

Efficient MAC for distributed multiuser MIMO systems

A. Michaloliakos, R. Rogalin, V. Balan, K. Psounis* and G. Caire
University of Southern California, Los Angeles, CA 90089

Abstract—A distributed multiuser MIMO system consists of several access points which are connected to coordinating servers and operate as a large multi-antenna access point. Thanks to joint decoding and precoding, all transmitted signal power is useful, rather than “interference”. The system has the potential to support constant rates as the number of clients increases, thus offering a tremendous bandwidth boost. Despite the high gains and potential applicability to real world setups like enterprise networks, this approach is regarded today mostly as a theoretical solution because of some serious implementation difficulties.

Motivated by our recent success in addressing synchronization issues in a real distributed multiuser MIMO testbed that we have developed, in this work we move one step further and design an efficient MAC scheme for such a system. First, we study and design optimal as well as practical user selection and scheduling schemes. Second, we investigate coding methods to realize the achievable rates at higher layers, and pay particular attention to recent rateless codes advancements. Last, we propose a specific MAC super-frame, and discuss how to make our system backward compatible with 802.11ac.

I. INTRODUCTION

The demand for higher and higher data rates in cellular networks and wireless LANs is relentless. While new standards are developed almost every couple of years with the goal of increasing data rates, e.g. 802.11n/ac and 4G-LTE, wireless bandwidth is never sufficient. At first glance, one reason for the scarcity of wireless bandwidth is its extremely high cost (see, for example, the sums paid from cellular providers in recent bandwidth auctions from governments worldwide). But more fundamentally, the problem is that wireless bandwidth is upper bounded by unsurpassed physical laws.

Both academia and industry have been active in trying to address this issue. Observing trends over the past decades, the evidence suggests that the way forward is increasing density. Specifically, while improvements in network protocols, and modulation and coding schemes have managed significant improvements, those improvements pale in comparison to those made by denser deployment of infrastructure, e.g. more WiFi access points. Nevertheless, this approach comes with its own limitations, as the larger the density the larger the interference.

* This research has been funded by the European Union (European Social Fund - ESF) and Greek national funds through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF) - Research Funding Program Thalis: Investing in knowledge society through the European Social Fund, while K. Psounis has been on leave at the Athens University of Economics and Business.

In theory, the best answer to this problem is multiuser MIMO: thanks to joint decoding and precoding all transmitted signal power is useful and there is no interference. As a result, an access point with enough antennas can simultaneously communicate with an increasing number of clients while keeping the per user rate the same. In practice, however, there are a number of challenges [1]. A major one is that the antennas on the access point have to increase as the number of clients increases, and they also have to be placed far enough apart for this multiplexing gain to be realized. A recent idea addressing this issue is to coordinate different access points to act as a distributed, “mega” access point. While this idea, termed *distributed (or “virtual”) MIMO*, provides some relief to the issue mentioned above, it makes implementation issues even more critical and challenging. For example, a rather challenging issue is how to synchronize remote access points to very high levels of accuracy such that beamforming using antennas which are driven by different clocks is possible. Recent work of ours [2], [3] as well as some parallel work [4], [5] has made good progress on that front and distributed multiuser MIMO testbeds have been implemented, see Section III for more details.

Motivated by this, in this paper we move one step forward and investigate how to design an efficient *Media Access Control* (MAC) scheme for distributed Multi User MIMO (MU-MIMO) systems. We assume that a number of access points (APs) are connected via ethernet wires to a coordinating server and are able to get timely channel state information (CSIT) and transmit synchronously the jointly precoded data which the server generates for them. We also assume that Time Division Duplex (TDD) is used to separate the uplink (UL) and the downlink (DL), and design a MAC scheme which supports MU-MIMO at the DL.

There are many different aspects of MAC that we tackle in this paper. First, we investigate the user selection and scheduling problem. Specifically, at every time slot the coordinating server needs to preselect some users to get CSIT from them, and then to schedule a subset of them based on some performance/optimal criteria. In our investigation we start by stating the optimal scheduling policy and then offer greedy/heuristic algorithms to tackle implementation and fairness problems. Second, we investigate how we can convert the PHY layer gains into rate gains for the higher layers. Specifically, in the absence of the theoretical Gaussian rates, one would have to use a discrete set of modulation/coding

schemes (MCS). However, this set has to be prohibitively large to prevent significant performance losses, which prompts us to focus our attention on rateless codes. Third, we propose a specific MAC super-frame, and we also comment on how we can make our MAC scheme compatible with the new 802.11ac standard which offers optional support for MU-MIMO.

The outline of the paper is as follows. In the next section we go briefly through related work. In Section III we briefly summarize our distributed MU-MIMO testbed and highlight PHY issues related to MAC. Then, Section IV discusses in detail the various scheduling schemes that we consider and investigates a number of practical limitations. In Section V we study using simulations the rate and delay performance of different approaches, which prompts us to further investigate the use of a rateless code (e.g. [6]) in the context of a distributed MU-MIMO system. Then, Section VI presents a MAC super-frame and discusses backward compatibility issues with both legacy clients and with 802.11ac. Last, Section VII concludes the paper.

