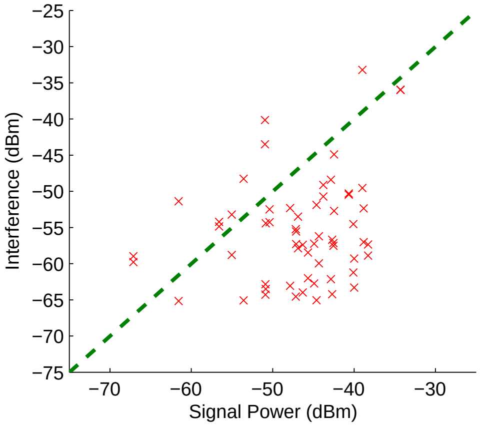

Figure 9: Signal power from each client's own AP plotted against interfering signal power from the other AP; each point is one receiver.

In our experiments involving concurrent transmission, each sending host transmits on its own to the two receivers, which record the samples they observe. Normally, each transmission is scaled by the receiver's AGC so that it fills the dynamic range of the ADC in the transceiver. Before combining the two signals, we revert this scaling by dividing each signal's samples by the gain applied by the AGC, in floating point to avoid losing precision. We do so both because of WARP's limited synchronization capabilities and because it allows us to take more accurate interference measurements, which on many occasions would otherwise have been lost below the noise introduced by the variable-gain amplifier.[5]

Rahul *et al.* [21] have demonstrated a method that can synchronize two hosts with an error of less than 20 ns, which is shorter than the sampling interval of a typical 802.11 transceiver. At the start of a reception, receivers use AGC to set the correct amplifier gain and Schmidl-Cox for synchronization. With the two transmitting hosts synchronized within tens of nanoseconds, both of these methods work correctly. The potential phase offsets between the interfering and useful transmissions at the receiver are irrelevant, since it needn't decode the interfering signal. One potential difficulty of combining the two transmissions is the addition of AWGN twice, but as the noise level of the interfering transmission is typically significantly lower than that of the useful one, doing so only underestimates the performance of our system.

To send multiple streams, hosts use the singular value decomposition of the channel and to null we project onto the appropriate nullspace. On the receiving side, hosts use a Minimum Mean Square Error filter to maximize the received power without amplifying noise.

We use the measured SINRs to calculate the uncoded BER [8] for each 802.11n modulation, from which we in turn calculate the coded BER for 802.11n's different coding rates [26]. From the frame error rates, we predict throughput achieved using the standard 4 ms transmit opportunity duration.

Our results include the appropriate MAC overhead for each scheme—ITS for concurrent schemes, CTS-to-self for CSMA, preamble, and ACK. During the ITS exchange, we include the transmission of CSI and precoding matrices once every 30 ms. Our experiments assume greedy unidirectional flows. We also assume that each AP already has knowledge of the channel between itself and the clients which would normally be measured implicitly from previous transmissions of data or control frames. If (fresh) CSI for a client is not available, its AP can either use the *No Data Packet* (NDP) mechanism or probe for a frame with staggered preambles to acquire it as is done in standard 802.11n/ac. Doing so would require the exchange of two short packets (20–30 μs) adding about 0.2% of overhead in an environment with a 30 ms coherence time; COPA and vanilla 802.11 would incur the same such overhead.

Selectively using subcarriers could problematically increase the *Peak to Average Power Ratio* (PAPR). In our experiments hosts only drop a few subcarriers; there are enough remaining and they have enough entropy from data scrambling that we do not observe any such problem.

