

Efficient MAC for distributed multiuser MIMO systems

V. Balan, A. Michaloliakos, R. Rogalin, 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 central servers and operate as a large distributed multi-antenna access point. Thanks to joint decoding and precoding, all transmitted signal power is useful, rather than “interference” in contrast to conventional random access. This has the potential to support constant rates as the number of clients increases, thus offering a tremendous bandwidth boost. This approach is particularly suited to the case of an enterprise network (e.g., a WLAN covering a conference center, an airport hall or a hotel), or to the case of home networks connected to the ISP infrastructure through the same DSL bundle. However, distributed multiuser MIMO 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 study how to design an efficient MAC scheme for such a system. Specifically, for the more challenging case of the downlink, we start by stating the optimal scheduling policy and then offer heuristic algorithms to tackle various practical problems. First, after reviewing the theoretically optimal scheme, we study a greedy zero-forcing policy whose performance is shown to be very close to optimal. Second, rather than assuming the full spectrum of rates is readily available, we use a discrete set of coding/modulation pairs and offer a practical, “quantized” version of our MAC. Third, we design a flat-spectrum version of our MAC algorithm in the spirit of existing practical schemes like 802.11n, such that legacy clients may be used in the context of our distributed MIMO system with minimal changes. Last, to guarantee some notion of fairness and avoid timeouts in higher layer protocols like TCP, we introduce a simple round robin step in the MAC design.

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 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 (that is, more WiFi access points, more cellular towers per square km and recently very small user-deployed femtocells). Nevertheless, this approach comes with its own limitations, as the larger the density the larger the effects of interference. In theory, the best answer to this problem is multiuser MIMO, since under a number of assumptions, 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, the story is very different. First, the number of 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. Then, the wireless channel between the clients and the access point needs to be rich enough to provide enough spatial diversity. Last, the implementation considerations of such systems appear overwhelming, as reported by numerous academic publications and industry white papers [1], [2], [3].

Very recently, a few researchers have investigated the possibility of coordinating different access points to act as a “mega, distributed” access point. This idea, termed *distributed (or “virtual”) MIMO*, provides some relief to the first two issues mentioned above, but it makes implementation issues even more critical and challenging. Since the access points have autonomous clocks and RF oscillators, and are connected through an existing data network (e.g., Ethernet, or a DSL bundle), there is no obvious way to transmit from them in a *time and phase* synchronous manner (a necessary requirement in order to achieve the promised multiuser MIMO gains). Nevertheless, we have recently experimented with a few synchronization algorithms in a software radio testbed of ours, and we have managed to synchronize in time and phase access points that sit in geographically different places and have separate clocks. Motivated by this, in this paper we move one step forward and investigate how to design efficient MAC schemes for distributed MIMO systems.

The design of a MAC layer that can fully exploit the characteristics of the underlying distributed multiuser MIMO physical layer is a challenging task. One may envision to use Time Division Duplex (TDD) to separate the uplink (UL) and the downlink (DL). In this paper we focus on the more challenging task of scheduling and resource allocation in the DL. We envision that the central server which is connected to all access points and is responsible to implement the joint encoding process for all of them, also keeps track of the packet

queues (data waiting to be transmitted on the DL) and other QoS information (e.g. priorities, fairness, etc.) for each user. Then, the central server would select a subset of users to beamform to at each DL time slot.

We investigate various approaches to the scheduling problem. We start by stating the optimal scheduling policy and then offer heuristic algorithms to tackle various practical problems. First, after reviewing the theoretically optimal scheme, we study a more practical greedy zero-forcing policy. Second, rather than assuming Gaussian coding with full rate flexibility, we use a discrete set of modulation/coding schemes (MCS) and offer a practical, “quantized” version of our MAC. Third, we design a flat-spectrum version of our MAC algorithm in the spirit of existing practical schemes like 802.11n, such that legacy clients may be used in the context of our distributed MIMO system with minimal changes. Last, to guarantee some notion of fairness and avoid timeouts in higher layer protocols like TCP, we introduce a simple round robin step in the MAC design.

The outline of the paper is as follows. In the next section we go briefly through related work. In Section II we summarize our prior work on the distributed MIMO testbed implementation and address the synchronization issues. Then, Section III introduces in detail the various scheduling schemes that we consider. Section IV compares the performance of the proposed scheduling schemes in the context of a distributed MIMO system. Section V discusses MAC protocol design issues. Section VI concludes the paper.

II. A DISTRIBUTED MIMO SYSTEM

In this section we summarize our work on building a distributed MIMO system based on a software-defined radio platform. We start by showing why time and phase synchronization is required in order to materialize the promised gains. We then describe how we have addressed this issue in our testbed. In a nutshell, we lock the phase of all access points using a common reference broadcasted over the air by an access point which we call the master, the main, and the primary transmitter interchangeably.

A. The need for time and phase synchronization

Consider a distributed multiuser MIMO-OFDM scenario with K users and M access points, each with a single antenna. Let the bandwidth of the channel be divided into N orthogonal subcarriers, and let the discrete-time baseband complex equivalent channel impulse response, of a single transmit/receive pair, $\tilde{\mathbf{h}} = (\tilde{h}_0, \dots, \tilde{h}_L)^T$ (indices are omitted here for simplicity) have at most $L + 1$ significant taps. The DFT of this impulse response is given by $\mathbf{H} = (H(0), H(1), \dots, H(N-1))^T$ where $\mathbf{H} = \sqrt{N}\mathbf{F}(\tilde{\mathbf{h}}^T, \mathbf{0}^T)^T$. Here \mathbf{F} is the DFT matrix given by $[F]_{n,t} = \frac{1}{\sqrt{N}}e^{-j2\pi tn/N}$ and $(\tilde{\mathbf{h}}^T, \mathbf{0}^T)^T$ is the channel impulse response padded with zeros to length N .