II. RELATED WORK

A. Precoding schemes

The seminal papers of Foschini [7] and Telatar [8] showed the gains of multiple antennas on the transmitter and the receiver for a point-to-point channel, introducing the notion of *Multiple Input Multiple Output* (MIMO) communication systems. A few years later the results were extended for multiple receiving users in [9] and the solution of its information theoretic capacity region was presented in [10]. The bottom-line of this theory is that both in the single-user and in the multiuser MIMO cases the capacity of a system with M jointly processed transmit antennas and K receive antennas (across all receiving users), with sufficiently rich fading, behaves like $\min\{M, K\} \cdot \log \text{snr} + O(1)$ bit/s/Hz, for large snr, as if $\min\{M, K\}$ “parallel” virtual channels were created in the spatial domain. This capacity multiplication factor is referred to as the *multiplexing gain* of the MIMO channel.

Extensive research has been done on practical precoding schemes able to achieve a large fraction of this capacity with low complexity. Under the assumption of full CSIT, the capacity achieving scheme consists of Dirty-Paper Coding (DPC) [11] combined with optimal power allocation. DPC is essentially an information theoretic tool for realizing perfect interference pre-cancellation with no power penalty whose practical implementation is notoriously difficult [12], [13]. A simpler alternative to such non-linear precoding schemes is represented by *linear* precoding, i.e., beamforming. Among linear beamforming schemes, the simplest and best known is Zero Forcing Beamforming (ZFBF) [14], consisting of a column-normalized version of the right pseudo-inverse of the channel matrix. It has been shown [15] that when the number of users is larger than the number of the transmit antennas, ZFBF may use user selection to approach the system’s sum rate capacity.

B. Software defined radio testbeds

A number of recent system implementations have made forays into the topics of MU-MIMO and distributed, slot aligned OFDM transmissions. MU-MIMO ZFBF as a precoding scheme in a centralized setting has been examined in [16], for a system consisting of a single AP with multiple antennas hosted on the same radio board. A system using a common clock source to drive a large number of radio boards is presented in [17]. This system uses conjugate beamforming [18], a completely decentralized precoding scheme which requires a significantly larger number of antennas in order to provide rate gains comparable to the ones of centralized precoding schemes [17]. The use of interference alignment and cancellation as a precoding technique, which does not require slot synchronization or phase synchronization of the transmitters, has been illustrated in [19]. While this solution achieves some spatial multiplexing, realizing the full spatial multiplexing gain using precoding schemes such as ZFBF requires tight phase synchronization between the jointly precoded transmitters [20], [21].

Slot alignment¹ was used in SourceSync [22] in conjunction with space-time block coding in order to provide a diversity gain in a distributed MIMO downlink system. In Fine-Grained Channel Access [23], a similar technique allows for multiple independent clients to share the frequency band in fine increments, without a need for guard bands, resulting in a flexible OFDMA (OFDM with orthogonal multiple access) uplink implementation. Last, full spatial multiplexing is achieved by tightly synchronizing different access points (APs) towards the creation of a distributed MU-MIMO systems in recent works by us and others [2], [3], [4], [5]. In these papers an AP is elected as the *master* and its phase and timing is broadcasted and tracked as reference for synchronization from the other APs.

C. MAC protocols for MU-MIMO and rateless systems

There is a large body of work on communication and information theory venues focusing on the user selection problem introduced by precoding schemes when the number of antennas on the transmitters exceeds that of the receivers (see [24], [25], [15], [26] and Section IV). In these mostly theoretical papers, practical issues concerning modulation and coding schemes, delays that users might incur and overall overhead of the signaling required are neglected.

The few papers that treat the MAC layer design from a networking perspective focus on MAC protocol issues while opting for practical scheduling decisions agnostic to optimality results. For example, in [27] and [28] the authors focus on extending the 802.11 distributed coordination function (DCF) to support MU-MIMO operations by including the extra signaling required for MU-MIMO in extended versions of the RTS and CTS messages. Last, a special mention should be

¹Slot alignment refers to the situation where transmissions from different nodes align within the cyclic prefix of OFDM, which is required to avoid inter-block interference.

made to the recently introduced 802.11ac [29] protocol that will eventually replace 802.11n. 802.11ac includes optional MU-MIMO support by modifying MAC layer fields to accommodate the extra signaling required. In Section VI we investigate how to make our proposal backward compatible with the standard.

III. A DISTRIBUTED MU-MIMO SYSTEM

In this section we briefly describe the distributed MU-MIMO testbed over which we have implemented several precoding schemes. We discuss the two major obstacles of realizing such a distributed system: the synchronization of the remote APs, and the requirements of the wired backhaul connecting the APs with a coordinating server (see Figure 2). We briefly describe our synchronization method which is part of our prior work [2], [3]. We proceed by identifying the wired networking requirements of such a system.

A. Synchronization

Most precoding schemes designed for MU-MIMO transmission assume that multiple transmit antennas are hosted on the same AP and share the baseband circuitry and the clocking circuitry which produces the passband signals. For schemes exploiting full CSIT it is essential that the signal amplitudes and the relative phase and timing offsets of the signals received from different antennas remain unchanged between channel estimation and the actual transmission. Since full CSIT schemes achieve the full multiplexing gains of the distributed MIMO system, and in high SNR the advantage in spectral efficiency can be very significant, it is imperative to achieve a high level of synchronization between remote APs.

We have implemented a synchronization method in the FPGA of the WARP [30] software radio platform that achieves high accuracy and enables full CSIT precoding in the context of distributed multiuser MIMO. In a nutshell, one AP broadcasts a reference signal onto which the rest of the APs lock. Thus, synchronization is achieved over the air. For more information on the synchronization see [2], [3].

B. Wired Network Requirements

After achieving synchronization, a coordinating server may compute the precoded data using the CSIT information it collects from the APs and then distribute it to the APs, or, alternatively, the coordinating server may distribute the CSIT information and the data for the users to all APs so that the computation can take place locally at every AP. Figure 2 shows a setup with four users and four APs, connected to a coordinating server through individual Ethernet links.

