

Exploiting Hybrid Channel Information for Downlink Multi-User MIMO Scheduling

Wenzhuo Ouyang*, Narayan Prasad† and Sampath Rangarajan†

*Ohio State University, Columbus, OH

†NEC Labs America, Princeton, NJ

e-mail: ouyangw@ece.osu.edu, {prasad, sampath}@nec-labs.com

Abstract—We investigate the downlink multi-user MIMO (MU-MIMO) scheduling problem in the presence of imperfect Channel State Information at the transmitter (CSIT) that comprises of coarse and current CSIT as well as finer but delayed CSIT. This scheduling problem is characterized by an intricate ‘exploitation - exploration tradeoff’ between scheduling the users based on current CSIT for immediate gains, and scheduling them to obtain finer albeit delayed CSIT and potentially larger future gains. We solve this scheduling problem by formulating a frame based joint scheduling and feedback approach, where in each frame a policy is obtained as the solution to a Markov Decision Process. We prove that our proposed approach can be made arbitrarily close to the optimal and then demonstrate its significant gains over conventional MU-MIMO scheduling.

I. INTRODUCTION

Multiple Input Multiple Output (MIMO) technology is essential for the emerging 4G-LTE wireless communication systems. In the downlink of such a system, which typically has several active users, multiple antennas enable simultaneous transmissions to multiple users by allowing the transmitter (base-station) to transmit (along directions in a signal space) in a manner which ensures that each user can receive its intended signal along at least one interference-free dimension (a.k.a. the Multi-user MIMO principle) [1]. The number of active users is generally greater than the maximum supportable number of simultaneous transmissions, which in turn is equal to the number of transmit antennas at the base-station (BS). Consequently, only a subset of users can be selected for the MU-MIMO transmission and hence proper user scheduling is important to achieve a desired network utility (e.g., throughput, fairness).

The usual assumption made in existing literature on MU-MIMO scheduling is that the BS can obtain the channel state information from all users with sufficient accuracy and with negligible delay. Such information, referred to as the Channel State Information at the Transmitter (CSIT), is crucial to ensure that each scheduled user is not dominated by co-channel interference. Typically, the BS obtains CSIT by broadcasting a sequence of pilot symbols, and the users in turn estimate their CSI and feedback their quantized estimates to the BS. This feedback process introduces two sources of imper-

fections to the CSIT. (1) Estimation and quantization errors (due to limited training and finite codebooks); (2) Delays (due to user processing speeds and less flexible scheduling on the feedback channel). The impact of erroneous CSIT on MU-MIMO performance has been analyzed in [2] and utility maximization for MU-MIMO with erroneous CSIT has been considered in [4]. Delay in the CSIT has hitherto been addressed by using prediction based approaches but their drawback is that they have to assume a model for channel evolution, which is significantly difficult to obtain in practice and they also require the delay to be small enough to allow for useful prediction.

For the scenario where the number of users is small enough so that user scheduling is unnecessary, referred to here as the static scenario, Maddah-Ali and Tse proposed a scheme, namely the MAT scheme [5], that utilizes CSIT that is error-free albeit completely outdated. Their seminal work revealed that the outdated CSI is an important resource that, when combined with the eavesdropped information at the users, can provide a considerable performance gain in terms of degrees of freedom. Recently, the MAT scheme was extended (for the static scenario) to the hybrid CSIT case by also incorporating coarse and current CSIT [6] to obtain further system gains. However, in the ubiquitous setting where user scheduling is important, such hybrid CSIT needs to be exploited wisely since it is costly to obtain even delayed but error-free CSI feedback *from all users* for making the scheduling decisions. Indeed, the problem is quite different and more challenging than the static case. User scheduling for the MAT scheme has been considered in [3] but their suggested method is akin to the myopic approach discussed later in this paper.

In this paper, we study MU-MIMO downlink scheduling with hybrid CSIT, erroneous as well as delayed, where the time axis is divided into separate scheduling intervals. We consider the realistic scenario where current and coarse CSIT is obtained from all users while more accurate (not necessarily perfect) but delayed CSIT is obtained *only from the scheduled users*. The scheduling problem is hence characterized by an intricate ‘exploitation - exploration tradeoff’, between scheduling the users based on current CSIT for immediate gains,

and scheduling them to obtain finer albeit delayed CSIT and potentially larger future gains. The contributions of the paper are listed as follows.