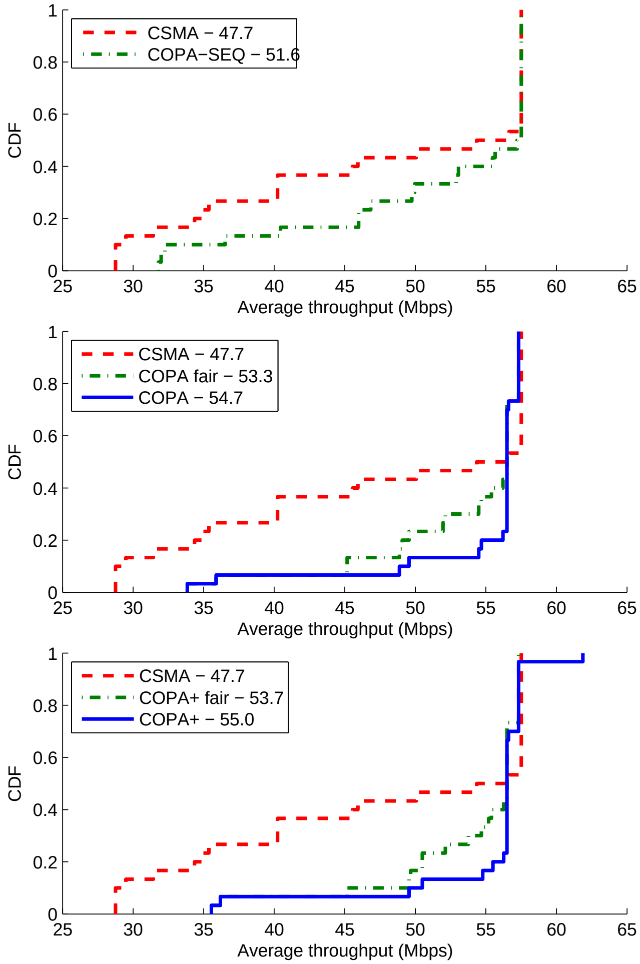

Figure 10: Throughput CDF (across topologies) for two single-antenna AP, single-antenna client pairs.

Limitations of our methodology. The WARP platform incurs a latency cost when downloading samples received and uploading samples to send. In our experiments, there is thus a 2–3 second delay between measuring the CSI and using the resulting precoding matrix and power allocation for concurrent transmissions. However, measurements of our indoor testbed environment (omitted in the interest of brevity) show that the channel coherence time in our experiments exceeds the WARP's latency penalty, and so the results we present are representative of a system operating at line speed. As noted above, our experimental evaluation of the throughput achieved by COPA, however, includes the overhead of COPA's ITS exchanges (see Table 1) in an environment with a much shorter 30 ms coherence time—we do so to account for COPA's costs in a more dynamic environment.

Finally, we cannot measure the performance of mercury/waterfilling in live experiments, as computing its concurrent power allocation solution requires tens of seconds (typically 30–50 s) in four-stream scenarios. For these results, we instead use an emulated channel.

4.2 Single-Antenna Scenario

Although we don't anticipate concurrent sending to work well very often with single-antenna APs, we investigate in order to understand the effect and limitations of power al-

---

[5] Although this process results in reduced quantization noise for the weaker transmission than what it would normally experience, it doesn't affect our results. If the interfering transmission is the weaker one, we do not decode it. If the weaker transmission is the useful one, it will in any case be undecodable because of its low SINR.

location and subcarrier selection. In Figure 10, we present three CDF graphs showing the effects of the different strategies in 30 topologies—the top graph shows non-concurrent variants, the middle shows practical COPA variants (including concurrent and sequential transmission), and the lower graph shows the best but impractical COPA variants.

In just under half the topologies regular CSMA (*i.e.,*, no concurrent senders) can achieve throughput of 57.5 Mbps—the maximum achievable rate when transmitting at 65 Mbps with a 4 ms transmit opportunity. The mean throughput though is 47.7 Mbps because the remaining receivers fare fairly poorly. The COPA-SEQ curve shows the effect of Equi-SINR power allocation and subcarrier selection, without concurrent senders. COPA-SEQ achieves a mean throughput of 51.6 Mbps. We have investigated whether this improvement comes from subcarrier selection or from power allocation: either one, by itself gives about 60-70% of the improvement, but both are needed together for the full benefits to be seen. CSMA plus mercury and water filling, for example, gives 50.1 Mbps in this scenario.