Let the “frequency-domain” transmit symbols be grouped into blocks of length N , such that $\mathbf{x} = (x(0), \dots, x(N-1))$. The OFDM modulator performs an inverse DFT, $\tilde{\mathbf{x}} = \mathbf{F}^H \mathbf{x}$,

and cyclic prefixing (CP), i.e., it appends the last L time-domain symbols in front of each block $\tilde{\mathbf{x}}$, such that the block of actually transmitted time-domain symbols is $\tilde{\mathbf{x}}_{\text{cp}} = (\tilde{x}_{N-L}, \dots, \tilde{x}_{N-1}, \tilde{x}_0, \dots, \tilde{x}_{N-1})$. The OFDM demodulator collects the corresponding block of time-domain signal samples

$$\tilde{y}_{\text{cp},i} = \sum_{\ell=0}^L \tilde{h}_\ell \tilde{x}_{\text{cp},i-\ell} + z_i$$

for $i = 0, \dots, N-1$ (i.e., discarding the guard time due to CP), and apply DFT to the corresponding vector, obtaining the frequency domain received symbol vector $\mathbf{y} = \mathbf{F}\tilde{\mathbf{y}}$. Thanks to CP, the discrete-time convolution operated by the channel is turned into a cyclic convolution, such that the inverse DFT precoding and DFT processing at transmitter and receiver perfectly diagonalize the frequency selective channel into N parallel frequency domain channels given by

$$y(n) = H(n)x(n) + z(n), \quad n = 0, \dots, N-1,$$

where $z(n)$ are frequency-domain complex circularly symmetric AWGN samples assumed to have unit variance, after a suitable normalization of the received signal. For sufficiently small A/D sampling rate $1/T_s$, the frequency-domain channel response \mathbf{H} can be well approximated by the Fourier transform of the *continuous-time* channel impulse response $h(\tau)$, sampled in the frequency domain at frequencies $f_n = \frac{n}{NT_s}$ modulo $[-1/(2T_s), 1/(2T_s)]$, for $n = 0, \dots, N-1$. Notice that $h(\tau)$ includes both the physical propagation channel and the effects of the transmit and receive filters and D/A pulse-shaping. Notice also that, while OFDM can be applied for any N, L and sampling rate $1/T_s$. However, \mathbf{H} may be very different from the sampled continuous-time channel transfer function, and is given in general by the sampling (in frequency) of the spectral folded version of the channel transfer function. For conceptual simplicity, we omit this detail here.

In a distributed multiuser MIMO setting, the above modulation/demodulation is operated at each transmit antenna and at each user receiver. In this case, the received downlink signal at user k on subcarrier n is given by

$$y_k^{\text{dl}}(n) = \mathbf{h}_k^H(n)\mathbf{x}(n) + z_k(n), \quad (1)$$

where $\mathbf{h}_k(n) = (H_{1,k}(n), \dots, H_{M,k}(n))^T$ is the vector of frequency domain channel coefficients from the M access point antennas to user k receiver on subcarrier n , $\mathbf{x}(n)$ is the vector of frequency domain transmit symbols, and $z_k(n) \sim \mathcal{CN}(0, 1)$ is AWGN.

For nomadic users (typical of WLAN scenarios), the channel changes in time quite slowly, so that we may assume that the channel impulse response is locally invariant with time. For this reason, without loss of generality, we can consider the channel matrices as independent of the OFDM symbol index k over a Time Division Duplexing (TDD) frame of $B \gg 1$ OFDM symbols, comprising an UL slot, a guard time, and a DL slot. We wish to estimate the channel matrix between the access point antennas and the users’ antennas, for each subcarrier $n = 0, \dots, N-1$ using TDD. In order to do so,

each user sends a training signal, simultaneously over the UL slot. The uplink channel at access point m on any TDD slot and subcarrier n is given by

$$\mathbf{y}_m^{\text{ul}}(n) = \sum_{k=1}^K H_{m,k}(n) \mathbf{x}_k^{\text{ul}}(n) + \mathbf{z}_m^{\text{ul}}(n) \quad (2)$$

where all vectors have dimension $B_{\text{tr}} \times 1$ and $\mathbf{x}_k^{\text{ul}}(n)$ is the vector of training symbols sent by user k on subcarrier n , independent of t since it is repeated at every slot. Notice that here we focus only on channel estimation via the uplink, by exploiting the TDD reciprocity, hence we consider only the training phase of uplink slots. As long as $B_{\text{tr}} \geq K$, it is possible to assign to the K users a set of orthogonal training sequences such that $\mathbf{X}^{\text{ul}}(n) = [\mathbf{x}_1^{\text{ul}}(n), \dots, \mathbf{x}_K^{\text{ul}}(n)]$ is a scaled unitary matrix. Hence, each m -th access point can estimate its own channel vector $(H_{m,1}(n), \dots, H_{m,K}(n))$ simply as

$$(\hat{H}_{m,1}(n), \dots, \hat{H}_{m,K}(n))^{\text{T}} = \alpha (\mathbf{X}^{\text{ul}}(n))^{\text{H}} \mathbf{y}_m^{\text{ul}}(n),$$

where α is a suitable scaling coefficient. Each access point sends its channel estimates to the central server, which computes the DL MIMO precoder. For each subcarrier n , let

