

Combating Inter-cell Interference in 802.11ac-based Multi-user MIMO Networks

Hang Yu, Oscar Bejarano, and Lin Zhong

Department of Electrical and Computer Engineering, Rice University, Houston, TX
{Hang.Yu, obejarano, lzhong}@rice.edu

Abstract

In an 802.11ac-based MU-MIMO network comprised of multiple cells¹, inter-cell interference allows only a single AP to serve its clients at the same time, significantly limiting the network capacity. In this work, we overcome this limitation by letting the APs and clients in interfering cells coordinately cancel the inter-cell interference using their antennas for beamforming. To achieve such coordinated interference cancellation in a practical way, we propose a novel two-step optimization. First, without requiring any channel knowledge, each AP and client optimizes the use of its antennas for either data communication or inter-cell interference cancellation, in order to maximize the total number of deliverable streams in the MU-MIMO network. Second, with only partial channel knowledge, each AP and client optimizes their beamforming weights after the optimal antenna usage has been identified in the first step. Our solution, CoaCa, integrates this two-step optimization into 802.11ac with small modifications and negligible overhead, allowing each AP and client to locally perform the two-step optimization. Our experimental evaluation indicates that for a MU-MIMO network with two cells, by cancelling the inter-cell interference CoaCa can convert the majority of the expected number of streams increase (50%-67%) into network capacity improvement (41%-52%).

Categories and Subject Descriptors

C.2.1 [COMPUTER-COMMUNICATION NETWORKS]: Network Architecture and Design - Wireless Communication

Keywords

Inter-cell interference; multi-user MIMO; 802.11ac; CoaCa

1. INTRODUCTION

For 802.11 networks comprised of multiple cells, inter-cell interference has become a key factor that limits the network

capacity because it prevents the APs in neighboring cells from serving their clients concurrently. To cancel inter-cell interference with *beamforming*, the interfering AP and client are required to (i) be aware of the channel between them, and (ii) coordinate to determine their duty of cancellation. Yet, neither of the two requirements is directly supported by existing 802.11 protocols. There have been efforts in supporting them in 802.11n that features single-user MIMO (SU-MIMO). 802.11n+ proposed in [10] seeks to enable concurrent links across multiple cells. n+ cancels the inter-cell interference by (i) letting nodes in one cell overhear the transmissions from nodes in the other cell to acquire necessary channel knowledge, and (ii) coordinating nodes in the two cells by opportunistically starting the concurrent link in the overhearing cell afterwards. Notably, the interference cancellation coordination in n+ is one-way: nodes in the overhearing cell must use their spare antennas to cancel the inter-cell interference. The number of spare antennas should be no smaller than that of the ongoing streams in the overheard cell. Since an 802.11n link with SU-MIMO usually includes one or two streams, only one or two spare antennas are needed for each node in the overhearing cell.

For the newer 802.11ac supporting multi-user MIMO (MU-MIMO), an AP can transmit to multiple, say K (e.g., $K=4$) clients simultaneously. When a cell is congested with more clients than antennas on the AP, the AP delivers a single stream to each served client so that the number of streams is equal to the number of clients, denoted as the *multiplexing gain* of MU-MIMO. To increase the multiplexing gain by starting concurrent streams in the overhearing cell, n+ would require each AP and client in it to have at least K spare antennas to cancel the inter-cell interference, which is usually infeasible in practice. The fundamental reason why n+ does not effectively extend to 802.11ac is its one-way coordination for inter-cell interference cancellation, where nodes in the overhearing cell have to contribute all the required antennas to solely carry the burden of cancellation. Our key insight in this work is that when two congested MU-MIMO cells “jointly coordinate,” the number of required antennas on their nodes can be reduced and more streams can be delivered concurrently.

Let us consider the example in Figure 1. With three antennas on each AP and three clients in each cell, n+ only allows a single AP (AP1 or AP2) to serve its clients at a time, therefore activating only up to three streams in the network. Instead of solely relying on the nodes in one cell to cancel the inter-cell interference like in n+, we do it in the following alternative way. First, each AP uses two antennas to serve two clients (Client 1 and Client 3 in Cell 1, Client4 and Client6 in Cell 2), and the third antenna to cancel the inter-cell interference to one client in the other cell (Client1 in Cell 1, Client4 in Cell 2). Then, the other client in each cell (Client3 in Cell 1, Client6 in Cell 2) uses its three

¹We use *cell* to denote the domain of an access point (AP) and its associated clients.

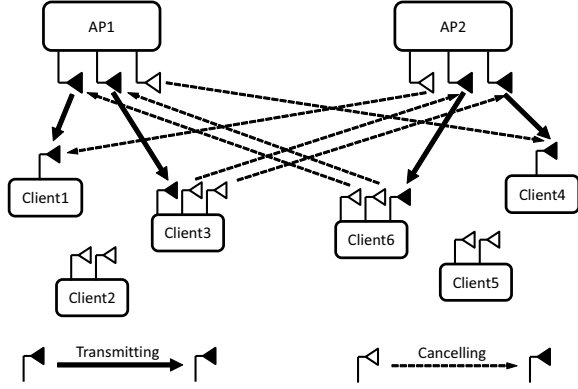


Figure 1: An example with two APs each serving up to three clients. Jointly coordinating the two cells to cancel the inter-cell interference delivers the maximum number of (four) streams.

antennas to cancel the inter-cell interference from the two streams sent by each AP. Consequently, we can activate four simultaneous streams in the network. Observe that both APs and clients have properly shared the responsibility of canceling the interference between them. Therefore, compared to $n+$ this is a more effective way to use the antennas on each AP and client. In fact, such joint cell coordination maximizes the number of streams or clients. The central question we seek to answer is: *how to practically achieve this cell coordination in a MU-MIMO network composed of multiple interfering cells?*

Determining the beamforming weights for each AP and client that cancel inter-cell interference with joint coordination is a non-trivial process. The theoretically optimal solution that maximizes network capacity can be empirically found by jointly optimizing the beamforming weights for all APs and clients. The optimal solution may not completely eliminate inter-cell interference since interference below the noise power is often no longer considered the capacity-limiting factor. However, identifying the optimal solution is neither practical nor compatible with 802.11ac because (i) it is computationally intractable due to the lack of an analytical solution, and (ii) it has to be done in a centralized way with full channel knowledge of the entire MU-MIMO network.

In this work, we present a novel solution that allows distributed cell coordination for inter-cell interference cancellation, and can be practically integrated into 802.11ac. The key idea in our solution is that the process of identifying the beamforming weights can be broken into two separate steps, namely *antenna usage optimization* and *beamforming weight optimization*. The first step determines how each AP and client antenna should be used: data communication or inter-cell interference cancellation. The optimal antenna usage of each AP and client collectively maximizes the number of streams in the MU-MIMO network. The second step determines the beamforming weights of each AP and client based on its optimized antenna usage. Given the use of antennas, it is possible to adopt practical beamforming techniques with a closed-form solution such as zero-forcing beamforming. The feasibility of such two-step optimization is based on an important heuristic we use to simplify the problem: we strive to completely eliminate inter-cell interference and maximize the multiplexing gain of MU-MIMO. Only with this heuristic the antenna usage can be reduced into a binary form including data communication and inter-cell interference cancellation, and optimized in a separate step prior to the optimization of beamforming weights.

The separation of the antenna usage optimization and beamforming weight optimization greatly simplifies the cell coordination effort, allowing us to practically integrate such two-step optimization into 802.11ac. Our proposed protocol, called CoaCa (Coordinated optimization of the AP and Client antennas), leverages the channel sounding process in 802.11ac to let each AP and client locally perform the two-step optimization in a distributed way. CoaCa includes two key designs. First, by interleaving the channel sounding process from all the APs, each node can easily acquire the necessary global information to optimize the use of their antennas. Such information including the number of antennas on each node only needs a few bits to be represented and can be explicitly shared by each AP. Second, by reporting and overhearing the beamformed channels in the interleaved channel sounding, each node can obtain just enough channel knowledge to optimize its beamforming weights. CoaCa incurs negligible overhead over 802.11ac and compatibly works with unmodified 802.11ac clients. While the current design of CoaCa only allows downlink MU-MIMO which is consistent to 802.11ac, the optimized beamforming weights for the APs and clients can be also used for uplink MU-MIMO leveraging channel reciprocity. However, realizing uplink MU-MIMO faces a new set of challenges such as misaligned symbol timing and clock frequency offset between clients, which are studied by other prior work, e.g., [17] and outside the scope of this work.

We realize a prototype of CoaCa on the WARP platform [14], and evaluate its performance in realistic indoor wireless environments. Our experimental results show that on average CoaCa is able to improve the capacity of a two-cell MU-MIMO network by 41% to 52%, with no more than four antennas on each AP and client. Even though the capacity gain is lower than the multiplexing gain increase (50% to 67%), CoaCa considerably outperforms existing solutions that often only allow a single cell to operate. While our evaluation does not include a large-scale MU-MIMO network, a *cell clustering* technique can be adopted to transform a large-scale network into several small clusters where each of them includes only two to three cells. In fact, the distributed nature of CoaCa makes it hard to gracefully scale with the number of cells. Unlike centralized solutions such as Network-MIMO [2,12] that can convert interference into signals, CoaCa must cancel the interference, requiring a much larger number of antennas on the interfering AP and clients. By transferring a large-scale MU-MIMO network into small clusters, we can apply CoaCa to each cluster independently, requiring only a small thus practical number of antennas on the APs and clients.