In the middle graph we can see the improvement over CSMA when we allow COPA to do concurrent transmission. The "COPA Fair" curve shows how COPA performs when we restrict it to be incentive compatible, whereas "COPA" just aims for maximum total throughput.

"COPA Fair" achieves 3% more than COPA-SEQ; the improvement is due to concurrent transmission being selected in some topologies, even though nulling is not possible with a single antenna. In about 15% of the topologies there is a slight drop in throughput when using "COPA Fair" due to the increased MAC overhead.

The "COPA" curve shows how COPA performs when it does not need to be incentive-compatible. Aggregate throughput now rises to 55 Mbps, and concurrent transmission is possible in more cases, even though without nulling it does not greatly improve throughput. Here COPA has selected a form of OFDMA, with some subcarriers being used by only one AP at a time. In these few cases each subcarrier is used by the AP that can best make use of it, despite unfairness. The difference between the "COPA" and "COPA Fair" curves is the price of fairness. In this single-antenna case, COPA often gives *all* the wireless capacity to one receiver.

The "COPA+" curves in the lower graph in Figure 10 shows what happens when we include iterated mercury/waterfilling (including subcarrier selection) among the strategies COPA can select. These curves are trace-driven emulation based on real CSI measurements because of the high processing time of this algorithm. Although COPA+ is impractical, the curves illustrate the additional gains that a more optimal power allocation scheme might achieve. At the top right of the graph we can see true concurrent transmission occurring in two topologies, with the same subcarrier successfully being used concurrently by both APs.

4.3 Constrained Nulling Scenario

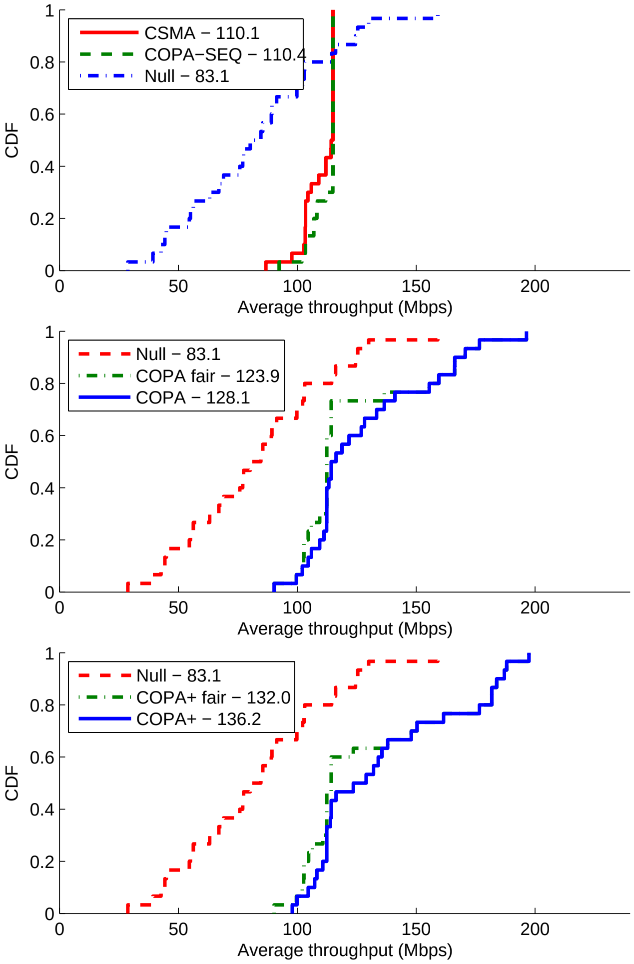

Figure 11: Throughput CDF (across topologies) for two four-antenna AP, two-antenna client pairs.