- We tackle the aforementioned ‘exploitation - exploration tradeoff’ by formulating a frame based joint scheduling and feedback approach, where in each frame a policy is obtained as the solution to a Markov Decision Process (MDP), the latter solution being determined via a state-action frequency approach [10][11].

- We consider a general utility function and associate a virtual queue with each user that guides the achieved utility for that user. Based on MDP solutions and virtual queue evolutions, we show that our proposed frame-based joint scheduling and feedback approach can be made arbitrarily close to the optimal.

In the following we use $(\cdot)^T, (\cdot)^\dagger$ for the transpose and conjugate transpose, respectively. Moreover, $[\mathbf{A}, \mathbf{B}]$ and $[\mathbf{A}; \mathbf{B}]$ are used to denote column-wise and row-wise concatenation of matrices \mathbf{A} and \mathbf{B} , respectively. $\|\mathbf{A}\|$ is used to denote the Frobenius norm of the matrix \mathbf{A} .

II. SYSTEM MODEL AND PROBLEM FORMULATION

We consider the downlink MU-MIMO scheduling problem with one Base Station (BS) and N users. The BS is equipped with M_t transmit antennas and employs linear transmit precoding. Each user is equipped with a single receive antenna. Time is divided into intervals and we let $\mathbf{h}_i[k] \in \mathbb{C}^{1 \times M_t}$, $i = 1, \dots, N$ denote the channel state vector seen by user i in interval k . In each interval, a subset of users can be simultaneously scheduled. Further, since each user has only one receive antenna, it can achieve at-most one degree of freedom (i.e., its average data rate per channel use can scale with SNR as $\log(\text{SNR})$). On the other hand, the system can achieve at-most M_t degrees of freedom in that the total average system rate can scale with SNR as $M_t \log(\text{SNR})$. For notational convenience we assume that in each interval *two users* can be simultaneously served, hence limiting the achievable system degrees of freedom to 2. All results can however be extended to the general case without this restriction.

A. Conventional MU-MIMO scheme

Conventional MU-MIMO scheme relies on estimates of the user channel states (that are available at the BS) for the current interval. Indeed, perfect CSIT for the current interval enables the BS to transmit simultaneously to both scheduled users without causing interference at either of them. However, in the absence of perfect CSIT such complete interference suppression via transmitter side processing is no longer possible and when only very coarse estimates for the current interval are available, conventional MU-MIMO breaks down and in fact becomes inferior to simple single-user per interval transmission.

B. Joint Scheduling and Channel Feedback

We consider a joint scheduling and channel feedback scheme that builds upon a variant of the extended MAT technique [6]. Specifically, we assume that coarse quantized channel state estimates from all users for the current interval are available to the BS, along with limited finer albeit outdated quantized channel state estimates. In this context we note that in the FDD downlink only quantized estimates are available to the BS and henceforth unless otherwise mentioned, we will use ‘‘estimates’’ to mean ‘‘quantized estimates’’. The time duration of interest is divided into intervals with each interval comprising of 3 slots each. The three slots are mutually orthogonal time-bandwidth slices. For convenience, we assume that all three slots in an interval are within the coherence time and coherence bandwidth window so that the channel seen by each user remains constant over the three slots in an interval. At the beginning of the k^{th} interval, whose corresponding slots are denoted by $[k, 1], [k, 2]$ and $[k, 3]$, the scheduler broadcasts a short sequence of pilot symbols to all the users. This sequence enables a coarse estimation of the wireless channel at each of the N users, which is fed back to the BS after quantization and is denoted by $\hat{\mathbf{H}}[k] = \{\hat{\mathbf{h}}_i[k], i = 1, \dots, N\}$, where $\hat{\mathbf{h}}_i[k]$ denotes the coarse channel estimate obtained from user i for interval k . Based on these coarse estimates, along with its past scheduling and channel state history (formally introduced next), the scheduler chooses a pair of users to schedule in the current interval, where in the first slot a linear combination of new packets is sent for the selected user pair. Data transmission to the selected user pair in the current interval also contains additional pilots that enable a finer estimation of the channel states seen by that user pair over the current interval. Note that such finer estimation is crucial for data detection. However, due to user processing and feedback delays, we assume that (quantized versions of) such finer estimates are not available to the BS during the current interval itself. Because of this constraint, instead of performing the transmissions in slots 2 and 3 for interference resolution for the packets sent in Slot 1 of the current interval, as would be done in the extended MAT scheme [6], the BS performs transmissions for interference resolution for packets sent in Slot 1 of the *prior most recent interval* when the selected user pair was scheduled. The scheduling model is illustrated in Fig. 1.