In summary, this work makes the following contributions:

- A novel solution that allows the APs and clients in multiple interfering MU-MIMO cells to coordinately cancel the inter-cell interference in two separate steps, including antenna usage optimization and beamforming weight optimization.
- An algorithm that efficiently identifies the optimal antenna usage for each AP and client in the MU-MIMO network to maximize the multiplexing gain.
- An analytical study of the channel knowledge requirement for each AP and client to locally optimize the beamforming weights based on its optimal antenna usage.
- A protocol that integrates antenna usage optimization and beamforming weight optimization into 802.11ac with small modifications and negligible overhead.

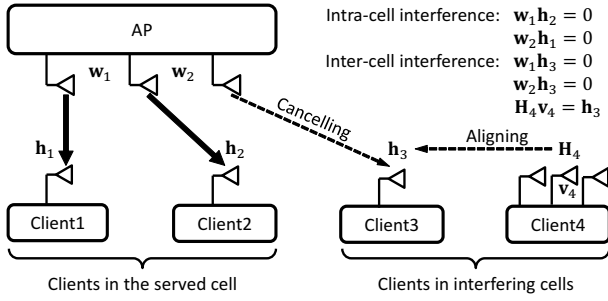


Figure 2: Intra-cell and inter-cell interference cancellation in a MU-MIMO network. The AP uses three antennas to cancel the intra-cell interference between Client1 and Client2, and the inter-cell interference to Client3. Given two additional antennas Client3 could cancel the inter-cell interference from the AP, saving one antenna for the latter. Client4 employs interference alignment to align its channel to that of Client3, so that the inter-cell interference from the AP is naturally eliminated when the AP cancels the interference to Client3.

2. BACKGROUND

In this section first we discuss relevant MU-MIMO techniques for interference cancellation. Then we present an overview of the supported MU-MIMO feature in the IEEE 802.11ac protocol.

2.1 Interference Cancellation in MU-MIMO

MU-MIMO improves network capacity by achieving a *multiplexing gain*, defined as the number of concurrent streams or simultaneously served clients. To appreciate the multiplexing gain in a MU-MIMO network composed of multiple interfering cells, both intra-cell and inter-cell interference must be sufficiently suppressed with the *beamforming* technique which we introduce below.

Intra-cell interference cancellation. To cancel the intra-cell interference, a MU-MIMO AP uses transmit beamforming to *precode* the data stream to each client. The downlink channels of the simultaneously served clients must be orthogonalized such that each client only receives its own stream without interference. The precoding strategy of the AP that achieves such orthogonality is known as *zero-forcing beamforming* (ZFBF). In zero-forcing beamforming, the transmit beamforming weight vectors, a.k.a. the precoding vectors of the AP, \mathbf{w}_j , are chosen to orthogonalize the channel vectors from the AP to the clients, \mathbf{h}_k , i.e., $\mathbf{w}_j \mathbf{h}_k = 0$ ($j \neq k$). To satisfy the orthogonality constraints, the dimension of \mathbf{w}_j , which is the number of antennas on the AP, must be no smaller than the number of served clients. This is the maximum multiplexing gain that can be achieved in a single MU-MIMO cell.

Inter-cell interference cancellation. The zero-forcing beamforming technique can be extended to cancel the inter-cell interference. That is, if there are any spare antennas on the AP after cancelling the intra-cell interference, they can be used to orthogonalize the channel vectors of the clients in other cells in the same way. Note that in the proper context without introducing ambiguity, we use *antenna* to refer to the *Degree of Freedom* (DoF) provided by a physical antenna on an AP or a client. In Figure 2, since there are three antennas, the AP uses two of them to cancel the intra-cell interference between Client1 and Client2, and the third (spare) one to cancel the inter-cell interference to Client3 in the interfering cell.

Inter-cell interference can be alternatively cancelled by a client if the client features multiple antennas for receive beamforming, a.k.a., post-combining [23]. Note that receive beamforming actually allows a client to separate and recover both the intended




Figure 3: The channel sounding process in 802.11ac. First the AP sends a NDP-A frame and a NDP frame for all specified clients to estimate their downlink channels, and then each client sequentially replies with a BF-R frame containing the estimated channels.

and interference streams. For ease of explanation we simply use the term *interference cancellation* to denote such capability of a receive beamforming client. To cancel the inter-cell interference, the client j chooses its receive beamforming weight vector, a.k.a. the post-combining vector, \mathbf{v}_j , such that $\mathbf{w}_k \mathbf{H}_j \mathbf{v}_j = 0$ for all served clients k by the AP, where \mathbf{H}_j is the channel matrix from the AP to client j . For example, in Figure 2 if Client3 had two spare antennas they could be used to cancel the two streams from the AP to Client1 and Client2, saving the third antenna on the AP. This would provide the AP with the flexibility to use its third antenna to potentially serve another client. Without inter-cell interference, multiple client antennas can be used to increase the SNR of the received stream via maximum ratio combining (MRC) [18].

Another technique for dealing with inter-cell interference is *interference alignment* [6, 8, 10, 11]. The key idea in interference alignment is to align the channel vectors of multiple clients, i.e., $\mathbf{H}_j \mathbf{v}_j = \mathbf{H}_k \mathbf{v}_k$, so that the interference between the AP and these clients traverses a single aligned channel and requires fewer AP antennas to be cancelled. Such alignment can be conceptually understood as if the client used its own antennas to cancel the interference saving the antennas on the AP. Note that interference alignment is more commonly assumed for uplink transmission, e.g., multiple clients transmit to an AP they interfere with on an aligned channel. In this work we leverage channel reciprocity to apply interference alignment to downlink MU-MIMO in which the AP transmits to multiple clients it interferes with on the same aligned channel. In Figure 2, without interference alignment, the AP would need two spare antennas to cancel the inter-cell interference to both Client3 and Client4. When Client4 aligns its channel to that of Client3 by setting $\mathbf{H}_4 \mathbf{v}_4 = \mathbf{h}_3$, the AP only needs one antenna to cancel the inter-cell interference.

2.2 MU-MIMO in 802.11ac

Now we briefly introduce the MU-MIMO feature supported by 802.11ac [4], the latest amendment to the 802.11 protocol family. 802.11ac allows the AP to use MU-MIMO techniques to simultaneously transmit downlink streams to up to four of its served clients. For APs in the interference range of each other, 802.11ac does not allow them to transmit at the same time; instead, the APs contend to access the medium using CSMA/CA.

Channel knowledge is necessary for the AP to calculate the transmit beamforming weights that cancel the intra-cell interference. To acquire channel knowledge, 802.11ac mandates an explicit channel sounding process, which we show in Figure 3. To sound the channel, the AP first broadcasts a *Null Data Packet Announcement* (NDP-A) frame. The purpose of the NDP-A frame is to specify the set of clients the AP is about to serve, and notify them to prepare for estimating and reporting their downlink channels. After the NDP-A frame, the AP sends a *Null Data Packet* (NDP) frame that allows the clients to estimate their downlink channels leveraging the training symbols in the frame. Then, the

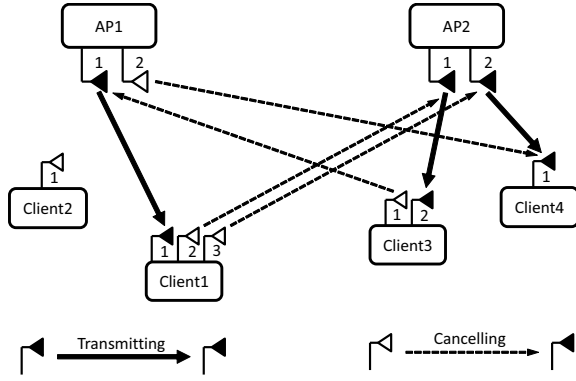


Figure 4: The optimal use of the AP and client antennas in our illustrative example. The interference from AP2 to Client1 is cancelled by Client1, while the interference from AP1 to Client3 and Client4 are cancelled by Client3 and AP1, respectively.

specified clients sequentially report their estimated channels to the AP with a *Beamforming Report* (BF-R) frame. The NDP-A frame designates a client that must immediately reply with the BF-R frame after the NDP frame, while other clients must wait for the *Beamforming Report Poll* (BF-P) frame from the AP to respond. The explicit channel sounding process in 802.11ac does not require channel reciprocity which is needed by implicit channel estimation [16]. The inter-frame interval in the channel sounding process is *SIFS* (16 μ s), which is shorter than *DIFS* (34 μ s) and thereby provides the APs and clients guaranteed medium access without being intervened by other 802.11 nodes.