The most natural approach is to make the server responsible for the joint encoding of the transmitted signals and for passing the resulting waveforms, in the form of frequency domain soft symbols, to the transmitter radios. This approach requires sending full waveforms over the ethernet connection for every antenna on each AP. Using the specifications of the WARP board, which are also typical for consumer-grade WIFI, at the APs we have dual 16-bit DACs (Digital-to-Analog Converters)

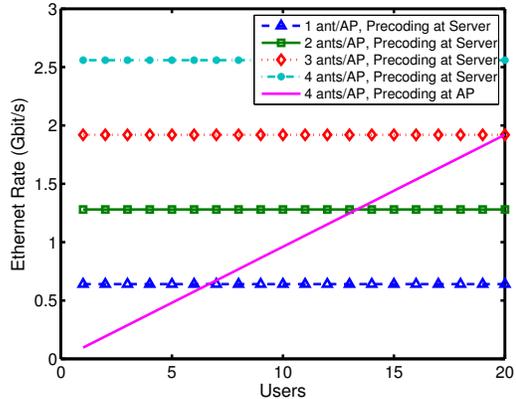


Fig. 1. Comparison of precoding at the APs and precoding at the server for various number of antennas per AP and users.

and we transmit at 20 MHz bandwidth. This suggests using 32 bits per sample (16 bits for the real and 16 bits for the imaginary part) and gives a requirement of 640 Mbps per AP antenna. It is obvious that the more antennas an AP has, the larger the wired bandwidth requirement. On the positive side, the bandwidth requirement is independent from the number of users concurrently served by the system.

A second approach is to push the precoding computation to the APs, by distributing the CSIT and the user data to them from the coordinating server, instead of the full waveforms. Assuming that CPU calculations at the APs can be done fast enough, we focus again on the wired bandwidth requirement. For typical WIFI specifications (dual 14-bit Analog-to-Digital converters, 48 data OFDM subcarriers, and 20 MHz bandwidth) and CSIT updates taking place every 100 ms, the wired network requirements depend mostly on the number of users served and the modulation and coding scheme used. In a worst case scenario where every user is served at an 8 bits/symbol rate (256-QAM) the required backhaul per AP is 96 Mbps times the number of users served. Notice that the CSIT information that has to be exchanged between the APs for 64 OFDM subcarriers using 28 bits per sample resolution every 100ms adds only a 17.920 Kbps additional bandwidth requirement per AP antenna. Thus, for all practical purposes, in this case the backhaul requirement does not depend on the number of antennas per AP.

The wired bandwidth requirement is plotted in Figure 1 for the cases of 1 to 4 antennas per AP and up to 20 users being concurrently served. As expected, the precoding at the APs is almost insensitive to the number of antennas per AP (only the case of 4 antennas per AP is plotted since the other lines almost coincide with it). Nevertheless, as the number of users increases, this approach is quickly dominated by the precoding at the server scenario.

IV. OPTIMAL AND PRACTICAL SCHEDULING

The design of a MAC layer that can fully exploit the characteristics of the underlying distributed MU-MIMO physical

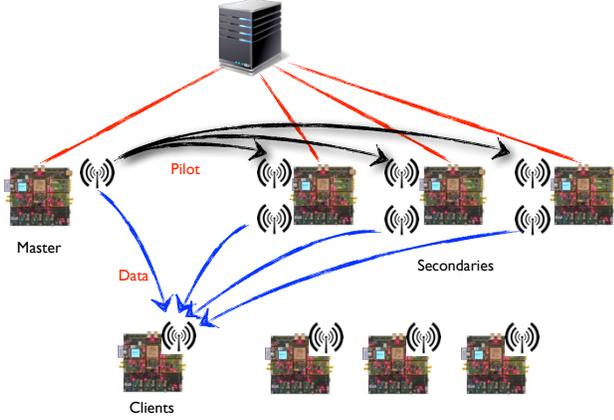


Fig. 2. Testbed diagram. The coordinating server is connected to four transmitters, the main transmitter on the left and the three secondary transmitters on the right. Four receivers act as clients.

layer is a challenging task.

First, let us consider the issue of allocating the resources (air time, frequency) between the UL (communication from the users to the access point) and the DL (communication from the access point to the users). We choose to use time division duplex (TDD) for a number of reasons. First, TDD is ideally suited for the transport of asymmetric traffic, as is typical in an enterprise WiFi environment, since it allows to allocate different bandwidth to each direction. Second, as already mentioned, MU-MIMO systems require CSIT for the downlink, and, when using TDD we can exploit channel reciprocity to acquire CSIT for the downlink through pilots sent from the users in the uplink channel.

A. Optimal Scheduling

We will start by discussing the issue of optimally scheduling users in the DL of the distributed MU-MIMO system model we described. Notice that, after time and phase synchronization between the decentralized APs of our system has been achieved, we can treat this setup as if all antennas are located on the same AP. Without receiver cooperation, as is the case in our scenario, successful utilization of the channel requires careful scheduling and precoding of the independent signals at the transmitter side.