$$\begin{aligned} \mathbf{H}(n) &= [\mathbf{h}_1(n), \dots, \mathbf{h}_K(n)] \\ &= \begin{bmatrix} H_{1,1}(n) & \cdots & H_{1,K}(n) \\ \vdots & \ddots & \vdots \\ H_{M,1}(n) & \cdots & H_{M,K}(n) \end{bmatrix} \end{aligned} \quad (3)$$

denote the uplink channel matrix. The corresponding assumed downlink matrix is $\mathbf{H}^{\text{H}}(n)$. For simplicity and clarity of exposition, we consider DL precoding based on linear Zero-Forcing (ZF) beamforming in the case of $K \leq M$, for which the precoding matrix $\mathbf{V}(n)$ is such that $\mathbf{H}^{\text{H}}(n)\mathbf{V}(n) = \mathbf{\Lambda}(n) = \text{diag}(\Lambda_1(n), \dots, \Lambda_K(n))$, for some real positive $\Lambda_1(n), \dots, \Lambda_K(n)$, and $\mathbf{V}(n)$ has unit norm columns in order to preserve the total transmit power. In general, $\mathbf{V}(n)$ is obtained by normalizing the columns of the Moore-Penrose pseudoinverse [4] of $\mathbf{H}^{\text{H}}(n)$. Finally, the transmitted DL signal for OFDM symbols t in the DL slot is given by

$$\mathbf{x}^{\text{dl}}(n, t) = \mathbf{V}(n) \mathbf{u}(n, t) \quad (4)$$

$$n = 0, \dots, N-1, \quad t = 1, \dots, B_{\text{dl}} \quad (5)$$

where $\mathbf{u}(n, t)$ denotes the information symbol vector destined to the users and B_{dl} denotes the number of information-bearing OFDM symbols in the DL slot. We assume that $B_{\text{tr}} + B_{\text{dl}} < B$, in order to include a small guard interval between UL and DL slots.

If the UL and DL slots in the same TDD frame have exactly the same timing reference and the same carrier phase reference, the received signal at the users' receivers, on each subcarrier $n = 0, \dots, N-1$, can be written as

$$\mathbf{y}^{\text{dl}} = \mathbf{H}^{\text{H}}(n) \mathbf{V}(n) \mathbf{u}(n) + \mathbf{z}^{\text{dl}}(n) = \mathbf{\Lambda}(n) \mathbf{u}(n) + \mathbf{z}^{\text{dl}}(n) \quad (6)$$

Since the overall channel matrix $\mathbf{\Lambda}(n)$ is diagonal, we have achieved complete user separation, so that the DL can serve the K users on the same DL slot, without interference.

If timing and carrier phase synchronization was perfect, then the DL channel from access point m to user k would have impulse response $h_{m,k}^*(-\tau)$ (time-reversal of the corresponding UL impulse response $h_{m,k}(\tau)$). Instead, due to misalignment of the sampling clocks and phases introduced by the RF oscillators, the impulse response is $h_{m,k}^*(-\tau - \tau_m - \delta_k) e^{j(\phi_m + \theta_k)}$ where τ_m, ϕ_m are the timing and carrier phase shifts between UL and DL slots at access point m and δ_k, θ_k are the timing and carrier phase shifts between UL and DL slots at user k . From well-known rules of linearity and time-shift of DFT, we arrive at the following expression for the effective DL channel matrix:

$$\tilde{\mathbf{H}}^{\text{H}}(n) = \mathbf{\Theta}(n) \mathbf{H}^{\text{H}}(n) \mathbf{\Phi}(n) \quad (7)$$

Where

$$\mathbf{\Theta}(n) = \text{diag} \left(e^{j(f_n \delta_k + \theta_k)} : k = 1, \dots, K \right) \quad (8)$$

and

$$\mathbf{\Phi}(n) = \text{diag} \left(e^{j(f_n \tau_m + \phi_m)} : m = 1, \dots, M \right) \quad (9)$$

We notice that the diagonal matrix of phasors $\mathbf{\Theta}(n)$ multiplying the nominal channel matrix from the left poses no problems, since these phase shifts can be recovered individually by each user terminal as in standard coherent communications [5]. In contrast, the diagonal matrix $\mathbf{\Phi}(n)$ multiplying from the right poses a big problem. In fact, the server computes the MIMO precoder $\mathbf{V}(n)$ on the basis of $\mathbf{H}^{\text{H}}(n)$, so that when applied to the effective channel in (7) the matrix multiplication $\tilde{\mathbf{H}}^{\text{H}}(n) \mathbf{V}(n)$ is no longer diagonal in general. We conclude that the presence of timing and carrier phase misalignment between the UL and DL slots, at each individual access point, yields residual multiuser interference which may completely destroy the performance of a distributed multiuser MIMO system.

B. A distributed MIMO testbed

Software Radio Implementation. We have implemented a time and phase synchronization scheme as a digital circuit in the FPGA of the WARP radio platform [6]. The WARP radio is a modular software radio platform composed of a central motherboard hosting an FPGA and several daughterboards containing radio frequency (RF) front-ends. The entire timing of the platform is derived from only two reference oscillators, hosted on a separate clock board: a 20 MHz oscillator serving as a source for all sampling signals and a 40 MHz oscillator which feeds the carrier clock inputs of the transceivers present on the RF front-ends.