3. OVERVIEW

To improve the capacity of an 802.11ac-based MU-MIMO network by allowing more streams in its cells, we propose a novel solution that practically achieves coordinated inter-cell interference cancellation with AP and client beamforming. The key idea in our solution is that the process of determining the beamforming weights for each AP and client can be broken into two steps, namely *antenna usage optimization* and *beamforming weight optimization*. We are motivated by the following important observation: when we strive to completely eliminate interference, the two optimizations can be executed sequentially. This is because the constraint of completely eliminating inter-cell interference reduces the antenna usage into a binary form. That is, one antenna can be used for either transmitting or receiving streams, or cancelling inter-cell interference. It is therefore plausible to optimize the beamforming weights solely based on the given optimized antenna usage.

Such two-step optimization significantly reduces the cell coordination effort for cancelling the inter-cell interference. First, coordinately optimizing the antenna usage by each AP and client merely requires the information of the number of antennas on all nodes in the network. Without needing any channel knowledge, cell coordination in this step is simplified since the required information can be represented with only a few bits and explicitly shared by each AP with negligible overhead. Second, given the optimized antenna usage, an AP or a client is fully aware of its duty toward cancelling the inter-cell interference. To determine the optimal beamforming weights that fulfill this duty, the AP or client only needs a subset of channel knowledge in the network. Such reduction of the required channel knowledge further makes cell coordination easier. With much simplified cell coordination,

Table 1: The optimal use of the AP and client antennas in our illustrative example, where “ \rightarrow ” and “ \leftarrow ” indicate data communication and inter-cell interference cancellation, respectively.

	Antenna 1	Antenna 2	Antenna 3
AP1	\rightarrow Client1	\leftarrow Client4	\times
AP2	\rightarrow Client3	\rightarrow Client4	\times
Client1	\leftarrow AP1	\leftarrow AP2	\leftarrow AP2
Client3	\leftarrow AP1	\leftarrow AP2	\times
Client4	\leftarrow AP2	\times	\times

the two-step optimization can be integrated into 802.11ac retaining the distributed nature of the protocol. Our proposed solution, called CoaCa, leverages *interleaved channel sounding* and *channel reporting and overhearing* to let each AP and client optimize their antenna usage and beamforming weights in a local but coordinated way.

4. ANTENNA USAGE OPTIMIZATION

In this section, we provide an algorithm that identifies the best antenna usage for each AP and client to maximize the multiplexing gain of the MU-MIMO network. Recall that an AP or a client antenna is used for either delivering streams or cancelling inter-cell interference. Therefore, our algorithm finds the optimal allocation of the AP and client antennas for such two uses. In the following, we first use a simple but illustrative example with two APs and four clients to demonstrate the process of finding the optimal antenna usage, and then provide the algorithm that applies to MU-MIMO networks with arbitrary number of cells.

4.1 Illustrative Example

Our example shown in Figure 4 includes two MU-MIMO cells where each AP is equipped with two antennas and serves up to two clients simultaneously. To find the optimal antenna usage, our algorithm needs to identify the best set of clients in each cell that can be simultaneously served by their corresponding AP. In other words, the selected clients must be capable of coordinately cancelling the inter-cell interference with their interfering AP.

Our algorithm starts by letting a single AP, say AP1, serve both of its clients, and tries to add concurrently served clients in the other cell similarly to $n+$. Clearly, with only two antennas, AP2 cannot cancel the interference to both clients in Cell 1 while serving either of its own clients. Then, unlike $n+$ which simply stops and lets AP1 serve its two clients, our algorithm asks AP1 to serve only a single client, say Client1, and again seeks to add concurrently served clients in Cell 2. Noticeably, now AP2 can serve both of its clients if the inter-cell interference is cancelled in the following way. First, Client1 uses its three antennas to cancel the two interfering streams from AP2. Second, Client3 uses its two antennas to cancel the interfering stream from AP1. Last, while Client4 with a single antenna cannot cancel the interfering stream from AP1, observe that AP1 has a spare antenna that can be just leveraged to cancel the interference to Client4. This way, we have found the best set of clients to serve in each cell that collectively achieves a multiplexing gain of three. We illustrate the optimal use of each AP and client antenna in Table 1.

4.2 Network of Two Cells

We next present our algorithm that identifies the best antenna usage for a two-cell MU-MIMO network with arbitrary number of clients and arbitrary number of antennas on each AP and client.

Our algorithm is motivated by the optimization process for the illustrative example in Section 4.1.

Since we are only interested in congested MU-MIMO networks where the number of associated clients is no smaller than that of the AP antennas, we assume in Cell 1 AP1 has N antennas and the N clients have P_n ($n = 1, 2, \dots, N$) antennas, and in Cell 2 AP2 has M antennas and the M clients have Q_m ($m = 1, 2, \dots, M$) antennas. We further assume the clients in each cell are sorted based on their number of equipped antennas:

$$P_1 \leq P_2 \leq \dots \leq P_N, \\ Q_1 \leq Q_2 \leq \dots \leq Q_M.$$

Algorithm for optimizing the antenna usage. Similar to the example in Section 4.1, to obtain the optimal antenna usage, our algorithm needs to determine the optimal set of served clients in both cells given that the inter-cell interference can be coordinately cancelled. Initially, our algorithm allows only Cell 1 to operate, by letting AP1 serve all its N clients. Then, it seeks to add clients in Cell 2 that can be concurrently served by AP2. The maximum number of added clients in Cell 2, denoted as L , is constrained by the inter-cell interference from AP2 to clients in Cell 1 and that from AP1 to clients in Cell 2. First, AP2 may use up to L' antennas to transmit L' streams to L' clients in Cell 2. While these L' data streams as inter-cell interference may be cancelled by a few clients in Cell 1 with enough ($\geq L'+1$) antennas, the remaining clients in Cell 1 that do not have enough ($\leq L'$) antennas must rely on the spare $M-L'$ antennas on AP2 to cancel the inter-cell interference. We can keep increasing L' until those clients in Cell 1 that must cancel their inter-cell interference do not have enough antennas:

$$L' = \max(n : P_{M-n+1} \geq n + 1). \quad (1)$$

Second, up to L'' clients in Cell 2 can be served by AP2 where these clients must have enough ($\geq N+1$) antennas to cancel the interference from AP1. We can keep increasing L'' until these clients do not have enough antennas:

$$L'' = M - \min(m : Q_m \geq N + 1) + 1. \quad (2)$$

Note, we define $P_{N+1} = Q_{M+1} = +\infty$ to ensure the correctness of Equation (1) and (2) in cases where $L'=0$ and $L''=0$. The maximum number of added clients in Cell 2 is then given by

$$L = \min(L', L''). \quad (3)$$

In the next step, we let AP1 remove clients from its served set. When AP1 removes K ($K = 1, 2, \dots, N$) clients, these clients must have the fewest (P_1, P_2, \dots, P_K) antennas. Afterwards, AP1 has K spare antennas that can be exploited to cancel the interference to clients in Cell 2. Naturally, up to K clients in Cell 2 can be served by AP2 not subject to the interference from AP1. Since such interference cancellation is performed by AP1, clients in Cell 2 with the fewest (Q_1, Q_2, \dots, Q_K) antennas can be picked to enjoy such benefit. Therefore, Cell 2 is left with $M-K$ clients with Q_{K+1}, \dots, Q_M antennas, and Cell 1 is left with $N-K$ clients with P_{K+1}, \dots, P_N antennas. Similar to the previous step where $K=0$, we have

$$L'(K) = \max(n : P_{K+M-n+1} \geq n + 1), \quad (4)$$

$$L''(K) = M - \min(m : Q_{K+m} \geq N - K + 1) + 1, \quad (5)$$

and

$$L(K) = \min(L'(K), L''(K)). \quad (6)$$

It is easy to verify that $L(N)=M$ where only a single cell (Cell 2) is operating similar to 802.11ac.

Finally, for each $K=0, 1, \dots, N$ our algorithm calculates $L(K)$, and finds the optimal set of clients ($N-K$ in Cell 1, L in Cell 2) that maximizes the number of streams $N-K+L(K)$. Given the optimal set of clients in each cell and their duty in cancelling the inter-cell interference, the optimal use of the AP and client antennas is meanwhile determined.

We must note that our algorithm is orchestrated to maximize the number of streams or clients without (i) considering fairness between clients, and (ii) avoiding channel hardening that may reduce network capacity when serving more clients [18, 22]. That is, our algorithm preferably selects clients with more antennas over those with fewer antennas, and lets each AP simultaneously serve all selected clients. To consider these two issues, client selection schemes can be combined with our algorithm without modification to the latter. First, to take client fairness into account, an AP can remove clients with lower priority from the selected set even though they have more antennas, and sort clients with the same number of antennas based on their relative priority. Second, to avoid channel hardening to occur, an AP can select clients further based on the historical observations of their channel orthogonality, which is the key to determine the capacity scaling toward the number of served clients. Additionally, even among the selected clients, an AP has the freedom to serve only a subset of them, after receiving their reported channels and more accurately evaluating their channel orthogonality. Clearly, such client selection schemes can be executed separately, before or after our algorithm outputs the best set of clients and their antenna usage.