In Figure 11 we show the performance achieved by two four-antenna APs transmitting to two two-antenna clients. In this scenario there are enough degrees of freedom that it ought to be possible to perform concurrent transmission of two MIMO streams from each AP to its client while nulling at the other client. Again, the top graph shows non-concurrent variants, the middle shows COPA with concurrent transmission, and the lower graph shows the impractical COPA+ results. We also show regular nulling without power allocation or subcarrier selection as a baseline in all graphs.

When we first obtained these results we were surprised at how poorly nulling performs with OFDM. Nulling only outperforms CSMA in 17% of our topologies, and even then, not by much. We examined the cause in Figure 4. Nulling works well for some subcarriers but not so well for others. This, superimposed on top of the variability in SNR when a transmitter tries to null an unintended receiver, leads to high SINR variability between subcarriers. Because 802.11 hosts use a single decoder, weaker subcarriers dominate BER, dragging down the achievable bitrate.

In this scenario, CSMA achieves two full MIMO streams in 60% of the topologies with only slightly reduced throughput for most of the remainder as the power budget with four

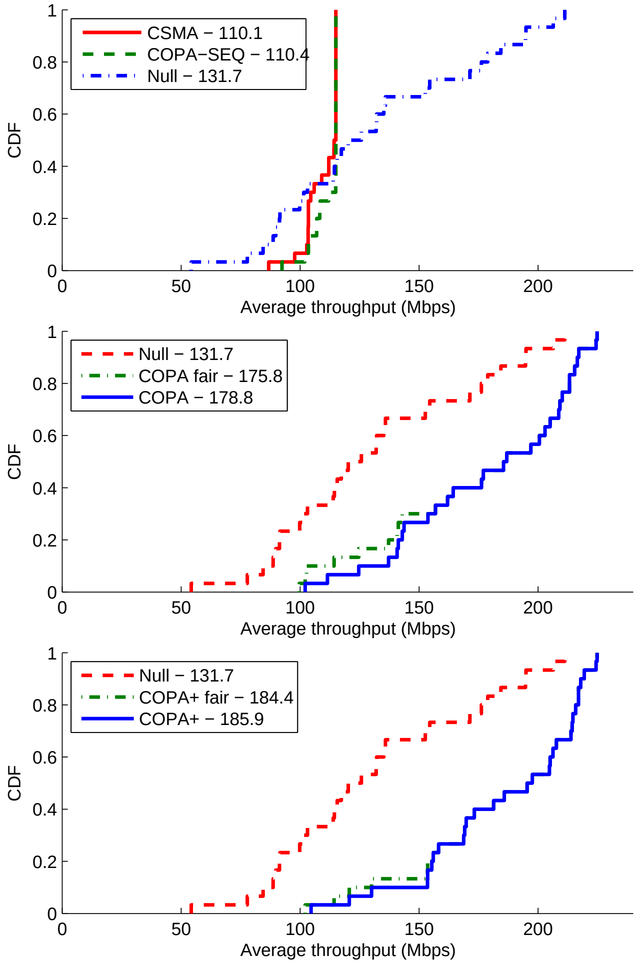

Figure 12: Throughput CDF (across topologies) for two four-antenna AP, two-antenna client pairs, when interference is *10 dB weaker* than empirically measured in our testbed.

antennas is 4x higher than in the previous scenario.

The “COPA” curve on the middle graph shows how much more effective nulling is when combined with power allocation and subcarrier selection. There is a mean of 54% improvement over vanilla nulling. If we use the incentive-compatible “COPA fair” variant that figure drops to 48%; the 6% difference between the two is the cost of fairness. In approximately 30% of cases “COPA” selects COPA-SEQ because no better concurrent solution was found. In all the remaining cases a concurrent nulled solution is chosen. The mean throughput improvements don’t tell the whole story though — nulling has higher variance than CSMA. Even though COPA improves things significantly, the variance is still high. Sometimes COPA gives negligible improvement over CSMA, but when it uses concurrent transmissions the gains are substantial. Finally, the COPA+ curves in the lower graph indicate that using iterated mercury/waterfilling might yield a further increase of about 10%.