As already mentioned, DPC is the optimal, capacity achieving strategy for the underlying MIMO Gaussian broadcast channel. Thus, it is the natural term of comparison for any practical low-complexity scheme. Further, ZFBF is a practical, and easily deployable choice for a broadcast channel like ours [24], [25], [15]. Thus, we will assume the use of ZFBF. The system model consists of M antennas in the transmitters forming a distributed MIMO system as described in the previous sections and U single-antenna users waiting to be serviced. When $U > M$ we can only serve $K \leq M$ out of the U users, otherwise the nulling of multiuser interference is not possible. User selection is a combinatorial problem, consisting of selecting, for each subcarrier n , a subset $S(n)$ of users (where $|S(n)| \leq M$) by maximizing a target utility

function under the transmit power constraint. Therefore, there exists a significant coupling between the decisions made at the MAC layer for the scheduling of the users and the precoding at the PHY layer. Assuming that each information stream is independently encoded with ideal capacity-achieving codes, the optimization problem to be solved at each scheduling slot is:

$$\begin{aligned}
 \max \quad & R(S) = \sum_{n=0}^{N-1} \sum_{k \in S(n)} W_k \log_2(1 + \text{SINR}_k(n)), \\
 \text{w. r. t.} \quad & \{S(n) \subseteq \{1, \dots, K\} : |S(n)| \leq M\}, \\
 & \{P_k(n)\} \\
 \text{s. t.} \quad & \frac{1}{N} \sum_{n=1}^N \sum_{k \in S(n)} P_k(n) \leq P_{\text{sum}} \quad (1)
 \end{aligned}$$

where $\text{SINR}_k(n)$ is the Signal to Interference plus Noise Ratio (SINR) for user $k \in S(n)$ on subcarrier n obtained by ZFBF, $\{W_k\}$ are the scheduling weights, $P_k(n)$ is the power allocated to user k on subcarrier n and P_{sum} is the total power constraint (transmitted power spectral density over the N subcarriers). Further, $\text{SINR}_k(n) = \Lambda_k^2(n)P_k(n)$ where $\Lambda_k(n)$ is the k -th ZFBF coefficient of the channel sub-matrix of the overall $U \times M$ channel matrix $\mathbf{H}(n)$ corresponding to the users in $S(n)$ [31]. Conceptually, this optimization problem can be solved by exhaustively searching over all feasible subsets of users S . In practice, greedy algorithms that add one *stream* at a time, where a stream is defined by a pair (k, n) of user and subcarrier index, have proven to provide excellent results at moderate complexity [26], [25]. We extend the greedy user selection algorithm of [25] for multiple subcarriers and we have the following algorithm where as R we denote the achievable sum rate using Gaussian rates for the scheduled user/subcarrier pairs.

Algorithm 1 Greedy ZF with Waterfilling and Gaussian rates (ZF-G)

Initialization: $S = \emptyset, R(S) = 0,$
 $S(n) = \emptyset \forall n$

while $|S| < U \cdot N$ **do**

$$\{k^*, n^*\} = \arg \max_{\{k, n\} \notin S, |S(n)| < M} R(S \cup \{k, n\})$$

if $R(S \cup \{k^*, n^*\}) \leq R(S)$ **then**

break;

else

$$S \leftarrow S \cup \{k^*, n^*\}$$

$$S(n^*) \leftarrow S(n^*) \cup \{k^*\}$$

end if

end while

To compute the maximum achievable sum rate for a given scheduled set of user/subcarrier pairs we use the waterfilling equation:

$$\sum_{n=0}^{N-1} \sum_{k \in S(n)} \left[\frac{W_k}{\mu} - \frac{1}{\Lambda_k^2(n)} \right]^+ = N \cdot P_{\text{sum}} \quad (2)$$

that can be derived from the convex optimization problem (1).

B. Practical Considerations

Though DPC achieves capacity and Algorithm 1 above (ZF-G) provides a convenient sub-optimal algorithm, there are practical considerations beyond simple ZF-G that are relevant to real systems. Firstly, the use of coding rates equal to the corresponding Gaussian channel capacity $\log(1 + \text{SINR})$ is overly idealized; by mapping the SINRs into a discrete set of modulation and coding schemes (MCSs) we can model a more realistic scenario. While we acknowledge that choosing the best among several discrete coding and modulation options (known as rate adaptation) is non-trivial, for the sake of simplicity we assume that we can choose the best scheme based on the received SINR, optimized through the Gaussian waterfilling power allocation. Table I provides one such mapping that corresponds to the 9 mandatory MCSs of 802.11ac [29], keeping in mind that mappings vary by vendor or may be dynamically chosen in practical scenarios. Using this scheme within the greedy ZF user selection, we examine the resulting method, dubbed greedy zero forcing with adaptive coding and modulation (ZF-ACM).

802.11ac MCS Index	Modulation	Code Rate	SNR Range
0	BPSK	1/2	$\geq 2\text{dB}$
1	QPSK	1/2	$\geq 5\text{dB}$
2	QPSK	3/4	$\geq 8\text{dB}$
3	16-QAM	1/2	$\geq 12\text{dB}$
4	16-QAM	3/4	$\geq 15\text{dB}$
5	64-QAM	2/3	$\geq 18\text{dB}$
6	64-QAM	3/4	$\geq 21\text{dB}$
7	64-QAM	5/6	$\geq 24\text{dB}$
8	256-QAM	3/4	$\geq 27\text{dB}$

TABLE I
MODULATION/CODING PAIRS FROM IEEE 802.11ac AND THE CORRESPONDING SNRS AT WHICH THEY CAN BE SELECTED

Since the power allocation step in ZF-G is time consuming, we may wish to make a simpler allocation decision. For example, we could schedule users one at a time, but then divide the total power constraint among the selected users. We call this scheme ZF-P, or greedy zero forcing power constrained.