4.3 Network of More than Two Cells

We next extend our algorithm to T ($T > 2$) cells. With T cells, we need to rewrite N and K as vectors, meaning that the $T-1$ APs serve ($N_1-K_1, \dots, N_{T-1}-K_{T-1}$) clients respectively. Then, the maximum number of clients that can be added in the last cell is given by

$$L'(K_1, \dots, K_{T-1}) = \max(n : P'_{M-n+1} \geq n + 1), \quad (7)$$

$$L''(K_1, \dots, K_{T-1}) = |\{Q'_m : Q'_m \geq \sum_{t=1}^{T-1} (N_t - K_t) + 1\}|, \quad (8)$$

and

$$L(K_1, \dots, K_{T-1}) = \min(L', L''), \quad (9)$$

where “ $|\cdot|$ ” represents the cardinality of a set, P'_n the number of spare antennas on client n , and Q'_m the equivalent number of antennas on client m , determined by K'_t , the number of spare antennas on AP t . The spare antennas on an AP or a client are defined as the remaining antennas after a few of them are used to cancel the inter-cell interference within the first $T-1$ cells. The equivalent number of antennas Q'_m is defined in the following way. For a given client in the last cell, if an AP in the previous $T-1$ cells uses one spare antenna to cancel the interference to the client, the client can be considered to equivalently have $N_t - K_t$ additional antennas for canceling this AP's interference. In other words, one can “transfer” the interference cancellation capability from an AP to a client it interferes with by providing the client such equivalent antennas. Given Q'_m , the best allocation of the spare antennas on the $T-1$ APs is given by successively assigning those of each AP to the K'_t clients in the last cell that have the smallest Q'_m .

Given the definitions of P'_n , Q'_m and K'_t , we next explain how Equation (7) and (8) are derived. For Equation (7), notice that the last AP can only cancel the interference to $M-L'$ clients in all previous $T-1$ cells. These clients who may belong to more than one cell must have the fewest spare antennas. The remaining

clients must have enough spare antennas to cancel the L' streams from the last AP. For Equation (8), the served clients in the last cell should not be subject to the interference from all previous $T-1$ APs. Since we have considered the interference cancellation from the $T-1$ APs with their spare antennas, we just need to find the maximum number of clients in the last cell with enough antennas, i.e., $Q'_m \geq \sum_{t=1}^{T-1} (N_t - K_t) + 1$. After calculating L for each (K_1, \dots, K_{T-1}) , we can find the optimal set of clients in each cell with the maximum multiplexing gain.

It is noticed that the above algorithm works in a recursive manner. Therefore, the complexity exponentially increases with the number of cells, T . To deal with this scalability issue, a cell clustering technique can be leveraged to convert a large-scale MU-MIMO network into a few small clusters, each of which includes up to three cells. We elaborate the cell clustering technique in Section 7.

4.4 Practical Implications

It is observed that our proposed algorithm only requires the information about the number of antennas on each AP and client to identify the optimal antenna usage. Such information is global in the MU-MIMO network, but can be compactly represented with only a few bits. Therefore, explicitly sharing such information in the network does not incur noticeable overhead, significantly simplifying the cell coordination. Given such information, each AP and client can execute the same algorithm in a synchronized way to achieve coordination. In section 6.1, we discuss how CoaCa leverages interleaved channel sounding to easily provide each AP and client such information with standard 802.11ac control frames.

5. CHANNEL ANALYSIS FOR BEAMFORMING WEIGHT OPTIMIZATION

In this section, we analyze the channel knowledge requirement for an AP or a client to optimize its transmit or receive beamforming weights. To calculate the beamforming weights that enable the optimal antenna usage, an AP or client must have certain channel knowledge based on which it cancels the intra-cell and inter-cell interference using the beamforming techniques presented in Section 2. In the following, we first study the two-cell example in Section 4.1 in order to simplify the analysis and obtain insightful findings, and then extend our analysis to a MU-MIMO network with arbitrary number of cells.

5.1 Illustrative Example

We reuse the example in Figure 4 to study the channel knowledge each AP and client needs to compute its optimal beamforming weights. We use $\mathbf{H}_{i \rightarrow j}$ ($\mathbf{h}_{i \rightarrow j}$) to denote the channel matrix (vector) from AP i to client j , and \mathbf{w}_j , \mathbf{v}_j to denote the AP's transmit beamforming weight vector for client j , and the receive beamforming weight vector of client j , respectively.

Channel knowledge for AP1. AP1 uses its two antennas to send a data stream to Client1 and cancel the inter-cell interference to Client4. To do this, \mathbf{w}_1 must be orthogonal to $\mathbf{h}_{1 \rightarrow 4}$, i.e., $\mathbf{w}_1 = \mathbf{h}_{1 \rightarrow 4}^\perp$ where “ \perp ” represents the null space of a vector. As a result, AP1 only needs the knowledge of $\mathbf{h}_{1 \rightarrow 4}$.

Channel knowledge for AP2. AP2 performs zero-forcing beamforming to simultaneously send streams to Client3 and Client4 without cancelling inter-cell interference. To do this, AP2 only needs the knowledge of $\mathbf{H}_{2 \rightarrow 3} \mathbf{v}_3$ and $\mathbf{h}_{2 \rightarrow 4}$. Note that $\mathbf{H}_{2 \rightarrow 3} \mathbf{v}_3$ is the beamformed channel of Client3: it combines the physical channel matrix from AP2 to Client3 ($\mathbf{H}_{2 \rightarrow 3}$) and the determined beamforming weight vector of Client3 (\mathbf{v}_3) as a single

channel vector ($\mathbf{H}_{2 \rightarrow 3} \mathbf{v}_3$). Compared to a physical channel matrix, a beamformed channel vector is much more efficient for a client to report since it needs fewer bits to be represented [18]. To cancel the intra-cell interference with zero-forcing beamforming, the knowledge of such beamformed channels is enough for AP2 by setting $\mathbf{w}_4 = (\mathbf{H}_{2 \rightarrow 3} \mathbf{v}_3)^\perp$ and $\mathbf{w}_3 = \mathbf{h}_{2 \rightarrow 4}^\perp$.

Channel knowledge for Client1. Client1 uses its three antennas to receive its stream from AP1 and cancel the inter-cell interference from AP2. Since AP2 sends two streams to Client3 and Client4, Client1 needs to cancel both of them. To do this, Client1 simply cancels the signals sent from the two physical antennas at AP2. In other words, \mathbf{v}_1 is chosen as $\mathbf{v}_1 = \mathbf{H}_{2 \rightarrow 1}^\perp$ where “ \perp ” here refers to the joint null space of all rows of the matrix. Consequently, the required channel knowledge for Client1 is restricted to its own channels from AP2, $\mathbf{H}_{2 \rightarrow 1}$.

Channel knowledge for Client3. Unlike Client4 who has a single antenna and needs not cancel the inter-cell interference, Client3 uses its two antennas to receive its stream from AP2 and cancel the inter-cell interference from AP1. Therefore, \mathbf{v}_3 must be orthogonal to the signal vector from AP1, i.e., $\mathbf{v}_3 = (\mathbf{w}_1 \mathbf{H}_{1 \rightarrow 3})^\perp$, which suggests \mathbf{w}_1 is needed for Client3 to calculate \mathbf{v}_3 . However, observe that $\mathbf{w}_1 = \mathbf{h}_{1 \rightarrow 4}^\perp$ so that the required channel knowledge for Client3 actually becomes $\mathbf{h}_{1 \rightarrow 4}$ and its own channel from AP1, $\mathbf{H}_{1 \rightarrow 3}$.

5.2 Network of Two Cells

Motivated by the findings from the illustrate example, we next analyze the channel knowledge requirement of the beamforming weight optimization for a two-cell MU-MIMO network with arbitrary configuration. In particular, we prove three key theorems regarding such requirement that can be summarized as follows: to calculate the optimal beamforming weights based on the optimal antenna usage, an AP or a client only needs the channel knowledge owned by a particular set of clients in the network. With this requirement, we can not only reduce the cell coordination effort, but also guarantee the optimality of the computed beamforming weights. We next elaborate the three theorems:

THEOREM 1. *To calculate the optimal beamforming weights, an AP only needs the channel knowledge owned by the clients it serves, and the clients it interferes with holding the interference cancellation responsibility.*

THEOREM 2. *To calculate the optimal beamforming weights, a client only needs the channel knowledge owned by certain clients in the same cell.*

THEOREM 3. *For clients in the same cell, there exists a proper order of them with which each client can calculate the optimal beamforming weights solely based on the channel knowledge owned by clients ranked before it.*

To prove these theorems, let us consider a two-cell network after antenna usage optimization, where AP1 having N antennas serves $N-K$ clients and AP2 having M antennas serves $M-J$ clients (see Section 4.2). Recall that the optimal use of an AP antenna is to either deliver a stream to a served client, or cancel the inter-cell interference to a client the AP interferes with. As a result, we can partition the antennas on AP1 into two sets: $N-K$ antennas used to serve $N-K$ clients in Cell 1, and K antennas used to cancel the interference to K clients in Cell 2. Similar antenna partitioning can be applied to AP2. Afterwards, clients in each cell can be also partitioned according to their responsibility of cancelling the inter-cell interference. In Cell 1 and Cell 2, J and K clients rely on

AP2 and AP1 to cancel the inter-cell interference, while the rest $N-K-J$ and $M-J-K$ clients use their own antennas to cancel the interference, respectively.