As mentioned above the scheduler obtains a finer estimate of the channel states seen by a user pair on the interval in which they are scheduled, at the end of that interval.¹ Let $\theta = (u_1, u_2, \kappa)$ represent the 3-tuple denoting the scheduling decision made for the current interval k such that u_1, u_2 denote the selected user

¹Arbitrary delays in obtaining such finer estimates are also considered later in the paper.




Fig. 1. Illustration of the scheduling process.

pair and κ denotes the index of the prior most recent interval over which that pair was scheduled. We let $\Gamma[k]$ be the collection of the most recently obtained finer channel estimates at the BS for each of the user pairs and their corresponding interval indices, at the start of interval k . Thus, the set $\Gamma[k]$ takes the form $\Gamma[k] = \{(\check{\mathbf{h}}_i[\kappa_{i,j}], \check{\mathbf{h}}_j[\kappa_{i,j}], \kappa_{i,j}), 1 \leq i < j \leq N\}$, where $(\check{\mathbf{h}}_i[\kappa_{i,j}], \check{\mathbf{h}}_j[\kappa_{i,j}])$ denote the finer estimates for interval $\kappa_{i,j}$ and $\kappa_{i,j}$ denotes the index of the prior recent-most interval on which pair i, j was scheduled. At the end of that interval (equivalently at the start of interval $k+1$) the set $\Gamma[k+1]$ is obtained by first setting it equal to $\Gamma[k]$ and then updating the 3-tuple corresponding to the pair (u_1, u_2) selected in interval k to $(\check{\mathbf{h}}_{u_1}[k], \check{\mathbf{h}}_{u_2}[k], k)$.

The set of user channel states are assumed to be i.i.d. across intervals and the channel states of any two distinct users are assumed to be mutually independent. Given a particular initial rough estimates of the channel states of the user pair selected in interval k , $(\hat{\mathbf{h}}_{u_1}[k], \hat{\mathbf{h}}_{u_2}[k])$, the distribution of the finer channel estimates in the same interval is described by the conditional distribution

$$P(\check{\mathbf{h}}_{u_1}[k], \check{\mathbf{h}}_{u_2}[k] | \hat{\mathbf{h}}_{u_1}[k], \hat{\mathbf{h}}_{u_2}[k]) \quad (1)$$

where the conditional probability depends on the types of channel estimators, quantization, training times and powers, etc. We let $\mathcal{C}_{\text{coarse}}$ ($\mathcal{C}_{\text{fine}}$) denote the finite sets or codebooks of vectors from which all coarse (fine) estimates are selected. Let $|\mathcal{C}_{\text{coarse}}|$ and $|\mathcal{C}_{\text{fine}}|$ denote their respective cardinalities and clearly $|\mathcal{C}_{\text{fine}}| \geq |\mathcal{C}_{\text{coarse}}|$.

C. Expected Transmission Rates (Rewards)

During the current interval k , formed by slots $[k, 1], [k, 2]$ & $[k, 3]$, once a pair of users is selected, the scheduler specifies transmit precoding matrices or vectors for each slot in the interval.

1) *Slot 1*: For slot 1, the overall transmit precoding matrix is denoted by the matrix $[\mathbf{W}_{u_1}[k], \mathbf{W}_{u_2}[k]]$, where $\mathbf{W}_{u_1}[k], \mathbf{W}_{u_2}[k] \in \mathbb{C}^{M_t \times 2}$. Let $\mathbf{x}_{u_1}[k] = \mathbf{W}_{u_1}[k] \mathbf{s}_{u_1}[k]$, $\mathbf{x}_{u_2}[k] = \mathbf{W}_{u_2}[k] \mathbf{s}_{u_2}[k]$, where $\mathbf{s}_{u_1}[k], \mathbf{s}_{u_2}[k]$ denote the 2×1 symbol vectors containing symbols formed using the new packets intended for user u_1 and u_2 , respectively, and where $E[\mathbf{s}_{u_i}[k] \mathbf{s}_{u_i}^\dagger[k]] = \mathbf{I}$, $i \in \{1, 2\}$. Then, the signal transmitted in slot-1 is $\mathbf{x}_{u_1}[k] + \mathbf{x}_{u_2}[k]$

so that the received signals at both users are

$$y_{u_1}[k, 1] = \mathbf{h}_{u_1}[k] (\mathbf{x}_{u_1}[k] + \mathbf{x}_{u_2}[k]) + n_{u_1}[k, 1], \quad (2)$$

$$y_{u_2}[k, 1] = \mathbf{h}_{u_2}[k] (\mathbf{x}_{u_1}[k] + \mathbf{x}_{u_2}[k]) + n_{u_2}[k, 1]. \quad (3)$$