4.4 Nulling with Weaker Interference

Other researchers have reported nulling works better in their environments than we find it to do in ours. We speculate that their building construction or choice of interferer location must result in lower cross-interference than we observe. To test this hypothesis we took the traces from all our 4x2 topologies, reduced the interference strength by 10dB, left the signal of interest unchanged, and ran emulated experiments. This results in emulated topologies where APs can still hear each other well enough to exchange ITS packets, but where before nulling the interference is on average about 20dB below the intended signal.[6]

Results are shown in Figure 12: with weaker interference, vanilla nulling works relatively well, beating CSMA in 65% of topologies. However, COPA greatly increases throughput. With weak interference, COPA almost never needs to fall back to COPA-SEQ. There is little difference between “COPA” and “COPA Fair” because both clients normally win from running COPA. Even when they don’t both win, the weaker client doesn’t suffer greatly. COPA beats CSMA by 62% and beats vanilla nulling by 36%. The win is biggest for the receivers that do worst with nulling, with many receivers getting more than 50% better throughput with COPA than with vanilla nulling. COPA+ does even better, beating vanilla nulling by 41%, but is unlikely to be practical.

4.5 Overconstrained Scenario

Between the single-antenna case and the full 4x2 case lie overconstrained scenarios, where some measure of nulling is possible, but there are not enough degrees of freedom to null completely. To explore this region, we examined the case where two three-antenna APs each have a two-antenna client. In this scenario we can normally only use CSMA to send two streams at a time since the transmitters do not have enough antennas to null towards their unintended receiver.

We have already seen that even when there are enough degrees of freedom, nulling has limited effectiveness when using OFDM. To improve that starting condition for COPA’s concurrent strategies, we have one of the two APs tell its receiver to shut down a receive antenna, so the problem is no longer overconstrained. Both APs then have enough degrees of freedom to proceed with their transmission while nulling towards their unintended receiver.

The “Null+SDA” curves in Figure 13 shows the effect of shutting down an antenna (SDA) and then performing otherwise-vanilla nulling. Thus provides some benefit to clients with good channels but, by itself, doesn’t come close to CSMA throughput. COPA does significantly better. “COPA Fair” beats CSMA by 13% and “COPA” beats CSMA by 17%. Around 40% of topologies can choose concurrent strategies, and those that do gain by between 20% and 40% over CSMA. Unsurprisingly, COPA+ does even better, with around 60% of topologies benefiting from using concurrent strategies.

4.6 Multiple Decoders

As we saw earlier, the variability of SINR across subcarriers results in diminished throughput. By dropping or al-

---

[6] To see this, take Figure 9 and move each point down by 10 dBm.