Proof of Theorem 1. We use AP1 for the proof of Theorem 1. To serve the $N-K$ clients in Cell 1 and cancel the interference to the K clients in Cell 2, AP1 only needs to know the channels from itself to these clients. The channels from AP2 to these clients are not needed since AP1 is not involved in the inter-cell interference from AP2. The channels from AP1 to the $M-J-K$ clients in Cell 2 are also unnecessary, since these clients use their own antennas to cancel the interference.

Proof of Theorem 2 and Theorem 3. We use clients in Cell 1 for the proof of Theorem 2 and 3. First, the J clients do not need any channel knowledge owned by other clients to optimize its beamforming weights. This is because their antennas do not contribute to inter-cell interference cancellation. Instead, they can be used to improve the client SNR by employing MRC based on their own channels $\mathbf{H}_{1 \rightarrow j}$. We call these J clients the *MRC clients*. Second, the $N-K-J$ clients need certain channel knowledge owned by other clients to cancel the interference from AP2. To do this, they perform interference alignment toward the channels of the other J clients, so that

$$\begin{aligned} & \text{span}(\mathbf{H}_{2 \rightarrow 1} \mathbf{v}_1, \dots, \mathbf{H}_{2 \rightarrow J} \mathbf{v}_J) = \\ & \text{span}(\mathbf{H}_{2 \rightarrow J+1} \mathbf{v}_{J+1}, \dots, \mathbf{H}_{2 \rightarrow N-K} \mathbf{v}_{N-K}). \end{aligned} \quad (10)$$

Through interference alignment, the beamformed channels of the $N-K-J$ clients, which we call the *IA clients*, are aligned to the channels of the MRC clients. When there are no MRC clients, the IA clients simply cancel the interference from the M physical antennas on AP2. When AP2 cancels the interference to the J MRC clients, the signal vector must be perpendicular to $\text{span}(\mathbf{H}_{2 \rightarrow 1} \mathbf{v}_1, \dots, \mathbf{H}_{2 \rightarrow J} \mathbf{v}_J)$, which meanwhile creates no interference to the $N-K-J$ IA clients. Clearly, the IA clients only need the knowledge of the beamformed channels, $\mathbf{H}_{2 \rightarrow j} \mathbf{v}_j$, from the MRC clients, and the optimal client order is given by ranking the MRC clients before the IA clients. The relative order between the MRC clients or between the IA clients does not have an impact.

5.3 Network of More than Two Cells

The above three theorems hold true for a MU-MIMO network with more than two cells, which we briefly explain as follows. First, Theorem 1 is self-explanatory given its proof for the two-cell network. Second, for Theorem 2, observe that for a given interfering AP, the process of partitioning the clients into MRC clients and IA clients is still feasible. Then, a client only needs the channel knowledge from those that are identified as MRC clients while the client itself is identified as an IA client. Apparently such channel knowledge is restricted to the client's own cell. For Theorem 3, the best client order is decided by the number of antennas the client carries in an increasing manner. This is because an IA client always has more antennas than a MRC client does. Then, after sorting all clients in a cell based on their number of antennas, a client only needs the channel knowledge from clients that are ranked before it.

5.4 Practical Implications

The analysis on the channel knowledge requirement tells us that the beamforming weight optimization for an AP or a client can be potentially performed in a distributed way due to reduced cell coordination. Theorem 1 indicates that an AP does not need the channel knowledge owned by all clients, and does not need to share its channel knowledge with other APs. Theorem 2 suggests

that a client does not need to acquire any channel knowledge from clients in other cells. Theorem 3 implies that even in a single cell, a client only needs the channel knowledge from a particular set of clients. In Section 6.2, we elaborate how CoaCa leverages channel reporting and overhearing to provide each AP and client the necessary channel knowledge for its beamforming weight optimization, without incurring any coordination overhead.

6. INTEGRATION WITH 802.11ac

We next present CoaCa, a protocol that integrates both antenna usage optimization and beamforming weight optimization into 802.11ac. CoaCa includes two key designs to achieve coordinated inter-cell interference cancellation, namely *interleaved channel sounding* and *channel reporting and overhearing*. With these two designs, each AP and client can locally perform the two-step optimization in a distributed but coordinated way.

6.1 Interleaved Channel Sounding

To provide the APs and clients with necessary information to optimize their antenna usage, CoaCa proposes *interleaved channel sounding*, in which the key idea is to let the APs in all cells send their NDP-A and NDP frame sequentially one after another, before each AP polls its served clients. The NDP-A frame can contain the information about the number of antennas on the AP and that on the clients the AP plans to serve. With interleaved channel sounding, such information can be timely broadcast to the entire MU-MIMO network, and used by all APs and clients to optimize their antenna usage with the same algorithm provided in Section 4 in a coordinated way. Moreover, the NDP frames allow the IA clients to estimate their channel matrix from the interfering APs, which is necessary to optimize their beamforming weights together with the overheard channel vectors from the MRC clients (see Section 6.2).

We illustrate the timeline of interleaved channel sounding in Figure 5 using the example in Figure 4. In the channel sounding process of 802.11ac, after an AP, say AP1, sends the NDP-A and NDP frame, the first served client, say Client1, will immediately respond with the BF-R frame after SIFS time. Unlike 802.11ac, in CoaCa no clients in Cell 1 is allowed to immediately respond; instead, AP2 sends its NDP-A and NDP frame SIFS time after AP1 sends its NDP frame. Only after both AP1 and AP2 send their NDP-A and NDP frame, each of them sequentially polls their clients with a BF-P frame and their clients respond with their BF-R frames in the same order.

Viability. The interleaved channel sounding ensures correct behaviour of the involved APs and their served clients. To interleave the channel sounding from multiple APs, an AP must be able to send its NDP-A and BF-P frame at the proper time with guaranteed medium access. In addition, the APs must sound their channels in a pre-determined order without introducing collision. For example, according to Figure 5, AP2 must (i) send its NDP-A frame immediately after AP1 sends its NDP frame, and (ii) poll its clients only after Client1 has sent its BF-R frame. CoaCa uses two techniques to ensure this coordinated behaviour, which we discuss based on Figure 5. First, CoaCa adopts CHAIN, a technique proposed in [24]. The key idea in CHAIN is that AP2 piggybacks its NDP-A or BF-P frame SIFS time after the ongoing frame from AP1 or Client1 finishes; this gives AP2 prioritized medium access since other 802.11 nodes must wait for at least DIFS time to contend for the medium. To determine their relative channel sounding order, AP1 and AP2 only need to coordinate once to initiate the transmissions in CHAIN. Second, to avoid collision between AP2 and Client1 who might also send its BF-R frame after

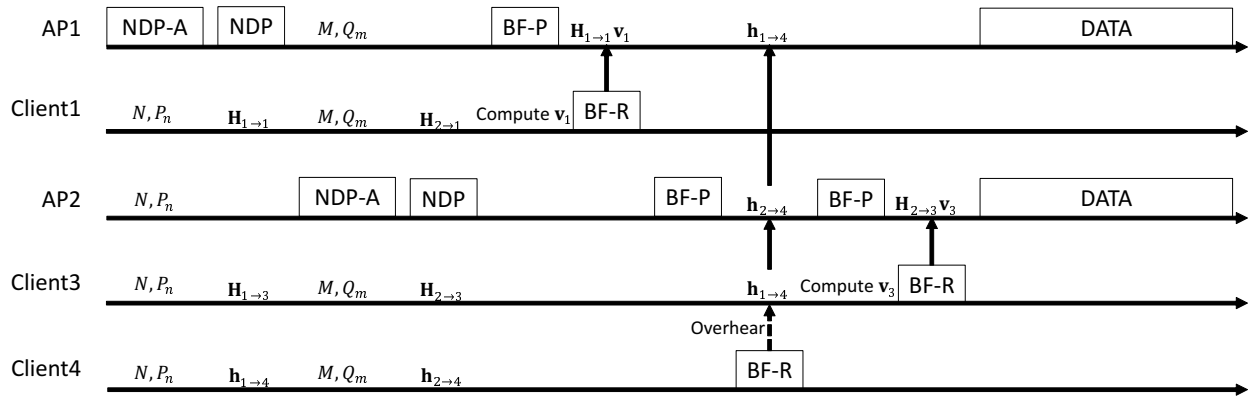


Figure 5: Timeline of CoaCa where two APs concurrently serve their clients. To provide each AP and client the necessary information to optimize its antenna usage, the first AP starts polling its clients with the optimal order only after all the APs have transmitted their NDP-A and NDP frames. To optimize its beamforming weights based on the optimal antenna usage, a client overhears the reported beamformed channels from other clients in the same cell that report before it.