Note that the allocated transmission power for scheduled user u_i is the norm $\|\mathbf{W}_{u_i}[k]\|^2$. We assume that the maximum average (per-slot) transmission power budget at the BS is P . Thus, the corresponding power constraint is $\|\mathbf{W}_{u_1}[k]\|^2 + \|\mathbf{W}_{u_2}[k]\|^2 \leq P$. Notice that the precoding matrix $[\mathbf{W}_{u_1}[k], \mathbf{W}_{u_2}[k]]$ seeks to facilitate the transmission of new packets to users u_1 and u_2 and thus must be designed based on the available coarse estimates $(\hat{\mathbf{h}}_{u_1}[k], \hat{\mathbf{h}}_{u_2}[k])$, since the corresponding finer estimates for that interval are not yet available to the scheduler. Accordingly, we assume that this precoding matrix can be obtained as the output of any arbitrary but fixed (time-invariant) mapping from $\mathcal{C}_{\text{coarse}} \times \mathcal{C}_{\text{coarse}}$ to $\mathbb{C}^{M_t \times 4}$, when the coarse estimates $(\hat{\mathbf{h}}_{u_1}[k], \hat{\mathbf{h}}_{u_2}[k])$ are given as an input. Note that assuming the mapping to be fixed is well suited to systems where the so-called ‘‘precoded pilots’’ are not available so that the choice of precoders needs to be signalled to the scheduled users. A fixed mapping (which is equivalent to one codebook of transmit precoders) then allows for efficient signaling.

2) *Slot 2*: In slot 2 of the interval, an interference resolving packet for a pending previous transmission involving users (u_1, u_2) , sent in interval $\kappa < k$, is transmitted. In particular, the transmitted signal vector over the M_t antennas is

$$z[k, 2] (\check{\mathbf{h}}_{u_1}[\kappa] \underbrace{\mathbf{W}_{u_2}[\kappa] \mathbf{s}_{u_2}[\kappa]}_{\mathbf{x}_{u_2}[\kappa]}),$$

where $z[k, 2] \in \mathbb{C}^{M_t \times 1}$ is a precoding vector. Note that $\check{\mathbf{h}}_{u_1}[\kappa] \mathbf{x}_{u_2}[\kappa]$ is a scalar, so the average power constraint $E[\|z[k, 2] \check{\mathbf{h}}_{u_1}[\kappa] \mathbf{x}_{u_2}[\kappa]\|^2] \leq P$ can also be written as $\|z[k, 2]\|^2 \|\check{\mathbf{h}}_{u_1}[\kappa] \mathbf{W}_{u_2}[\kappa]\|^2 \leq P$. The received signals in slot 2 at both users are therefore

$$y_{u_1}[k, 2] = \mathbf{h}_{u_1}[k] z[k, 2] (\check{\mathbf{h}}_{u_1}[\kappa] \mathbf{x}_{u_2}[\kappa]) + n_{u_1}[k, 2] \quad (4)$$

$$y_{u_2}[k, 2] = \mathbf{h}_{u_2}[k] z[k, 2] (\check{\mathbf{h}}_{u_1}[\kappa] \mathbf{x}_{u_2}[\kappa]) + n_{u_2}[k, 2]. \quad (5)$$

3) *Slot 3*: In slot 3 of the interval, similarly, the transmitted signal is

$$z[k, 3] (\check{\mathbf{h}}_{u_2}[\kappa] \underbrace{\mathbf{W}_{u_1}[\kappa] \mathbf{s}_{u_1}[\kappa]}_{\mathbf{x}_{u_1}[\kappa]}).$$

so that the power constraint is $\|z[k, 3]\|^2 \|\check{\mathbf{h}}_{u_2}[\kappa] \mathbf{W}_{u_1}[\kappa]\|^2 \leq P$. The received signals in slot 3 at both users are therefore

$$y_{u_1}[k, 3] = \mathbf{h}_{u_1}[k] z[k, 3] (\check{\mathbf{h}}_{u_2}[\kappa] \mathbf{x}_{u_1}[\kappa]) + n_{u_1}[k, 3] \quad (6)$$

$$y_{u_2}[k, 3] = \mathbf{h}_{u_2}[k] z[k, 3] (\check{\mathbf{h}}_{u_2}[\kappa] \mathbf{x}_{u_1}[\kappa]) + n_{u_2}[k, 3]. \quad (7)$$