Finally, because higher layer protocols such as TCP are subject to timeouts when users go unserved for long periods, we would like to ensure that users are served regularly to prevent timeouts. To accomplish this, we use round robin scheduling to choose the initial user in algorithm 1 and then schedule him on his best subcarrier and continue greedily for the remaining streams. This ensures that all users are served at least once every K slots. We designate this scheme ZF-RR, and expect that its long term performance should be close to that of ZF-G.

V. SIMULATIONS

In this section we will compare the schemes presented earlier with respect to their achievable sum rate and the delay users experience. We will evaluate the trade offs that a real system implementation may experience.

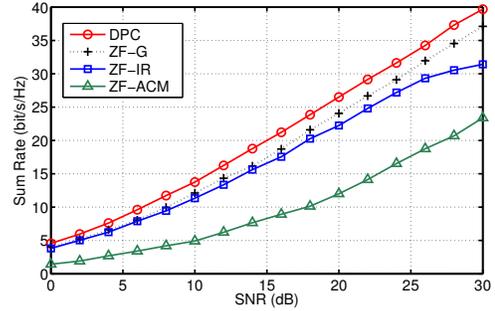


Fig. 3. Comparison of the sum-rates for various greedy zero-forcing schemes.

A. Sum Rate

We compute, through simulation, the achievable sum rate of the greedy algorithms presented earlier and compare it to the optimal sum rate, achieved by DPC. We have observed experimentally that in a typical indoor or short-range outdoor scenario the channel frequency response is strongly correlated, so that even though N may be large (e.g., $N = 64$), the system bandwidth spans at most 4 channel coherence bandwidths. This means that the same stream allocation can be replicated over 4 blocks of adjacent subcarriers, at minimal degradation in performance. For this reason, in these simulations we considered $N = 4$ independent subcarriers (equivalent to 4 channel coherence bands).

In Figure 3 the sum rates for the two schemes, greedy ZFBF with ideal rates (ZF-G) and the adaptive coding and modulation (ACM) scenario (ZF-ACM) described above are evaluated for multiple SNRs in the case of $U = 10$ clients and $M = 4$ total access points antennas. For purposes of reference, the optimal, capacity-achieving Dirty Paper Coding (DPC) [11] precoding technique is shown in the same plot. Note that ZF-P and ZF-RR sum rates overlap the ZF-G achievable one and thus are not depicted.

A few comments based on these results are in order. ZF-G is near optimal for the low SNR regime, and has a small, constant gap of 3 bits/s/Hz for high SNR for this level of multi-user diversity. For the medium to high SNR regime, the ZF-P achieves the same throughput with close to zero losses from the ZF-G with waterfilling. The ZF-RR strategy also provides almost full ZF-G multiuser diversity gains for all SNR. This is expected because only the first user in the greedy selection is indicated from the round-robin strategy, and thereafter waterfilling ZF-G is employed to schedule the remaining users. On the other hand, the ZF-ACM has a significant gap from the greedy ZF strategies with Gaussian rates, especially for high SNR values.

B. Realizing the Achievable Rates

The huge gap between ZF-ACM and the ZF versions using Gaussian rates motivates us to investigate how to efficiently translate the SNR gains at the PHY layer into rate gains for the higher layers. Note that the traditional ACM approach is particularly problematic in the case of MU-MIMO because

such a system serves multiple users in the same time slot, and an even larger set of rates and codes would have to be supported for efficiently using capacity.

Given the above, we turn our attention to more flexible ways of allocating the rates in the MU-MIMO scenario. Specifically, we use *rateless codes* (e.g., Raptor codes [32] and the recently proposed Spinal codes [6]) at the physical layer, in a so-called *Incremental Redundancy* (IR) configuration (see [33], [34], [35]) to decrease the signaling and retransmission overhead. In an ideal rateless coding adaptation scenario, we would achieve the *coded modulation capacity* of a fixed large QAM constellation. In Figure 3 the performance of greedy zero-forcing with such an ideal rateless code (ZF-IR) is also depicted for an ideal family of random rateless codes based on a 256-QAM constellation. It is immediately obvious that the gains of using this IR configuration are tremendous in comparison with classic ACM.

C. Delay Considerations

To evaluate the average delay of the greedy scheduling algorithms we implemented a queuing system with random arrivals for every user and updated the queues, based on the scheduling and power allocation of our algorithms, after every DL scheduling slot. For queue stability the queue backlogs are used as scheduling weights $\{W_k\}$. From Figure 3 we can already compute the symmetric arrival rate point, i.e., the point of the capacity region where all queues have the same arrival rate, for all scheduling schemes. We run our simulator for 10^5 scheduling slots for a range of symmetric arrival rates close to the achievable for every scheme and compute the average delay. For the sake of computational complexity, we choose to simulate the scenario of $U = 10$ users and $M = 4$ antennas with SNR equal to 10 dB and a single subcarrier, which represents a highly correlated channel in the MIMO-OFDM scenario.

As expected, for arrival rates that are significantly lower from the achievable, all schemes manage to keep delays at very low levels. The ZF-ACM is the first to give unbounded delays for an arrival rate of 0.5 bits/s. ZF-RR gives an achievable arrival rate of 1 bit/s which is very close to the 1.1 bits/s the ZF-G and ZF-P schemes can achieve for this SNR. The average delay for this symmetric arrival system is worse for the case of the ZF-RR compared to ZF-G. However, as mentioned earlier, ZF-RR greatly improves the per user delay in a non-symmetric arrival rate scenario. For example, in a situation where 9 users have an average arrival of 0.9 bits/s/time slot and the 10th has only 0.01, Figure 5 shows that the ZF-RR scheduler manages to significantly decrease the delay in comparison to the vanilla ZF-G for this user, without a noticeable increase to the delay of the other users.