SIFS time, CoaCa refrains Client1 from immediately responding. This is achieved by having AP1 specify a “fake” client with an invalid MAC address as the first responding client in the NDP-A frame. This way, all clients in Cell 1 yield to AP2.

Interoperability. CoaCa APs and clients can interoperate with unmodified 802.11ac clients. The key reason is that in the proposed interleaved channel sounding a CoaCa client still passively responds to the BF-P frame from its AP, similar to an unmodified 802.11ac client. In addition, CoaCa concatenates the number of antennas information (2 bytes) to the NDP-A frame and modifies the duration field, to ensure it can be decoded even by an unmodified 802.11ac client.

Overhead. The interleaved channel sounding introduces negligible overhead compared to 802.11ac. First, observe that in the interleaved channel sounding an extra BF-P frame from each AP is required to poll the first served client. Such overhead is not only negligible but also justified since (i) the BF-P frame is much shorter than the NDP-A and BF-R frame that constitute the major portion of channel sounding, and (ii) the extra BF-P frame eliminates the necessity for the first client to initiate the transmissions in CHAIN, which would otherwise require the client to overhear clients in other cells and negate the reduction of cell coordination in CoaCa. Second, the BF-R frame may have to contain the estimated channels from not only the associated AP but also the interfering APs. Such extra channel information seemingly increases the size of the BF-R frame in a proportional manner. However, due to the use of beamformed channels (see Section 6.2), such extra channel information can be more than compensated by the proportionally reduced information in each estimated channel, once the client is equipped with multiple antennas.

6.2 Channel Reporting and Overhearing

To provide the APs and clients with required channel knowledge to optimize their beamforming weights, similarly to the technique adopted by [18], in CoaCa a client reports the necessary beamformed channels in the BF-R frame and overhears other clients’ BF-R frames for their beamformed channels, as illustrated in Figure 5. First, reporting the beamformed channel $\mathbf{H}_{i \rightarrow j} \mathbf{v}_j$ instead of the physical channel $\mathbf{H}_{i \rightarrow j}$ can proportionally reduce the size of the BF-R frame which is known to incur substantial overhead [5, 21] to the channel sounding process of 802.11ac. This is because the beamformed channel is a vector while the physical channel is a matrix that needs many more bits to encode when

the client has multiple antennas. Reducing the size of the BF-R frame is especially beneficial for the MRC clients who must report the channels to not only its associated AP but also the interfering APs holding the interference cancellation responsibility (see Theorem 1). Such extra channel information as additional overhead can be more than compensated by the reduced size of the BF-R frame when the client has multiple antennas. Second, since the beamformed channels are explicitly contained in the BF-R frame, a client that overhears the BF-R frame can easily acquire such knowledge by decoding the frame. The channel knowledge is guaranteed accurate once the BF-R frame is successfully decoded.

Decodability. An AP or a client is able to decode the overheard BF-R frames with high probability, given the following two observations. First, Theorem 1 and Theorem 2 indicate that an AP only needs to overhear the BF-R frames from the MRC clients it interferes with, and a client only needs to overhear the BF-R frames from the MRC clients in the same cell. This significantly reduces the likelihood that an AP or a client is too distant from the client it seeks to overhear. Second, the BF-R frame is considered a control frame and commonly sent at base rate (6 Mbps) in order to improve its reliability [4]. This in turn extends its transmission range and reduces the possibility of frame decoding failure.

Sufficiency. Reporting and overhearing the beamformed channels can provide the APs and clients just enough channel knowledge. First, according to Section 5.2, the beamformed channel is sufficient for the APs to perform interference cancellation, and for the IA clients to perform interference alignment. Second, the AP can poll its clients in the optimal order to make sure a client has acquired enough beamformed channels before it determines the optimal beamforming weights and reports its own beamformed channel. This is because Theorem 3 indicates that if the clients in each cell send their BF-R frames following the optimal order, each client only needs the channel knowledge owned by clients ranked before it.

7. CELL CLUSTERING

In this section, we address the scalability issue of CoaCa toward the number of cells in the MU-MIMO network. In particular, the following reasons make it hard to apply CoaCa to a large-scale network with more than a few cells. First and most importantly, the number of required antennas on the APs and clients to appreciate the multiplexing gain improvement from

CoaCa significantly increases with more cells. Given that each AP and each client usually cannot afford more than eight and more than four antennas respectively, we find that CoaCa cannot provide obvious capacity improvement when the network includes more than three cells. Second, as revealed in Section 4.3, the complexity of the recursive algorithm to identify the optimal antenna usage exponentially increases with the number of cells. Third, the duration of the interleaved channel sounding scales up proportionally to the number of cells, which may lead to having outdated estimated channels. Finally, to receive the NDP-A frames and enable CHAIN, an AP must be in the overhearing range of all other APs, which becomes much more challenging with more cells.

We rely on a *cell clustering* technique to tackle the scalability issue of CoaCa. That is, we group all cells in the MU-MIMO network into clusters such that *within each cluster there are up to three cells*. Then, we apply CoaCa to each cluster, and employ standard CSMA/CA for the medium access between clusters. There are two important motivations for us to allow up to three cells in each cluster. First, cell clustering with such scale does not reduce the effectiveness of CoaCa. As mentioned earlier, when an AP has no more than eight antennas and a client has no more than four antennas, to maximize the number of streams CoaCa usually does not need to allow more than three APs to concurrently serve their clients. Therefore, even if the MU-MIMO network included more than three cells, they would not lead to a multiplexing gain increase. Notably, CoaCa seeks to most effectively leverage the available antennas from each AP and client in the network, by letting them coordinately cancel the inter-cell interference in a distributed way. Yet, it is incapable of going beyond such number of antenna constraint. To further increase the number of streams in the network, one must employ centralized solutions such as Network-MIMO [2, 12, 19] to convert interference into signals. Second, cell clustering with up to three APs can greatly simplify the operation of CHAIN in situations where only a subset of APs in a cluster intend to serve their clients. That is, in CHAIN, to ensure correct medium access the AP without packets to its clients may have to send an intermediate frame that triggers the NDP-A frame from the AP following it. With up to three APs in a cluster, this can be replaced by a much simpler design where the second and third AP use SIFS (16 μ s) and PIFS (25 μ s) in CHAIN to guarantee their relative priority, without the need of sending the triggering frame by the second AP.

Determining the included cells in a cluster can be accomplished by a cluster formation algorithm, which is outside of the scope of this work. We believe existing approaches based on assigning a master AP for managing each cluster, e.g., the one proposed by [25], can be adopted with appropriate adaptation.

8. EXPERIMENTAL EVALUATION

Finally we experimentally evaluate the performance of CoaCa in real-world indoor environments.

8.1 WARP-based Implementation

We implement CoaCa on WARP [14], a flexible software defined radio (SDR) platform. We choose the WARPLab framework [13] for our implementation and evaluation but have extensively modified it to improve its real-time capability. In WARPLab, through gigabytes Ethernet cables, multiple WARP boards configured as APs or clients are interconnected to and controlled by a central computer that runs MATLAB. By implementing part of the baseband processing in MATLAB, one can achieve great flexibility to develop and evaluate physical layer techniques such as beamforming; this feature of WARPLab makes

it a desirable framework for our implementation and evaluation of CoaCa.

Since the baseband processing occurs in MATLAB, we cannot implement the 802.11ac MAC satisfying the timing constraint. Therefore, our evaluation of CoaCa is focused on its PHY, i.e., the network capacity improvement from the increased multiplexing gain achieved by antenna usage optimization and beamforming weight optimization. We move the process of client polling, beamformed channel reporting and overhearing in the interleaved channel sounding into emulation. That is, after each AP sequentially sends its NDP-A and NDP frame that allows the clients to estimate their channels, we emulate the BF-P and BF-R frames on the central computer. To do this, we assume each AP and client possesses the same knowledge of the reported and overheard beamformed channels, as if they were obtained from actual BF-R frames over the air. Note that such emulation is close-to-reality since the BF-R frame delivers the channel knowledge in an explicit and reliable way. Then, based on the channel knowledge, the APs and clients calculate their optimal beamforming weights, with which they coordinately cancel the inter-cell interference, and transmit and receive the data frames. We calculate the network capacity based on the measured signal-to-noise-and-interference ratio (SINR) at each served client. While the actual throughput gain by MU-MIMO is known to be reduced by the channel sounding overhead [5, 21], we do not incorporate such overhead into our evaluation since it equally exists in 802.11ac according to Section 6.1.

8.2 Experimental Setup

In our experiments we construct a MU-MIMO network with two cells. We use each WARP board to serve as either an 802.11ac AP or an 802.11ac client so that the number of AP and client antennas can be up to four. To evaluate the capability of CoaCa to eliminate inter-cell interference, we deploy the APs and clients within a single interference domain so that the capacity of each served client is interference instead of noise limited. Our experiments are conducted in a three-floor campus building that is representative of indoor environments. To make sure the interleaved channel sounding does not exceed the channel coherence time, we run the experiments at nighttime when we observe the wireless environment is static. To limit the interference within our MU-MIMO network, we select a clean Wi-Fi channel (# 14) that does not have any ongoing traffic.