Notice that the precoding vectors $z[k, 2], z[k, 3]$ seek to facilitate the completion of a pending transmission

to users u_1 and u_2 and thus must be designed based on the available coarse estimates $(\hat{\mathbf{h}}_{u_1}[k], \hat{\mathbf{h}}_{u_2}[k])$, as well as the available estimates for interval κ which are $(\check{\mathbf{h}}_{u_1}[\kappa], \check{\mathbf{h}}_{u_2}[\kappa])$ and $(\hat{\mathbf{h}}_{u_1}[\kappa], \hat{\mathbf{h}}_{u_2}[\kappa])$. Accordingly, we assume that these two vectors can be obtained as the output of an arbitrary but fixed mapping from $\mathcal{C}_{\text{fine}}^2 \times \mathcal{C}_{\text{coarse}}^4$ to $\mathbb{C}^{M_t \times 2}$. An example of mapping rules to obtain the precoding matrices and vectors is given later in the section on simulation results.

Next, in order to compute the average rates (rewards) we assume that the channel state vectors $\mathbf{h}_{u_i}[k]$, $\mathbf{h}_{u_i}[k]$ are known perfectly to user u_i , $i \in \{1, 2\}$ (each user of course also knows the quantized estimates it has fed back to the base-station). In addition, user u_1 (u_2) is also conveyed the finer estimate $\check{\mathbf{h}}_{u_2}[\kappa]$, $(\check{\mathbf{h}}_{u_1}[\kappa])$ via feed-forward signaling before the start of interval k . For simplicity, the feedback and feedforward signaling overheads are ignored in this work. Then, by the end of slot 3, from (2), (4) and (6), at user u_1 , we have

$$y_{u_1}[k, 1] - \frac{y_{u_1}[k, 2]}{\mathbf{h}_{u_1}[k]\mathbf{z}[k, 2]} = \mathbf{h}_{u_1}[\kappa]\mathbf{x}_{u_1}[\kappa] + (\mathbf{h}_{u_1}[\kappa] - \check{\mathbf{h}}_{u_1}[\kappa])\mathbf{x}_{u_2}[\kappa] + n_{u_1}[k, 1] - \frac{n_{u_1}[k, 2]}{\mathbf{h}_{u_1}[k]\mathbf{z}[k, 2]},$$

$$y_{u_1}[k, 3] = \underbrace{(\mathbf{h}_{u_1}[k]\mathbf{z}[k, 3])}_{\delta_{u_1}[k]} \check{\mathbf{h}}_{u_2}[\kappa]\mathbf{x}_{u_1}[\kappa] + n_{u_1}[k, 3], \quad (8)$$

where the additive noise variables $n_{u_1}[k, 1], n_{u_1}[k, 2], n_{u_1}[k, 3]$ are i.i.d. circularly symmetric complex Gaussian variables with zero-mean and unit variance, $\mathcal{CN}(0, 1)$. Notice that the interference term $(\mathbf{h}_{u_1}[\kappa] - \check{\mathbf{h}}_{u_1}[\kappa])\mathbf{x}_{u_2}[\kappa]$ is independent of the desired signal as well as the additive noise. Letting $\mathbf{h}_{u_1}^{\text{error}}[\kappa] = \mathbf{h}_{u_1}[\kappa] - \check{\mathbf{h}}_{u_1}[\kappa]$, the noise plus interference covariance for user u_1 , denoted by $\mathbf{\Gamma}_{u_1}[k]$, is therefore

$$\begin{bmatrix} 1 + \|\mathbf{h}_{u_1}^{\text{error}}[\kappa]\mathbf{W}_{u_2}[\kappa]\|^2 + \frac{1}{|\mathbf{h}_{u_1}[k]\mathbf{z}[k, 2]|^2} & 0 \\ 0 & 1 \end{bmatrix}.$$