The Synchronization Circuit. Our synchronization circuit operates similarly to other OFDM-based, distributed transmission systems such as SourceSync [7] or Fine Grained Channel Access [8], but extends them by achieving phase synchronization among transmitters. An important component of those systems, essential in order to avoid leakage from one carrier to another during the decoding process, is the realization of frame alignment that arranges frame starting points at the receivers within an interval shorter than a CP

length. In other words, the overlap of the frames sent by different senders must be greater than the length of a frame without CP in order to allow the receiver to perform a full-length discrete Fourier transform on the received signal.

We achieve frame synchronization through a technique used in block boundary detection, namely the insertion of pseudo-noise (PN) sequences in the master access point's packet header in order to allow the secondary transmitters and the receivers to obtain a time reference. For reasons that will become clear, achieving frame synchronization within the length of the CP is a sufficient starting point for also achieving phase synchronization. We measure phase drift using pilot tone signals. In order to reduce self-interference at the secondary transmitters, the tone signals are placed outside the data band, from which they are separated by a large guard interval. The secondary transmitters place an analog baseband filter around their data band further limiting their interference with the pilots.

The secondary transmitters overhear a packet sent by the primary transmitter and use the initial PN sequence in order to determine the block boundary timing of this packet. Using a DFT the secondary transmitters decode the successive frames of the incoming packet. They then employ the CORDIC algorithm on the complex-valued received soft symbols in order to obtain their phases in radians. The phases of the out-of-band pilot signals are tracked throughout the entire packet transmission in order to estimate the phase drift from the primary sender. The measurements from the different pilots are averaged and passed through a simplified Kalman filter which maintains an accurate estimate and predicts, based on the current estimate, the phase drift after the passage of a few further frames. In addition, the header sent by the primary sender contains a number of channel estimation symbols, used to obtain an initial phase offset estimate for each subcarrier. Since the phase drift is almost identical for all carriers, these two measurements suffice in order to predict the phase rotation induced by the main transmitter on any subcarrier tone for the entire period of a packet.

The phase estimates are used in synthesizing a synchronized signal. The secondary transmitters use an inverse discrete Fourier transform, whose output frames are timed such that they align with the frames of the main sender's signal. For every subcarrier the secondary transmitters rotate the soft symbol to be sent by an angle corresponding to the subcarrier's estimated phase offset. The result is a tone that, while not having the same phase as the corresponding tone from the main transmitter, follows that tone at a fixed, pre-known phase difference.

Centralized joint encoding. By transmitting phase synchronous signals from multiple access points we have created the equivalent of a distributed MIMO transmitter, capable of employing multiuser MIMO precoding strategies in order to transmit to multiple users at the same time. However, the use of multiple access points complicates the design of the transmitter system. For most of the precoding schemes available, the encoding of the waveforms to be transmitted over the antennas

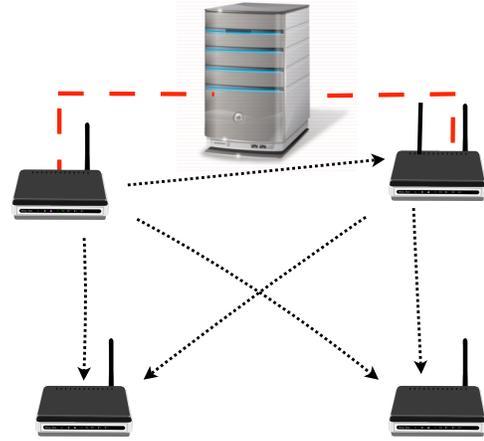


Fig. 1. Testbed diagram. The central server is connected to the two transmitters, the main transmitter on the left and the secondary transmitter on the right.

must be done jointly, since reaching a single user usually involves transmitting over multiple antennas. While in theory the joint encoding process could be duplicated at each access point given the binary information destined to each user, we chose to do the encoding only once, at a central server and send the quantized resulting waveforms to each access point for transmission¹.

C. Experimental results

Our system setup is presented in Figure 1. It consists of a primary transmitter, a secondary transmitter and two receivers. The main sender uses a single RF front-end configured in transmit mode, placing an 18 MHz shaping filter around the transmitted signal. The secondary sender uses an RF front-end in receive mode and a second RF front-end in transmit mode, with a 12 MHz shaping filter. As mentioned previously, the pilots used in phase tracking are outside the secondary's transmission band, therefore the secondary transmitter will not interfere with the pilot signals from the main transmitter. The series of experiments is intended to test the accuracy of the synchronization and the efficiency of channel separation.

Synchronization Accuracy. In this particular experiment we have placed the two transmitters and the two receivers at random locations. We placed a third RF front-end on the secondary sender and configured it in receive mode. The secondary transmitter samples its own synthesized signal over a wired feedback loop and compares it with the main transmitter's signal. The synchronization circuit measures and records the phase differences between these two signals. Since we use the primary transmission as a reference, in this experiment we do not broadcast the signal synthesized by the secondary transmitter in order to protect the primary transmission from unintended interference.

Our experiments reveal that the synchronization error between the secondary transmitter and the primary transmitter is

¹This approach is practical in enterprise networks where a number of access points are already connected to a common server.

quite negligible. The error is measured on a frame-to-frame basis using the feedback circuit. In decimal degree values, the average absolute phase error of the synchronized signal is 1.88 degrees while its the standard deviation is 2.37 degrees. The 95th percentile of the synchronization error is at most 4.5 degrees.