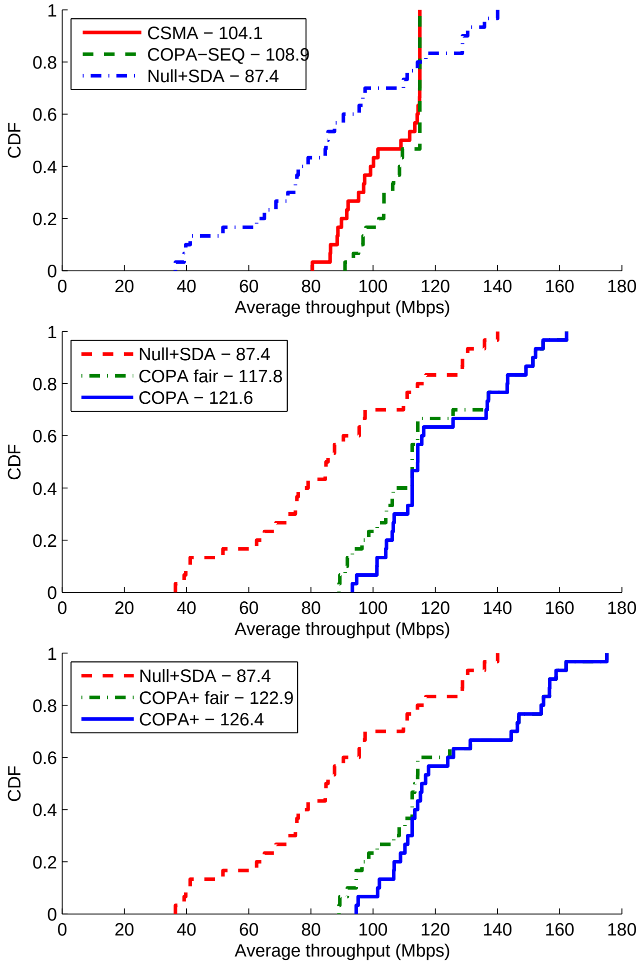

Figure 13: Throughput CDF (across topologies) two three-antenna AP, two-antenna client pairs.

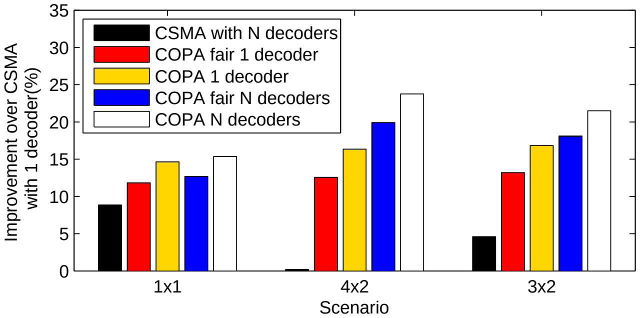

Figure 14: Potential percentage improvements from using per-stream error control coding with CSMA and COPA.

locating more power to subcarriers which experience more adverse channel conditions COPA helps the situation to a large extent. Although current hardware does not support it, a further improvement can come by selecting a bitrate for each subcarrier independently. This would mean that both AP and client would need to use multiple encoders and decoders for each stream (one for each coding rate supported). However, it would allow us to adapt the the SINR of each subcarrier fully, allowing for great variation across subcarriers.

In Figure 14 we examine the potential benefit of using one decoder per coding rate in our topologies.[7] The figures give the improvement of each scheme over what CSMA would achieve with one decoder. In the single-antenna case, using multiple decoders improves CSMA performance, but fails to greatly improve COPA performance as nulling isn't possible. In the four- and three-antenna AP cases, on the other hand, CSMA doesn't greatly benefit as it is already running at full speed. Using one decoder per channel increases throughput over baseline "COPA" and "COPA Fair" by about a further 10% with four antenna APs and by about 5% with three antenna APs. These gains are modest: even with a single decoder COPA has already realized most of the potential gains.

5. RELATED WORK

Channelization-based systems. WiFi-NC [4] splits a single wideband OFDM channel into multiple narrower subchannels that can operate independently from each other, generalizing earlier work on selecting fixed-size channel widths [3]. Neither, however, addresses the problem of interference from adjacent senders.

OFDMA-based systems. FARA [20] uses OFDMA on an AP's downlink but unlike COPA, does not take inter-AP interference into account. FICA [24] builds on FARA, adding uplink OFDMA and medium access refinements, but does not allow neighboring APs to reuse subcarriers to different hosts as COPA does. Yu *et al.* [29] propose iterative waterfilling for power assignment in Gaussian channels, while COPA's practical system design handles QAM signals.

Multi-antenna interference mitigation. BigStation [28], Geosphere [16], and SAM [25] use nulling or the Sphere decoder to separate users' signals on the uplink (SAM and Geosphere), or both uplink and downlink (BigStation). TIMO [5] uses the same technique to null cross-technology interference. But these systems support a number of users limited by the number of antennas at the AP, making no explicit provision for interference between adjacent APs.