Define $\mathbf{G}_{u_1}[k] = [\mathbf{h}_{u_1}[\kappa]\mathbf{W}_{u_1}[\kappa]; \delta_{u_1}[k]\check{\mathbf{h}}_{u_2}[\kappa]\mathbf{W}_{u_1}[\kappa]]$ and note that $\mathbf{G}_{u_1}[k] \in \mathbb{C}^{2 \times 2}$. Further, let $\mathbf{H}^{\text{csi}}((u_1, u_2), (\kappa, k)) = \{\check{\mathbf{h}}_{u_1}[\kappa], \check{\mathbf{h}}_{u_2}[\kappa], \hat{\mathbf{h}}_{u_1}[\kappa], \hat{\mathbf{h}}_{u_2}[\kappa], \check{\mathbf{h}}_{u_1}[k], \check{\mathbf{h}}_{u_2}[k]\}$ denote the set of channel state information at the scheduler for user pair u_1, u_2 over intervals κ, k . Then, using (8) the instantaneous information rate, denoted as $I_{u_1}[k]$ is given by

$$I_{u_1}[k] = \frac{1}{3} \log |\mathbf{I} + \mathbf{\Gamma}_{u_1}^{-1}[k]\mathbf{G}_{u_1}[k]\mathbf{G}_{u_1}^\dagger[k]|, \quad (9)$$

where the fraction $1/3$ is to account for the fact that three slots are needed to obtain this rate. Then, (an optimistic value for) the average information rate that can be achieved via rateless coding (cf. [9]) is given by

$$R_{u_1}^{\text{opt}}[k] = E [I_{u_1}[k] | \mathbf{H}^{\text{csi}}((u_1, u_2), (\kappa, k))]. \quad (10)$$

A more conservative rate that is appropriate for conventional coding, denoted as $R_{u_1}^{\text{conv}}[k]$, is given by

$$r_{\underline{\theta}, u_1} \left(1 - \Pr \left(I_{u_1}[k] < r_{\underline{\theta}, u_1} \mid \mathbf{H}^{\text{csi}}((u_1, u_2), (\kappa, k)) \right) \right), \quad (11)$$

where $r_{\underline{\theta}, u_1}$ denotes the rate assigned (using any fixed mapping) to user u_1 in $\underline{\theta}$ before transmission of new packets for the pair (u_1, u_2) in interval κ , based on the available coarse estimates $\hat{\mathbf{h}}_{u_1}[\kappa], \hat{\mathbf{h}}_{u_2}[\kappa]$. The rates corresponding to (10) or (11) can be derived in a similar manner for user u_2 .

Note that in deriving the average rate in (10) or (11) we have assumed a simple albeit sub-optimal filtering at the user to suppress the interference from the transmission intended for the co-scheduled user. For completeness, we provide the average rate expressions for the case when the user employs the optimal linear filter and for brevity we only consider the optimistic rate for user u_1 . Towards this end, we collect the observations received by user u_1 as

$$\begin{bmatrix} y_{u_1}[k, 1] \\ y_{u_1}[k, 2] \\ y_{u_1}[k, 3] \end{bmatrix} = \mathbf{F}_{u_1}[k]\mathbf{x}_{u_1}[\kappa] + \tilde{\mathbf{F}}_{u_1}[k]\mathbf{x}_{u_2}[\kappa] + \begin{bmatrix} n_{u_1}[k, 1] \\ n_{u_1}[k, 2] \\ n_{u_1}[k, 3] \end{bmatrix},$$

where

$$\mathbf{F}_{u_1}[k] = \begin{bmatrix} \mathbf{h}_{u_1}[\kappa] \\ \mathbf{0} \\ \delta_{u_1}[k]\check{\mathbf{h}}_{u_2}[\kappa] \end{bmatrix}, \quad \tilde{\mathbf{F}}_{u_1}[k] = \begin{bmatrix} \mathbf{h}_{u_1}[\kappa] \\ \mathbf{h}_{u_1}[k]\mathbf{z}[k, 2]\check{\mathbf{h}}_{u_1}[\kappa] \\ \mathbf{0} \end{bmatrix}$$

For this model, we can determine the instantaneous information rate that can be achieved via optimal filtering using (9) but where where $\mathbf{\Gamma}_{u_1}[k] = \mathbf{I} + \tilde{\mathbf{F}}_{u_1}[k]\mathbf{W}_{u_2}[\kappa]\mathbf{W}_{u_2}^\dagger[\kappa]\tilde{\mathbf{F}}_{u_1}^\dagger[k]$ and $\mathbf{G}_{u_1}[k] = \mathbf{F}_{u_1}[k]\mathbf{W}_{u_1}[\kappa]$. The average information rate can then determined as before using (10).