VI. PROTOCOL DESIGN

The crucial design constraint in distributed MU-MIMO is to provide the coordinating server with timely estimates of the CSIT in order to perform DL scheduling and MU-MIMO precoding. This is accomplished by pre-selecting a large set

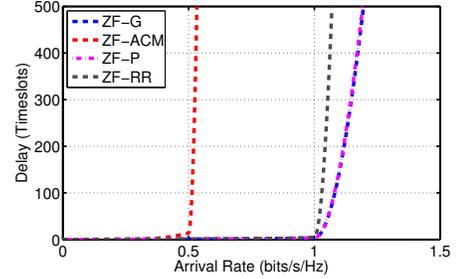


Fig. 4. Average queue delay for $U = 10$ users, $M = 4$ antennas, and $N = 4$ subcarriers and SNR = 10 dB.

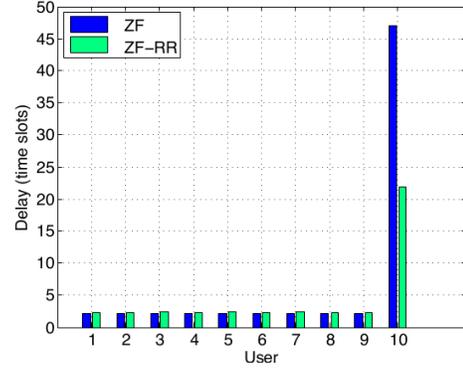


Fig. 5. Delay comparison of ZF-G vs. ZF-RR.

of clients (possibly in a round-robin fashion), requesting their UL pilot signals, and providing the UL estimated channels to the coordinating server. Thanks to TDD reciprocity, the coordinating server is able to infer DL CSIT, decide the users that will be scheduled and compute the DL precoder. It is mandatory to have the corresponding DL data transmission right after UL CSIT estimation, such that the channel is not outdated and MU-MIMO precoding is effective in spatially separating the client data streams.

With this in mind, we propose the super-frame structure shown in Fig. 6 which comprises of the following slots: SYNC slot for the APs synchronization protocol, DL-RFP slot to request pilots from the preselected clients, UL-P slot used by the clients to send UL pilots, DL-DATA slot for MU-MIMO data transmission, and UL-DATA slot during which clients may transmit ACK and data to the APs. A description for each of these components is provided in the following.

SYNC slot: The SYNC slot may include several orthogonal pilot signals used to synchronize the clocks of remote APs. As already mentioned, in our testbed the master AP broadcasts a pilot on which all other APs will lock on. This scheme requires all secondary APs to be within range of the master which may not be the case in an enterprise WiFi scenario. For this reason, one may envision a synchronization method where each AP exchanges pilots with all its neighboring APs and APs lock to a common clock reference that is computed from one AP or the coordinating server using information from all these pilots. While it is beyond the scope of this paper to

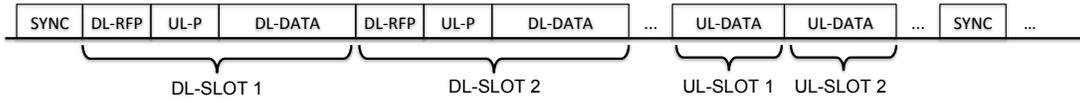


Fig. 6. The super frame structure of the distributed MU-MIMO MAC layer.

further discuss particular synchronization methodologies, it is important to note that the SYNC slot can accommodate any such methodology. Clearly, the larger the number of exchanged pilots, the larger the SYNC slot.

DL-RFP slot: In this slot, all APs simultaneously broadcast a control signal to inform the clients that will participate in the next UL slot and request their UL pilots. Since after the SYNC slot the APs are time-synchronous (within a cyclic prefix), they can transmit simultaneously and each client will see the data through a “virtual” multipath channel originated by the superposition of all the APs signals. Here, reliability is more important than spectral efficiency, thus we use Space-Time Block Codes (STBC) [36] in this slot.

UL-P slot: The pre-selected clients send their short UL pilot signals in this slot in an orthogonal manner (either in time or in frequency). From this set of clients, based on the inferred DL channels the MAC scheduler will select the users for the ensuing DL slot.

DL-DATA slot: After an UL-P slot, a DL-DATA slot takes place. The DL frame has the following fields. First, a packet header is broadcasted to all clients, including the scheduled ones, using STBC for resilience, to inform the clients that will be receiving data during this DL slot (this is needed since the scheduling decision involves a further selection of the pre-selected clients which have sent their UL pilots). This information is encoded into an allocation map, similar to the one found in the LTE standard [37], which assigns carriers to small groups of different clients and specifies the signal constellations used on each subcarrier. From this point on, the transmitters switch to MU-MIMO mode and send precoded symbols. Each data stream contains its individual pilots for coherent reception which must go through the concatenation of the MU-MIMO precoder and the DL channel.

UL-DATA slot: During this slot the clients will transmit acknowledgments (ACK) based on a successful cyclic redundancy checksum of their decoder’s output. The SMACK scheme [38] is used for efficient bandwidth usage for this task. In addition, UL-DATA slots are also used for legacy MAC protocols like 802.11n/ac. We choose to have these slots after all DL activity to ensure that UL data activity is not introducing time delays between user preselection decisions, UL pilot transmissions, and DL data transmissions based on UL channel estimation.

A. Backward compatibility with 802.11ac

It is clear from the discussion above that the APs that participate in the distributed MU-MIMO system cannot be legacy and therefore no restrictions need to be imposed on the design of the SYNC slot where only they participate. The

RTS/CTS procedure of 802.11 protocols can be used to silence other co-existing legacy APs and clients for this period.