OpenRF [11] is closest to COPA in design, but requires a centralized controller and is limited by the number of spatial degrees of freedom in the system. COPA introduces per-subcarrier power control in conjunction with interference nulling to overcome this limitation. DIRC [13] and Speed [14] use fixed-pattern directional transmissions and receptions, respectively, to improve spatial reuse, but suffer in the presence of indoor multipath. Distributed MIMO and distributed antenna systems such as MegaMIMO [22] and MIDAS [27] achieve high capacity but require tight cooperation between participating APs, making no provision for interference from the "outside."

Interference alignment (IA) IAC [6] describes a LAN IA design, while Suh and Tse [23] describe a cellular IA design. 802.11n+ [12] opportunistically nulls and aligns signals, adding a decentralized MAC design. Independently,

---

[7] 802.11 uses four coding rates and current 802.11ac transceivers usually have 2 decoders.

Gomadam *et al.* [7] propose iterative IA algorithms that require only local knowledge at each node to asymptotically approach the Shannon capacity of interference networks at high SNR, but do not present experimental results from a system implementation. While IA techniques hold great promise, many require large symbol extensions to acheive their asymptotic bounds, thus reducing practicality, while others limit the number of concurrent aligned streams [7]. The power control that COPA uses is a practical way of overcoming these limitations and is complementary.

Spatial reuse in cellular systems. GSM, LTE, and WiMAX leverage *hard reuse*, assigning time/frequency blocks based on users' distance to the base station [9]. *Soft fractional frequency reuse* tunes downlink base station transmit power resulting in theoretical capacity gains [18]. Compared to these approaches, COPA does not require a centralized controller to coordinate the reuse, as well as using more selective and precise transmit beamforming to null interference between access points, rather than transmit sectors, and performing joint power allocation and spatial nulling.

Subcarrier Switch-Off. Nitsche *et al.* [17] show that for a single transmission, in the absence of any interference, subcarrier selection can yield significant throughput gains. Punal *et al.* [19] show in simulation that in the same scenario, switching off subcarriers that face adverse channel conditions can improve the robustness of subcarrier power allocation. COPA extends these approaches to MIMO transmission by multiple, mutually interfering senders, and converges to a throughput-maximizing power allocation in the more complex case where decisions about one stream affect all others.

6. CONCLUSION

We have presented COPA, an approach for mitigating interference between loosely cooperating Wi-Fi APs and clients that leverages per-subcarrier power allocation, interference nulling, and multi-stream transmission to improve throughput. Our experiments show that interference nulling increases the variability of SINR across subcarriers at receivers, but COPA's cooperative power allocation mitigates this, materially boosting throughput.

7. ACKNOWLEDGEMENTS

Nikolaidis and Karp received support from EU FP7 Grant 287581 (OpenLab). The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement no. 279976.

8. REFERENCES

[1] 802.11 specification. http://standards.ieee.org/about/get/802/802.11.html.

[2] Maxim 2829 datasheet. http://datasheets.maximintegrated.com/en/ds/MAX2828-MAX2829.pdf.

[3] R. Chandra, R. Mahajan, T. Moscibroda, *et al.* A case for adapting channel width in wireless networks. *SIGCOMM* (2008).

[4] K. Chintalapudi, B. Radunovic, V. Balan, *et al.* WiFi-NC: Wi-Fi over narrow channels. *NSDI* (2012).

[5] S. Gollakota, F. Adib, D. Katabi, and S. Seshan. Clearing the RF smog: Making 802.11 robust to cross-technology interference. *SIGCOMM* (2011).

[6] S. Gollakota, S. Perli, and D. Katabi. Interference alignment and cancellation. *SIGCOMM* (2009).

[7] K. Gomadam, V. Cadambe, and S. Jafar. Approaching the capacity of wireless networks through distributed interference alignment. *GLOBECOM* (2008).