8.3 Accuracy of Interference Cancellation and Interference Alignment

Since CoaCa relies on interference cancellation and interference alignment to eliminate inter-cell interference, we first evaluate their accuracy on WARP under realistic indoor wireless channels. Note that interference cancellation and alignment can be inaccurate due to (i) errors in the estimated and overheard channels used to compute the beamforming weights, and (ii) channel noise and hardware imperfection that distort the received signals. Toward this, we use three micro-benchmarks with two or three nodes to study how much inter-cell interference between an AP and a client can be eliminated. To evaluate interference cancellation by the AP, we assume an AP having two to four antennas cancels the interference to a client via transmit beamforming; to evaluate interference cancellation by the client, we assume a client having two to four antennas cancels the interference from an interfering AP with receive beamforming; to evaluate interference alignment by the client, we assume an AP cancels the interference to one client via transmit beamforming, and another client aligns its




Figure 6: Interference reduction from interference cancellation by the AP (Left), interference cancellation by the client (Middle), and interference alignment by the client (Right). On average interference cancellation and interference alignment achieve 30 dB and 25 dB interference reduction, respectively. The number of AP or client antennas does not have an obvious impact on the cancellation or alignment accuracy.

channel to that of the first client via receive beamforming. In each micro-benchmark we assume the beamforming AP or client owns the necessary channel knowledge, which is obtained right before the node transmits or receives. For all the experiments, we measure the original interference power and residual interference power after cancellation or alignment to compute the interference reduction. Note that we have properly scaled the transmit power from the AP to make sure the interference power is always above the noise floor of the client even after cancellation. This is to avoid underestimation of the effectiveness of interference cancellation and alignment when the client becomes noise limited.

Figure 6 shows the interference reduction (dB) achieved by interference cancellation and interference alignment performed by the AP or client, with different number of antennas. We highlight the following findings from Figure 6. First, on average interference cancellation and alignment are able to reduce the interference by 30 dB and 25 dB, respectively. This is often sufficient to reduce the interference to below the noise floor in 802.11ac networks. Second, interference cancellation by the AP via transmit beamforming, and by the client via receive beamforming have similar accuracy. This is because with explicit channel estimation the AP and client own the same channel knowledge, and therefore the accuracy of interference cancellation is not subject to channel calibration errors that only exist in systems with implicit channel estimation such as [16]. Third, the number of antennas for interference cancellation and alignment does not have a clear impact on the accuracy. While somewhat counter-intuitive, this finding can be explained by the observation that additional antennas cannot reduce channel estimation error or hardware nonlinearity. Last, the accuracy of interference alignment is lower than that of interference cancellation. This is because interference alignment suffers more from channel estimation errors by leveraging the channel knowledge from more than one clients. Consistent findings regarding the effectiveness of interference cancellation and alignment are also reported in [10].

8.4 Network Capacity Improvement

We next evaluate the effectiveness of CoaCa on improving the capacity of a MU-MIMO network by achieving a higher multiplexing gain. We compare CoaCa with two existing schemes: (i) MACCO [18] which does not cancel inter-cell interference by allowing a single AP to transmit and using the client antennas to improve the channel orthogonality and capacity scalability; (ii) n+ [10] which allows an AP in one cell to serve its clients first and the AP in the other cell to opportunistically transmit at the same time. Note, given that each cell is congested, n+ cannot

Table 2: Number of antennas on the APs and clients for Case 1-4.

	AP1	AP2	Cell 1 clients	Cell 2 clients
Case 1	2	2	2/2	2/2
Case 2	2	2	1/2	1/3
Case 3	2	3	4/4	3/3/4
Case 4	4	4	1/2/4/4	1/2/4/4

increase the multiplexing gain and may only provide a diversity gain as explored in [7]. Therefore, for a given channel condition we compare CoaCa with either MACCO or n+, whichever achieves better network capacity. For simplicity we use “MACCO/n+” to denote the better scheme among the two of them. We select four representative cases with different numbers of antennas on the APs and clients, summarized in Table 2. For each case, we run multiple instances of the experiment by deploying the APs and clients in different locations. It is important to mention that CoaCa improves the network capacity by aiming to deliver more streams in the MU-MIMO network, which is known to be a suboptimal approach. Therefore, the actual capacity improvement of CoaCa is dependent on the channel condition, i.e., the orthogonality between the channels of multiple served clients. Conducting the experiment under various channel conditions allows us to observe not only the average but also the worst-case performance of CoaCa. To randomize the channel conditions, we arbitrate the AP and client locations in each experiment instance, since previous work has observed that in an indoor environment with rich multipath the channel condition does not have a clear dependence on the node locations [1, 18]. This way, our experiments cover channel conditions that are either beneficial or adverse to CoaCa.

Figure 7 shows our results. In each plot, one data point corresponds to the measured capacity for CoaCa (X axis) and MACCO/n+ (Y axis) under a single experiment instance. The expected multiplexing gain increase from CoaCa defined as the increase of the number of streams in the two MU-MIMO cells, is also plotted for comparison. We report the following important findings from Figure 7. First, when CoaCa delivers more streams it considerably improves the network capacity, i.e., on average by 40%, 52% and 41% for Case 2, 3 and 4, respectively. When CoaCa cannot deliver more streams due to insufficient antennas on the APs or clients, e.g., in Case 1, on average it achieves similar network capacity to MACCO/n+. Second, the network capacity improvement from CoaCa is lower than the expected multiplexing gain increase (50%, 67%, and 50% for case 2, 3, and 4, respectively). This is because CoaCa is theoretically suboptimal and cannot proportionally increase the network capacity due to its two-step optimization. The per-client capacity can be reduced when the antennas are used to cancel the inter-cell interference instead of enhancing the client SNR. Finally, even in Case 2, 3 and 4 where CoaCa allows more streams, there is a small probability that CoaCa does not increase the network capacity. This happens when the antenna usage optimization yields few spare AP and client antennas to enhance the SNR, and the channel orthogonality between clients are far from being orthogonal, such that the capacity of each cell suffers from serious channel hardening. Such worst-case performance confirms that due to the use of two-step optimization as a key heuristic CoaCa does not necessarily maximize the network capacity. However, given its considerable capacity gain on average, we believe the heuristic that makes CoaCa practical outweighs its worst-case performance.

Dependency on the client SNR. We next evaluate how the client SNR impacts the network capacity gain from CoaCa. We need to

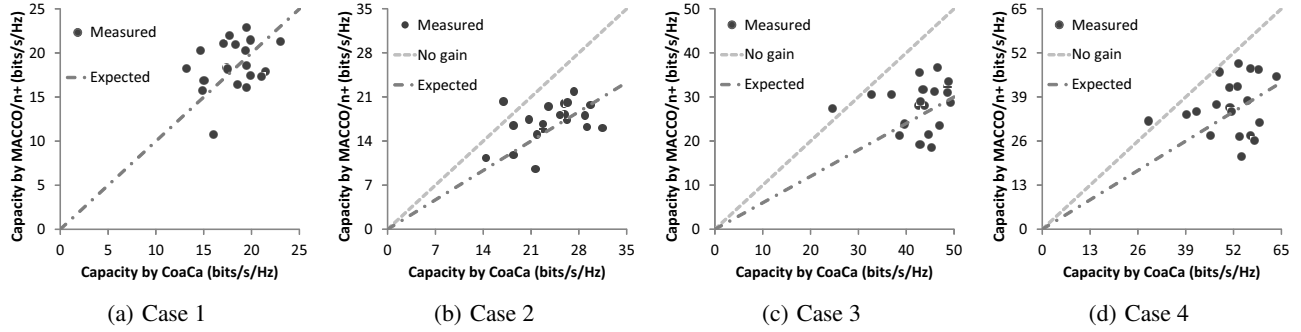


Figure 7: Network capacity achieved by MACCO/n+ and CoaCa for four cases with different numbers of antennas on APs and clients. CoaCa outperforms MACCO/n+ when it delivers more streams. However the capacity improvement from CoaCa is lower than the expected multiplexing gain increase due to imperfect channel orthogonality.




Figure 8: Network capacity achieved by n+, MACCO, and CoaCa with high (measured) and low (emulated) client SNR. CoaCa achieves more capacity improvement with higher client SNR, due to the logarithmic relationship between capacity and SNR.

mention that when the client capacity is interference-limited, the client often has a moderate or high SNR (about 15-25 dB according to our measurements). Therefore we emulate a low SNR regime for each case above, by reducing the transmit power on the two APs by approximately 15 dB. We demonstrate the results in Figure 8 where the average network capacity for n+, MACCO, and CoaCa are compared for both high (measured) SNR and low (emulated) SNR regimes. Clearly, when the client SNR is reduced, CoaCa achieves less network capacity gain: for Case 2-4 the improvement decreases to 28%, 35%, and 30%, respectively. This is because with lower SNR, the logarithmic relationship between capacity and SNR becomes closer to linear, and cancelling the inter-cell interference reduces more capacity with a diminished SNR.