Beamforming gain. Our second experiment was done using the complete four radio setup with the secondary transmitter broadcasting a secondary signal over the air. We measured the channel coefficients between the two transmitters and the receivers using standard downlink channel estimation techniques and arranged the amplitudes and the phases of the transmitted signals such that at one of the receivers the amplitudes of the two transmitted signals would be equal while the phases would align. The maximal theoretic power gain over transmitting the two signals independently is 3.01dB. We compared the average power of the individual transmissions from the two senders to the average power of a beamformed joint transmission. Our measurements show an average gain of 2.98 dB, which is consistent with the precision of the synchronization determined in the previous experiment.

This result shows that for all practical purposes we are able to achieve the full beamforming gain in our testbed.

Zero-Forcing Accuracy. The following experiment measures the amount of power which is inadvertently leaked when using ZF beamforming to non-targeted receivers due to synchronization errors. Again we have placed our radios at random locations in our testbed. We have estimated the channel coefficients and arranged for two equal amplitude tones from the two transmitters to sum as closely as possible to zero. The residual power is the leaked power due to angle mismatching. Our experiments indicate that the average power leaked is only -24.46 dB of the total transmitted power.

This demonstrates in a practical implementation setting that distributed ZF is capable of almost completely eliminating interference at non-targeted receiver locations, despite the two transmitters are not connected to the same precise timing and phase reference.

Zero-Forcing Beamforming Data Transmission. The final experiment transmits data to the two receivers. We have used symbols chosen independently from a QAM-16 constellation at similar power levels. The scattering plots in Figure 2 illustrates the received signal at the two receivers.

The SINR values at the two receivers are 29 dB and 26 dB respectively. It is evident that the testbed achieves the full MIMO multiplexing gain. This motivates us to move one step further and investigate scheduling/MAC issues in the rest of this paper.

III. OPTIMAL AND PRACTICAL SCHEDULING

The design of a Media Access Control (MAC) layer that can fully exploit the characteristics of the underlying distributed multiuser MIMO physical 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

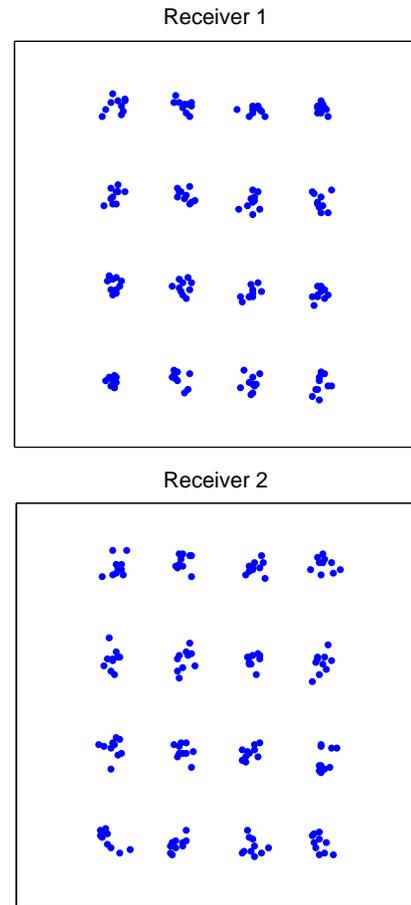


Fig. 2. Scattering Diagram. The scattering diagram for two independent data streams transmitted concurrently demonstrates that AirSync achieves complete separation of the user channels.

from the access point to the users). We can choose between two natural strategies for separating the uplink from the downlink: time division duplex (TDD) and frequency division duplex (FDD). The advantage of using TDD is that we can exploit channel reciprocity at the access point and measure the uplink channel (by using pilots from the users) to infer the downlink. On the other hand, in FDD, the uplink and downlink channels are in general different and therefore there is no reciprocity and the transmitter has to rely on feedback information that requires extra overhead as the number of access point antennas grows [9]. Moreover, TDD is ideally suited for the transport of asymmetric traffic, as is typical in an enterprise wifi environment, and it allows to allocate accordingly the bandwidth to each direction. An FDD system provides less flexibility for managing different traffic patterns and it is the scheme of choice when traffic is symmetric.

We will start by discussing the issue of optimally scheduling users in the DL of the distributed multiuser MIMO system model we described. Notice that, after time and phase synchronization between the decentralized access points of our system has been achieved, we can treat this setup as multiuser MIMO where all antennas are located on the same access point.

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. DPC, being the optimal, capacity achieving strategy for the underlying MIMO Gaussian broadcast channel, is the natural term of comparison for any practical low-complexity scheme. As anticipated earlier, ZF beamforming is a practical, and easily deployable choice for a broadcast channel like ours [10], [11], [12]. The system model consists of M antennas forming a distributed MIMO system as described in the previous sections and U single-antenna users waiting to be serviced. However, 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 $|S(n)| = M$) 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:

$$\begin{aligned}
& \text{maximize} && R(S) = \sum_{n=0}^{N-1} \sum_{k \in S(n)} W_k \log_2(1 + \text{SINR}_k(n)), \\
& \text{with respect to} && \{S(n) \subseteq \{1, \dots, K\} : |S(n)| \leq M\}, \\
& && \{P_k(n)\} \\
& \text{subject to} && \frac{1}{N} \sum_{n=1}^N \sum_{k \in S(n)} P_k(n) \leq P_{\text{sum}}
\end{aligned} \tag{10}$$

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 ZF beamforming, $\{W_k\}$ 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). From Section II, with due normalizations, we have

$$\text{SINR}_k(n) = \Lambda_k^2(n) P_k(n)$$

where $\Lambda_k(n)$ is the k -th ZF beamforming coefficient of the channel submatrix of the overall $U \times M$ channel matrix $\mathbf{H}(n)$ corresponding to the users in $S(n)$. 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 [13], [14], [11], [15]. We extend the greedy user selection algorithm of [11] for multiple subcarriers and we have the following algorithm where as R we denote the achievable rate sum using Gaussian rates for the scheduled user/subcarrier pairs.