More generally, there are two approaches to backward compatibility. One is to transmit in an 802.11ac compatible manner. Since the standard is not finalized yet we leave this as future work. Instead, we comment on the second approach which uses 802.11ac’s fields to inform legacy clients that a non-legacy packet is transmitted. One way of doing this is to use the VTH-SIG-A field in 802.11ac’s header (see Figure 7 as well as [29]) which is used to determine whether a data packet belongs to 802.11a/b/g, 802.11n, or 802.11ac. Specifically, this field composes of two OFDM symbols whose modulation determines the type of the packet, for example, if the first symbol is QPSK modulated, the WLAN receiver determines the received packet conforms to 802.11n, and if the two OFDM symbols are BPSK and then QPSK modulated, the receiver determines 802.11ac is used. Thus, all that is required is to extend the above signaling code by making, for example, two QPSK symbols imply our new frame is used. Our DL-DATA slot in this case would have to always start with the L-STF and L-LTF fields, followed by the L-SIG field which specifies the length of the packet, followed by the VTH-SIG-A field, and from then on legacy clients would know to stay silent for the duration of the frame.

VII. DISCUSSION AND FUTURE WORK

We are currently working on improving various aspects of our testbed related to the implementation of efficiently collecting CSIT and synchronizing many APs in a scalable manner. Also, motivated by the significant degradation between the rates achieved via DPC and those achieved in practice by an adaptive modulation and coding scheme (see Figure 3), we are currently studying the implementation of viable rateless coding schemes at the MAC/PHY layer.

There are a few main challenges related to the implementation of rateless codes. First, rateless coding does not schedule a service rate a priori, rather, the coding rate is a function of the sequence of future channel realizations. Therefore, the discussed schedulers which compute the scheduling weights on the basis of the allocated service rate required modification. Second, the most natural way of implementing rateless codes at the PHY layer consists of mapping the output of a Raptor encoder [32] onto modulation alphabets, chosen according to a quantized water filling scheme. A recent attractive alternative consists of using spinal codes [6]. A thorough comparison of these two options in the context of our distributed multiuser MIMO MAC/PHY is the subject of ongoing investigation.

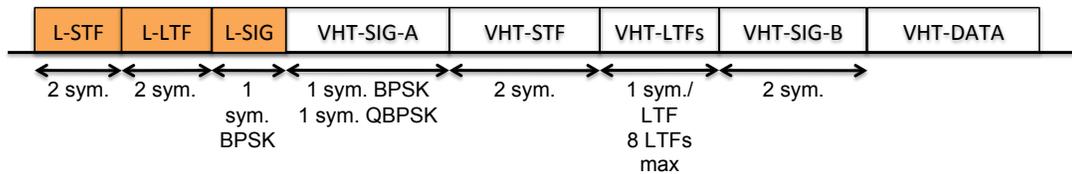


Fig. 7. 802.11ac frame.