9. RELATED WORK

MU-MIMO techniques such as zero-forcing beamforming have been demonstrated practical in real-world wireless environments using software-defined radio (SDR) platforms e.g., [1, 15–17, 22]. The authors in [1] investigated the feasibility of zero-forcing beamforming in real indoor environments, with a small number of antennas on the AP (≤ 4) and a few clients. The authors in [17] proposed a solution that allows uncoordinated, unsynchronized spatial multiple access from multiple clients. The authors in [16] and [22] study the base station (AP) architecture that features a large number of antennas (≥ 8) and serves many clients. All these work is restricted to a single cell including one AP and multiple associated clients. Inter-cell interference as an important capacity

limiting factor for MU-MIMO networks remains experimentally under-explored.

An effective but practically expensive way to tackle inter-cell interference in MU-MIMO networks is to let neighbouring APs collaborate with each other via a high-speed, low-latency connection, e.g., [2, 3, 8, 9, 12]. There are different ways of exploiting such AP connection. Network-MIMO [2, 12] allows multiple neighbouring APs to behave as a single massive AP, by synchronizing their time and frequency, and sharing their transmitted and received samples. IAC proposed in [8] only needs the APs to share the samples for joint encoding and decoding without synchronization. Robinhood proposed in [3] requires a set of APs to retransmit the received packets from the clients to the rest of APs, while cancelling the interference between the packets. OpenRF proposed in [9] uses a central controller to coordinate neighbouring APs to cancel inter-cell interference. Due to the requirement of AP connection and a centralized solution, the applicability of these approaches is restricted to carefully planned and centrally controlled networks such as an enterprise network. By leveraging multi-antenna clients to assist the interfering AP to coordinately cancel inter-cell interference, our solution is distributed and can be applied to any 802.11ac networks.

802.11n+ presented in [10] leverages beamforming to enable concurrent streams in 802.11n networks with SU-MIMO. It allows nodes in one cell to transmit and receive first, and nodes in other cells to opportunistically transmit and receive by cancelling the interference to and from the first cell. Due to the lack of joint coordination between multiple cells, n+ cannot work effectively in 802.11ac networks with MU-MIMO as we demonstrated in Section 1 using the example in Figure 1. By achieving better coordination, our solution can enable more streams than n+ does in 802.11ac networks, only requiring small protocol modifications and incurring negligible overhead. NEMOx proposed in [25] seeks to enable efficient spatial reuse in distributed MIMO networks, by proposing a scalable architecture that connects a smaller number of neighbouring APs for clustered Network-MIMO. Each cluster including multiple connected APs that form a giant virtual AP contends with each other via asynchronous CSMA. Within the interference range, no more than one virtual AP is allowed to operate at the same time. Our solution similarly clusters the APs and assumes CSMA between clusters, but employs a fundamentally different approach to address the interference inside a cluster. Instead of connecting the APs to employ Network-MIMO, our solution relies on the AP and clients to coordinately cancel the inter-cell interference.

In this work, we exploit the multiple antennas on clients to combat the inter-cell interference in multiple MU-MIMO cells and to achieve a higher multiplexing gain. Alternatively, MACCO proposed in [18] seeks to use the client antennas to improve the channel orthogonality between the clients, in order to make the capacity of a single MU-MIMO cell better scale with the number of clients. Our work and [18] actually represent two different ways of leveraging the client antennas to achieve either a SNR gain or a multiplexing gain, and therefore need to address completely different technical challenges. More importantly, the applicability of [18] is restricted to a single cell or multiple cells where inter-cell interference is insignificant; in this work we focus on situations where inter-cell interference stands as the bottleneck of the MU-MIMO network capacity and proposes a solution to overcome it.

10. CONCLUSION

In this work, we tackle the inter-cell interference problem in 802.11ac-based MU-MIMO networks by seeking to enable coordinated interference cancellation in a practical way. To achieve this, we propose a solution consisting of two separate optimizations: antenna usage optimization and beamforming weight optimization. With the separation, both optimizations can be integrated into 802.11ac with small modifications and negligible overhead. The experimental evaluation of our protocol, CoaCa, demonstrates its effectiveness to eliminate inter-cell interference and improve the network capacity in realistic indoor wireless environments.

ACKNOWLEDGEMENTS

This work was funded in part by NSF grants MRI 0923479, NetSE 101283, MRI 1126478, CNS 1218700 and CNS 1314822. The authors would like to thank Ashutosh Sabharwal for many insightful suggestions. The authors also thank the anonymous reviewers and the paper shepherd for their constructive comments that have considerably improved the paper.

11. REFERENCES

- [1] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly. Design and Experimental Evaluation of Multi-User Beamforming in Wireless LANs. In *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2010.
- [2] H. V. Balan, R. Rogalin, A. Michaloliakos, and K. Psounis. Achieving High Data Rates in a Distributed MIMO System. In *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2012.
- [3] T. Bansal, B. Chen, K. Srinivasan, and P. Sinha. RobinHood: Sharing the Happiness in a Wireless Jungle. In *Proc. ACM Int. Workshop Mobile Computing Systems and Applications (HotMobile)*, 2014.
- [4] O. Bejarano, E. W. Knightly, and M. Park. IEEE 802.11ac: From Channelization to Multi-User MIMO. *IEEE Communication Magazine*, 2013.
- [5] O. Bejarano, E. Magistretti, O. Gurewitz, and E. W. Knightly. MUTE: Sounding Inhibition for MU-MIMO WLANs. In *Proc. IEEE Int. Conf. Sensing, Communication and Networking (SECON)*, 2014.
- [6] V. R. Cadambe and S. A. Jafar. Interference Alignment and Degrees of Freedom of the K-User Interference Channel. *IEEE Trans. on Information Theory*, 2008.
- [7] B. Chen, K. C. Lin, and H. Wei. Harnessing Receive Diversity in Distributed Multi-User MIMO Networks. In *Proc. ACM SIGCOMM (Poster)*, 2013.
- [8] S. Gollakota, S. D. Perli, and D. Katabi. Interference Alignment and Cancellation. In *Proc. ACM SIGCOMM*, 2009.
- [9] S. Kumar, D. Cifuentes, S. Gollakota, and D. Katabi. Bringing Cross-Layer MIMO to Today's Wireless LANs. In *Proc. ACM SIGCOMM*, 2013.
- [10] K. C. Lin, S. Gollakota, and D. Katabi. Random Access Heterogeneous MIMO networks. In *Proc. ACM SIGCOMM*, 2011.
- [11] M. A. Maddah-Ali, A. S. Motahari, and A. K. Khandani. Communication Over MIMO X Channels: Interference Alignment, Decomposition, and Performance Analysis. *IEEE Trans. on Information Theory*, 2008.
- [12] H. Rahul, S. Kumar, and D. Katabi. JMB: Scaling Wireless Capacity with User Demands. In *Proc. ACM SIGCOMM*, 2012.
- [13] Rice University. WARPLab. <http://warp.rice.edu/trac/wiki/WARPLab>.
- [14] Rice University. Wireless Open Access Research Platform (WARP). <http://warp.rice.edu/trac>.
- [15] W. Shen, Y. Tung, K. Lee, K. C. Lin, S. Gollakota, D. Katabi, and M. Chen. Rate Adaptation for 802.11 Multiuser MIMO Networks. In *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2012.
- [16] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong. Argos: Practical Many-antenna Base Stations. In *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2012.
- [17] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. M. Voelker. SAM: Enabling Practical Spatial Multiple Access in Wireless LAN. In *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2009.
- [18] *Withheld for blind review*. Achieving Better Channel Orthogonality for Improved User Scaling of Multi-user MIMO. In *submission to MobiCom (submission ID: 107)*, 2014.
- [19] S. Venkatesan, A. Lozano, and R. Valenzuela. Network MIMO: Overcoming Inter-cell Interference in Indoor Wireless Systems. In *Proc. IEEE Asilomar Conf. on Signals, Systems and Computers (ACSSC)*, 2007.
- [20] X. Xie and X. Zhang. Scalable User Selection for MU-MIMO Networks. In *Proc. IEEE Conf. Computer Communications (INFOCOM)*, 2014.
- [21] X. Xie, X. Zhang, and K. Sundaresan. Adaptive Feedback Compression for MIMO Networks. In *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2013.
- [22] Q. Yang, X. Li, H. Yao, J. Fang, K. Tan, W. Hu, J. Zhang, and Y. Zhang. BigStation: Enabling Scalable Real-time Signal Processing in Large MU-MIMO Systems. In *Proc. ACM SIGCOMM*, 2013.
- [23] H. Yu, L. Zhong, A. Sabharwal, and D. Kao. Beamforming on Mobile Devices: A First Study. In *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2011.
- [24] Z. Zeng, Y. Gao, K. Tan, and P. R. Kumar. CHAIN: Introducing Minimum Controlled Coordination into Random Access MAC. In *Proc. IEEE Conf. Computer Communications (INFOCOM)*, 2011.
- [25] X. Zhang, K. Sundaresan, M. A. Khojastepour, S. Rangarajan, and K. G. Shin. NEMOx: Scalable Network MIMO for Wireless Networks. In *Proc. ACM Int. Conf. Mobile Computing and Networking (MobiCom)*, 2013.