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

Algorithm 1 Greedy ZF with Waterfilling and Gaussian rates (GZF)

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

equation:

$$\sum_{n=0}^{N-1} \sum_{k \in S(n)} \left[\frac{W_k}{\mu} - \frac{1}{\Lambda_k^2(n)} \right]^+ = N \cdot P_{\text{sum}} \tag{11}$$

that can be derived from the convex optimization problem (10).

Though DPC achieves capacity and GZF provides a convenient sub-optimal algorithm, there are practical considerations beyond simple GZF 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 coding/modulation pairs, 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 III provides one such mapping, keeping in mind that mappings may vary or be dynamically chosen in practical scenarios. Using this scheme within the greedy ZF user selection, we examine the resulting method, dubbed greedy zero forcing-quantized (GZF-Q).

802.11n MCS Index	Modulation	Code Rate	SNR Range
0	BPSK	1/2	$\geq 0.5\text{dB}$
1	QPSK	1/2	$\geq 3.5\text{dB}$
2	QPSK	3/4	$\geq 6.2\text{dB}$
3	16-QAM	1/2	$\geq 8.9\text{dB}$
4	16-QAM	3/4	$\geq 12.3\text{dB}$
5	64-QAM	2/3	$\geq 16.1\text{dB}$
6	64-QAM	3/4	$\geq 17.5\text{dB}$
7	64-QAM	5/6	$\geq 19.0\text{dB}$

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

Since the power allocation step in GZF 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 GZF-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 GZF-RR, and expect that its long term performance should be close to that of GZF.

IV. SIMULATIONS

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

A. Sum Rate

We computed, 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). We considered a system with $U = 10$ users and $M = 4$ antennas. In Figure 3 we see that DPC clearly dominates all other schemes as expected.

Nevertheless, GZF is near optimal for the low SNR regime, and has a constant gap of 3 bits/s/Hz for high SNR for this level of multi-user diversity. Notice that, for the medium to high SNR regime, the GZF-P achieves the same throughput with close to zero losses from the GZF with waterfilling. The GZF-RR strategy is also proven to provide almost full GZF multiuser diversity gains for all SNR. This was expected because only the first user in the greedy selection is indicated from the round-robin strategy, and thereafter waterfilling GZF is employed to schedule the remaining users. On the other hand, the GZF-Q has a significant gap from the GZF strategies with Gaussian rates. Nevertheless, this gap remains constant at 13 bits/s/Hz for high SNR and significantly smaller for lower SNR values. This gap can be reduced by selecting a more powerful family of MCSs, and possibly introducing a continuum of rates through the use of rateless codes and incremental redundancy at the MAC/PHY layer, a topic that we are currently investigating from the viewpoint of a practical implementation on our software-defined radio testbed.

B. Average Delay

To evaluate the average delay of the greedy scheduling algorithms we implemented a queuing system with random

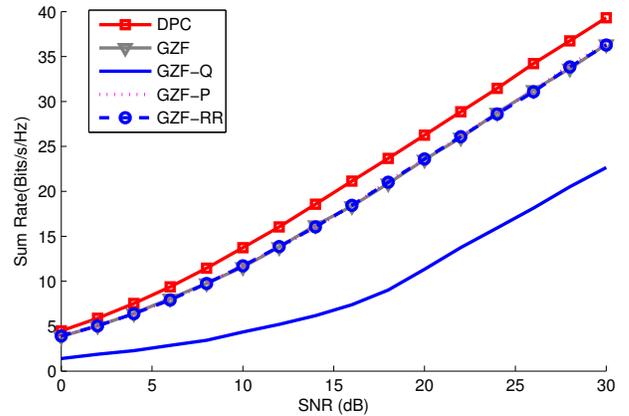


Fig. 3. Sum Rate vs SNR for $U = 10$ users, $M = 4$ antennas, and $N = 4$ subcarriers.

arrivals for every user and updated the queues, based on the scheduling and power allocation of our algorithms, after every DL scheduling slot. 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 using Little's law. For the sake of computational complexity, we choose to simulate it for 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 GZF-Q is the first to give unbounded delays for an arrival rate of 0.5 bits/s. GZF-RR gives an achievable arrival rate of 1 bit/s which is very close to the 1.1 bits/s the GZF and GZF-P schemes can achieve for this SNR. The average delay for this symmetric arrival system is worse for the case of the GZF-RR compared to GZF. However, as mentioned earlier, it will greatly improve the per user delay in a non-symmetric arrival rate scenario. Imagine a situation 9 of users have an average arrival of 0.9 bits/s/time slot in their queues and the 10th has only 0.01. In Figure 5 we observe that the GZF-RR scheduler manages to significantly decrease the delay in comparison to the vanilla GZF for this user, without a noticeable increasing the delay of other users. Therefore, in order to avoid unexpected timeouts of upper layer protocols running over this scheme, it makes sense to choose the GZF-RR policy where even users with small queue arrival rates (for instance a single email) get served at least once per U scheduling slots.