REFERENCES

- [1] S. Häne, D. Perels, D. S. Baum, M. Borgmann, A. Burg, N. Felber, W. Fichtner, and H. Bölcskei, "Implementation aspects of a real-time multi-terminal MIMO-OFDM testbed," in *RAWCon*, Sep. 2004.
- [2] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, "Achieving high data rates in a distributed mimo system," in *Proceedings of the 18th annual international conference on Mobile computing and networking*, ser. Mobicom '12. New York, NY, USA: ACM, 2012, pp. 41–52. [Online]. Available: <http://doi.acm.org/10.1145/2348543.2348552>
- [3] —, "Airsync: enabling distributed mimo with full spatial multiplexing," *IEEE Transactions on Networking (to appear)*.
- [4] H. S. Rahul, S. Kumar, and D. Katabi, "JMB: scaling wireless capacity with user demands," in *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, ser. SIGCOMM '12. New York, NY, USA: ACM, 2012, pp. 235–246.
- [5] F. Quitin, M. Mahboob Ur Rahman, R. Mudumbai, and U. Madhoo, "Distributed beamforming with software-defined radios: frequency synchronization and digital feedback," in *Global Telecommunications Conference, 2012. GLOBECOM '12. IEEE*, 2012.
- [6] J. Perry, H. Balakrishnan, and D. Shah, "Rateless spinal codes," in *ACM HotNets*, Cambridge, Massachusetts, 2011. [Online]. Available: <http://doi.acm.org/10.1145/2070562.2070568>
- [7] G. Foschini and M. Gans, "On limits of wireless communications in a fading environment when using multiple antennas," *Wireless Personal Communications*, vol. 6, pp. 311–335, 1998. [Online]. Available: <http://dx.doi.org/10.1023/A:1008889222784>
- [8] E. Telatar, "Capacity of multi-antenna gaussian channels," *European Transactions on Telecommunications*, vol. 10, no. 6, pp. 585–595, 1999. [Online]. Available: <http://dx.doi.org/10.1002/ett.4460100604>
- [9] G. Caire and S. Shamai, "On the achievable throughput of a multiantenna gaussian broadcast channel," *IEEE Trans. Inf. Theory*, vol. 49, no. 7, pp. 1691 – 1706, Jul. 2003.
- [10] S. Vishwanath, N. Jindal, and A. Goldsmith, "Duality, achievable rates, and sum-rate capacity of gaussian mimo broadcast channels," *IEEE Trans. Inf. Theory*, vol. 49, pp. 2658–2668, 2003.
- [11] M. Costa, "Writing on dirty paper (corresp.)," *IEEE Trans. Inf. Theory*, vol. 29, no. 3, pp. 439–441, May 1983.
- [12] U. Erez and S. ten Brink, "A close-to-capacity dirty paper coding scheme," *Information Theory, IEEE Transactions on*, vol. 51, no. 10, pp. 3417 –3432, oct. 2005.
- [13] A. Bennatan, D. Burshtein, G. Caire, and S. Shamai, "Superposition coding for side-information channels," *Information Theory, IEEE Transactions on*, vol. 52, no. 5, pp. 1872 – 1889, may 2006.
- [14] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. New York, NY: Cambridge University Press, 2005.
- [15] T. Yoo and A. Goldsmith, "On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 3, pp. 528 – 541, Mar. 2006.
- [16] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly, "Design and experimental evaluation of multi-user beamforming in wireless LANs," in *ACM MobiCom*, Chicago, IL, 2010. [Online]. Available: <http://doi.acm.org/10.1145/1859995.1860019>
- [17] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong, "Argos: practical many-antenna base stations," in *Proceedings of the 18th annual international conference on Mobile computing and networking*, ser. Mobicom '12. New York, NY, USA: ACM, 2012, pp. 53–64.
- [18] T. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," *Wireless Communications, IEEE Transactions on*, vol. 9, no. 11, pp. 3590 –3600, november 2010.
- [19] S. Gollakota, S. D. Perli, and D. Katabi, "Interference alignment and cancellation," in *ACM SIGCOMM*, Barcelona, Spain, 2009. [Online]. Available: <http://doi.acm.org/10.1145/1592568.1592588>
- [20] A. Lapidoth, S. Shamai, and M. A. Wigger, "On the capacity of fading MIMO broadcast channels with imperfect transmitter side-information," in *43rd Annual Allerton Conference*, Monticello, IL, 2005.
- [21] C. S. Vaze and M. K. Varanasi, "The degrees of freedom regions of MIMO broadcast, interference, and cognitive radio channels with no CSIT," *CoRR*, vol. abs/0909.5424, 2009.
- [22] H. Rahul, H. Hassanieh, and D. Katabi, "SourceSync: a distributed wireless architecture for exploiting sender diversity," in *ACM SIGCOMM*, New Delhi, India, 2010.
- [23] K. Tan, J. Fang, Y. Zhang, S. Chen, L. Shi, J. Zhang, and Y. Zhang, "Fine-grained channel access in wireless LAN," in *ACM SIGCOMM*, New Delhi, India, 2010.
- [24] H. C. Huang, S. Venkatesan, and H. Viswanathan, "Downlink capacity evaluation of cellular networks with known interference cancellation," in *DIMACS Workshop on Sig. Proc. for Wireless Commun.*, Oct 2002.
- [25] G. Dimic and N. Sidiropoulos, "On downlink beamforming with greedy user selection: performance analysis and a simple new algorithm," *IEEE Trans. Signal Process.*, vol. 53, no. 10, pp. 3857 – 3868, Oct. 2005.
- [26] M. Kobayashi and G. Caire, "Joint beamforming and scheduling for a multi-antenna downlink with imperfect transmitter channel knowledge," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 7, pp. 1468 –1477, Sep. 2007.
- [27] L. Cai, H. Shan, W. Zhuang, X. Shen, J. Mark, and Z. Wang, "A distributed multi-user mimo mac protocol for wireless local area networks," in *Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE*, 30 2008-dec. 4 2008, pp. 1 –5.
- [28] J. Mirkovic, J. Zhao, and D. Denteneer, "A mac protocol with multi-user mimo support for ad-hoc wlangs," in *Personal, Indoor and Mobile Radio Communications, 2007. PIMRC 2007. IEEE 18th International Symposium on*, sept. 2007, pp. 1 –5.
- [29] "IEEE Standard for Information technology–Telecommunications and information exchange between systems Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," *IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007)*, pp. 1 –2793, 2012.
- [30] Rice University, "Rice university warp project." [Online]. Available: <http://warp.rice.edu>
- [31] A. Wiesel, Y. Eldar, and S. Shamai, "Zero-forcing precoding and generalized inverses," *IEEE Trans. Signal Process.*, vol. 56, no. 9, pp. 4409 –4418, 2008.
- [32] A. Shokrollahi, "Raptor codes," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2551 –2567, Jun. 2006.
- [33] G. Caire and D. Tuninetti, "The throughput of hybrid-arq protocols for the gaussian collision channel," *IEEE Trans. Inf. Theory*, vol. 47, no. 5, pp. 1971 –1988, Jul. 2001.
- [34] S. Sesia, G. Caire, and G. Vivier, "Incremental redundancy hybrid ARQ schemes based on low-density parity-check codes," *IEEE Trans. Commun.*, vol. 52, no. 8, pp. 1311 – 1321, Aug. 2004.
- [35] C. Lott, O. Milenkovic, and E. Soljanin, "Hybrid ARQ: Theory, state of the art and future directions," in *IEEE ITW on Information Theory for Wireless Networks*, Solstrand, Norway, Jul. 2007.
- [36] S. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 8, pp. 1451–1458, 1998.
- [37] S. Sesia, I. Toufik, and M. Baker, Eds., *LTE – The UMTS Long Term Evolution*. John Wiley & Sons, Ltd, 2009.
- [38] A. Dutta, D. Saha, D. Grunwald, and D. Sicker, "SMACK: a Smart ACKnowledgment scheme for broadcast messages in wireless networks," in *SIGCOMM*, 2009.