[8] D. Halperin, W. Hu, A. Sheth, and D. Wetherall. Predictable 802.11 packet delivery from wireless channel measurements. *SIGCOMM* (2010).

[9] N. Himayat, S. Talwar, A. Rao, and R. Soni. Interference management for 4G cellular standards. *IEEE Communications Magazine*, 2010.

[10] H. Jung, K. Cho, Y. Choi, *et al.* React: Rate adaptation using coherence time in 802.11 wlans. *Computer Communications*, **34**(11): pp. 1316–1327, 2011.

[11] S. Kumar, D. Cifuentes, S. Gollakota, and D. Katabi. Bringing cross-layer MIMO to today's wireless LANs. *SIGCOMM* (2013).

[12] K. Lin, S. Gollakota, and D. Katabi. Random access heterogeneous MIMO networks. *SIGCOMM* (2011).

[13] X. Liu, A. Sheth, M. Kaminsky, *et al.* DIRC: Increasing indoor wireless capacity using directional antennas. *SIGCOMM* (2009).

[14] —. Pushing the envelope of indoor wireless spatial reuse using directional access points and clients. *MobiCom* (2010).

[15] A. Lozano, A. Tulino, and S. Verdú. Optimum power allocation for parallel gaussian channels with arbitrary input distributions. *IEEE Trans. on Information Theory*, **52**(7): pp. 3033–3051, 2006.

[16] K. Nikitopoulos, J. Zhou, B. Congdon, and K. Jamieson. Geosphere: Consistently turning MIMO capacity into throughput. *SIGCOMM* (2014).

[17] T. Nitsche and J. Widmer. Sub-carrier switch off in ofdm-based wireless local area networks. *Sensor, Mesh and Ad Hoc Communications and Networks (SECON), 2013 10th Annual IEEE Communications Society Conference on*, pp. 442–450 (IEEE, 2013).

[18] T. Novlan, R. Ganti, A. Ghosh, and J. Andrews. Analytical evaluation of fractional frequency reuse for OFDMA cellular networks. *IEEE Trans. on Wireless Comms.*, **10**(12): pp. 4294–4305, 2011.

[19] O. Punal, H. Escudero, and J. Gross. Power loading: Candidate for future wlans? *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2012 IEEE International Symposium on a*, pp. 1–4 (IEEE, 2012).

[20] H. Rahul, F. Edalat, D. Katabi, and C. Sodini. Frequency-aware rate adaptation and MAC protocols. *MobiCom* (2009).

[21] H. Rahul, H. Hassanieh, and D. Katabi. Sourcesync: a distributed wireless architecture for exploiting sender diversity. *ACM SIGCOMM Computer Communication Review*, **41**(4): pp. 171–182, 2011.

[22] H. Rahul, S. Kumar, and D. Katabi. JMB: Scaling wireless capacity with user demand. *SIGCOMM* (2012).

[23] C. Suh and D. Tse. Interference alignment for cellular networks. *Allerton* (2008).

[24] K. Tan, J. Fang, Y. Zhang, *et al.* Fine-grained channel access in wireless LAN. *SIGCOMM* (2010).

[25] K. Tan, H. Liu, *et al.* SAM: Enabling practical spatial multiple access in wireless LAN. *MobiCom* (2009).

[26] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication* (Cambridge University Press, 2005).

[27] J. Xiong, K. Sundaresan, K. Jamieson, *et al.* MIDAS: Empowering 802.11ac networks with multiple-input distributed antenna systems. *CoNEXT* (2014).

[28] Q. Yang, X. Li, H. Yao, *et al.* BigStation: Enabling scalable real-time signal processing in large MU-MIMO systems. *SIGCOMM* (2013).

[29] W. Yu, W. Rhee, S. Boyd, and J. Cioffi. Iterative water-filling for Gaussian vector multiple-access channels. *IEEE Trans. on Information Theory*, **50**(1): pp. 145–152, 2004.