V. PROTOCOL DESIGN

In this section we discuss the design of data packets with preambles which support the MAC schemes we have investigated. The MAC layer packet design and the protocol's sequence of actions are tuned for enabling multiuser MIMO broadcasts. The crucial design constraint is to provide the

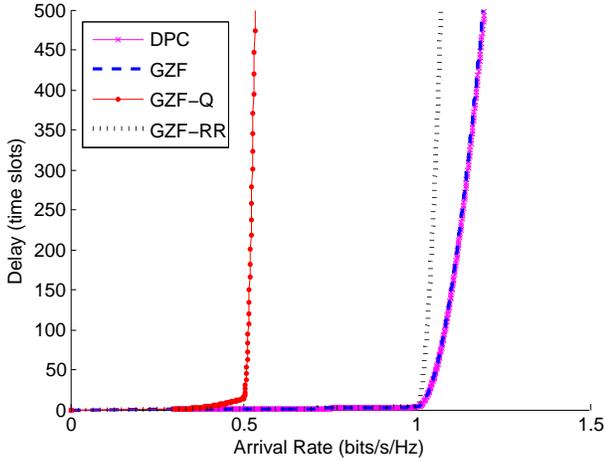


Fig. 4. Average queue delay for $U = 10$ users, $M = 4$ antennas, and $N = 4$ subcarriers and $\text{SNR} = 10$ dB.

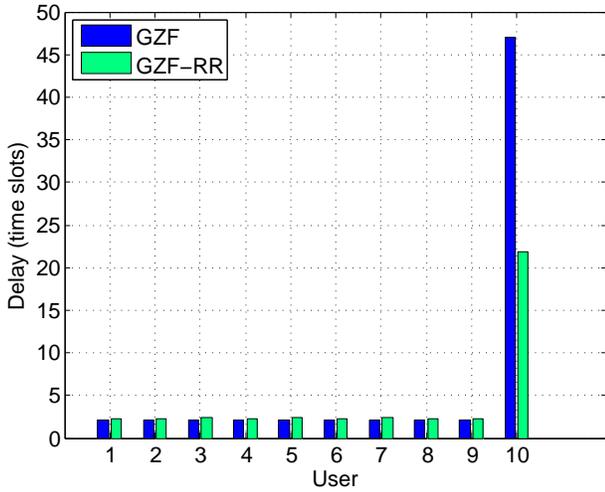


Fig. 5. Delay comparison of GZF vs. GZF-RR.

central server with timely estimates of the channel state information for all clients to which it is about to transmit or which are considered for the next round of transmissions. For this purpose, we schedule downlink transmissions to closely follow uplink acknowledgments and require the clients to provide the server with channel estimates during the uplink period. The mechanism through which this is achieved will be described in the following paragraphs. The central server uses the uplink estimates to select a set of clients for the following transmission slots, according to the scheduling algorithms introduced earlier.

The downlink packet starts with a transmission from the main sender containing a pseudo-noise sequence used to achieve frame alignment by the transmitters and for block boundary detection by the receivers. The master access point then transmits the first set of channel estimation pilots which are used by the other access points to determine the initial phases of the subcarrier tones, as described in Section II. After this point, all access points take part in the downlink

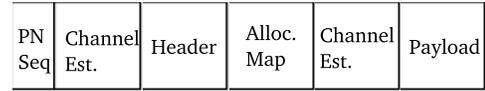


Fig. 6. Downlink data packet design.

transmission. The packet header that follows is broadcast to all clients, including the non-targeted ones, using the Alamouti encoding [16]. Due to phase alignment between transmitters, the clients do not need to track the secondary senders in order to decode this header. The MAC addresses of the hosts targeted in the current transmission and the MAC addresses of the clients that are required to provide the server with channel estimates during the next acknowledgment are the most important pieces of information contained in the header fields. The positions of the addresses in the header fields create an implicit ordering of the clients, which will be used in the uplink period. The following part of the header is an allocation map, similar to the one found in the LTE standard, which assigns carriers to small groups of different clients and specifies the constellations used in broadcasting to them. The header is followed by a second set of channel estimation pilots, transmitted this time around by all access points using ZFBF, which are used by all clients in order to obtain the channel estimates for their individual downlink channels. The clients use the downlink estimates together with the synchronization pilot tones in order to gain a lock on the subcarriers. The downlink transmission continues with payload transmission. Figure 6 presents a simplified schematic of downlink data packets along the lines of the discussion above.

In current 802.11 MIMO implementations, the channel estimates are obtained using downlink pilots which are in turn quantized by the receivers and communicated back in numerical form to the transmitter. The quantization and communication steps incur a large overhead. Using the reciprocity property of wireless channels, we can reduce the complexity of the channel estimation process significantly. First, we prefer to perform uplink channel estimation since uplink estimates can be received simultaneously by all access points, reducing the number of pilot transmissions needed by a factor equal to the total number of access point antennas. Second, uplink estimates are sent using analog pilot signals in an unquantized form, leaving the quantization step to the access points. This reduces the overhead of the transmission significantly. Third, while the usual estimation pilots are full OFDM frames, we choose to send pulse-like signals, measure the channel response, and fill the non-significant taps with zeros before taking a Fourier transform in order to determine the frequency domain response. This insures that our pilots need to be spaced only by an interval that can accommodate a long channel response, i.e. the length of a cyclic prefix.

After the downlink transmission has finished, the clients who have been requested to send their channel estimates start sending these short estimation pilots in quick succession. We note that there is a large degree of similarity between the

functioning of the downlink channel estimation for receive decode purposes and the uplink channel estimation step. The timing of the system remains unchanged during the uplink slot and the roles of the transmitters and the receivers are switched. The uplink pilots are followed by smart acknowledgments for the data packets sent using the technique detailed in [17].

Overhead. A note on the overhead of the above MAC is in order. As those familiar with the PHY/MAC details of the 802.11 family of protocols would have recognized already, the overhead of our MAC is not more than that of 802.11n. The additional signaling overhead comes from requiring a few frames to predict the initial phase, and a few frames to dictate the MAC addresses of the nodes from which we wish to request channel state information for the next time slot. Even with very conservative estimates this will be less than a 20% increase in header time duration over that of a traditional 802.11 system. Note, however, that we get a bandwidth increase that grows almost linearly in the number of clients. This means that our overhead, normalized such that we consider the total control bits over the total data bits transmitted during a fixed airtime slot, is much less than in a traditional 802.11 system.

VI. CONCLUSION AND FUTURE WORK

While this paper reports admittedly some work in progress, we would like to point out here the steps of our system testbed implementation that we are currently undertaking. Based on the described timing and phase synchronization over the air, the next step in the testbed implementation will include the actual UL channel estimation and DL ZF beamforming with scheduling and dynamic power/rate allocation.

Given the significant degradation observed here, between the theoretical optimum Gaussian dirty-paper coding and a quantized rate system that makes use of a family of state of the art MCSs, we are currently studying the implementation of viable rateless coding schemes at the MAC/PHY layer. There are a few main challenges to be addressed: 1) 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, we have to modify the scheduler in order to recursively compute the scheduling weights not on the basis of the allocated service rate, but on the basis of the *mutual information* accumulated at each receiver (see for example [18]; 2) the most natural way of implementing rateless codes at the PHY layer consists of mapping the output of a Raptor encoder [19] onto modulation alphabets, chosen according to a quantized water filling scheme. This requires the use of soft demapping in order to produce log-likelihood ratios for the raptor decoder as in a BICM configuration [20]. A recent attractive alternative consists of using spinal codes [21]. A thorough comparison of these two options in the context of our distributed multiuser MIMO MAC/PHY is the subject of current investigation; 3) in both cases, the choice of the modulation alphabet is very important. In the Raptor/BICM case, we can choose variable size modulation alphabets depending on the instantaneous channel gains resulting from

the GZF with relative power allocation. In contrast, spinal codes make use of a single large modulation alphabet. While going for larger and larger modulation alphabets and lower underlying binary coding rate is a good strategy in principles, and can approach the Gaussian channel capacity up to the so-called shaping loss of 0.25 bits per dimension, in practice large modulation alphabets suffer from increased sensitivity to imperfect synchronization and mismatch in the demodulator. Therefore, it is not a priori clear which scheme should be preferred in practical conditions.

REFERENCES

- [1] S. Gollakota, S. D. Perli, and D. Katabi, "Interference alignment and cancellation," in *ACM SIGCOMM*. New York, NY, USA: ACM, 2009, pp. 159–170.
- [2] S. Jayaweera, "Energy efficient virtual mimo-based cooperative communications for wireless sensor networks," in *ICISIP*, jan. 2005, pp. 1 – 6.
- [3] 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.
- [4] 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.
- [5] J. Proakis and M. Salehi, *Digital communications*. New York, NY: McGraw-Hill, 2007.
- [6] Rice University, "Rice university warp project." [Online]. Available: <http://warp.rice.edu>
- [7] H. Rahul, H. Hassanieh, and D. Katabi, "SourceSync: a distributed wireless architecture for exploiting sender diversity," in *ACM SIGCOMM*, New Delhi, India, 2010. [Online]. Available: <http://doi.acm.org/10.1145/1851182.1851204>
- [8] 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. [Online]. Available: <http://doi.acm.org/10.1145/1851182.1851202>
- [9] J. Jose, A. Ashikhmin, P. Whiting, and S. Vishwanath, "Channel estimation and linear precoding in multiuser multiple-antenna TDD systems," *IEEE Trans. Veh. Technol.*, vol. 60, no. 5, pp. 2102–2116, Jun. 2011.
- [10] H. C. Huang, S. Venkatesan, and H. Viswanathan, "Downlink capacity of cellular networks with known interference cancellation," in *DIMACS Workshop on Sig. Proc. for Wireless Commun.*, Oct 2002.
- [11] 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.
- [12] T. Yoo and A. Goldsmith, "On the optimality of multi-antenna broadcast scheduling using zero-forcing beamforming," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 3, pp. 528 – 541, Mar. 2006.
- [13] 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.
- [14] —, "A practical approach for weighted rate sum maximization in MIMO-OFDM broadcast channels," in *IEEE ACSSC*, Pacific Grove, CA, Nov. 2007.
- [15] —, "Joint beamforming and scheduling for a MIMO downlink with random arrivals," in *IEEE ISIT*, Seattle, WA, Jul. 2006.
- [16] S. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 8, pp. 1451–1458, 1998.
- [17] A. Dutta, D. Saha, D. Grunwald, and D. Sicker, "SMACK: a SMart ACKnowledgment scheme for broadcast messages in wireless networks," in *ACM SIGCOMM*, Barcelona, Spain, 2009. [Online]. Available: <http://doi.acm.org/10.1145/1592568.1592572>
- [18] H. Shirani-Mehr, H. Papadopoulos, S. Ramprasad, and G. Caire, "Joint scheduling and arq for mu-mimo downlink in the presence of inter-cell interference," *IEEE Trans. Comput.*, vol. 59, no. 2, pp. 578–589, february 2011.
- [19] A. Shokrollahi, "Raptor codes," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2551–2567, Jun. 2006.

- [20] G. Caire, G. Taricco, and E. Biglieri, "Bit-interleaved coded modulation," *IEEE Trans. Inf. Theory*, vol. 44, no. 3, pp. 927–946, may 1998.
- [21] 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>