We assume that either conventional coding is employed for all users or rateless coding is employed and accordingly let $R_{u_i}[k]$, $1 \leq i \leq 2$ denote the average rate, henceforth referred to also as the service rate, obtained over interval k . We also note here that the scheduling scheme (policy) is preceded by an initial set-up phase comprising of $N(N-1)/2$ intervals in which new packets are transmitted successively to each user pair without any accompanying interference resolution packets. For notational convenience, we assume that the scheduling policy starts operating from interval with index 0 using the initial set $\Gamma[0]$ determined by the set-up phase.

D. Incorporating one-shot transmissions and feedback delays

We first consider the case of one-shot transmissions. To enable one-shot transmission of packets to any pair in any interval k , we define an action $\underline{\theta}$ in which u_1, u_2 is the pair but $\kappa = \phi$ to capture the fact that the intended transmission is one-shot and hence does not seek to resolve any pending previous transmission.

Then, in all three slots of that interval transmission is done as in conventional MU-MIMO relying only on the available current estimates $\hat{H}[k]$. In particular, a transmit precoder $[\mathbf{w}_{u_1}[k], \mathbf{w}_{u_2}[k]] \in \mathbb{C}^{M_t \times 2}$ is formed based on $\{\hat{\mathbf{h}}_{u_1}[k], \hat{\mathbf{h}}_{u_2}[k]\}$ using a technique such as zero-forcing [8]. Defining $I_{u_1}^{\text{one-shot}}[k] = \log(1 + |\hat{\mathbf{h}}_{u_1}[k] \mathbf{w}_{u_1}[k]|^2 / (1 + |\hat{\mathbf{h}}_{u_2}[k] \mathbf{w}_{u_2}[k]|^2))$, the corresponding average rates obtained for user u_1 (similarly for user u_2) are given by

$$E \left[I_{u_1}^{\text{one-shot}}[k] \mid \hat{\mathbf{h}}_{u_1}[k], \hat{\mathbf{h}}_{u_2}[k] \right], \quad (12)$$

or

$$r_{\underline{\theta}, u_1} \left(1 - \Pr \left(I_{u_1}^{\text{one-shot}}[k] < r_{\underline{\theta}, u_1} \mid \hat{\mathbf{h}}_{u_1}[k], \hat{\mathbf{h}}_{u_2}[k] \right) \right).$$

In addition at the end of interval k , we simply set $\Gamma[k+1] = \Gamma[k]$ since no pending packets are completed or introduced.

Recall that so far we have assumed that upon choosing action $\underline{\theta}$ for interval k , the finer estimates $\hat{\mathbf{h}}_{u_1}[k], \hat{\mathbf{h}}_{u_2}[k]$ are available at the start of interval $k+1$ (representing a unit delay). In practical systems there can be a delay of several intervals in obtaining such finer estimates. Assuming that these delays are fixed and known in advance, they can be accommodated by expanding the definition of a state. In particular, we can define 4-tuples such as $(i, j, \kappa_{i,j}, d_{i,j})$ where $d_{i,j} \geq 0$ measures the remaining delay after which finer estimates $\hat{\mathbf{h}}_i[\kappa_{i,j}], \hat{\mathbf{h}}_j[\kappa_{i,j}]$ will be available. At any interval k selecting the action $(i, j, \kappa_{i,j}, d_{i,j})$ with $d_{i,j} > 0$ ($d_{i,j} = 0$) constrains the interference resolution to be based only on the coarse estimates $\hat{\mathbf{h}}_i[\kappa_{i,j}], \hat{\mathbf{h}}_j[\kappa_{i,j}], \hat{\mathbf{h}}_i[k], \hat{\mathbf{h}}_j[k]$ (on both coarse and fine estimates $\mathbf{H}^{\text{csi}}((i, j), (\kappa_{i,j}, k))$). Upon selecting this action the 4-tuple in $\Gamma[k+1]$ corresponding to the pair i, j is set to be $(i, j, k, d_{i,j} = D_{i,j})$ where $D_{i,j}$ is the maximum delay (starting from $k+1$) after which the finer estimates will be available. If that action is not selected, it is updated in $\Gamma[k+1]$ as $(i, j, \kappa_{i,j}, d_{i,j} = \max\{0, d_{i,j} - 1\})$. For convenience in exposition the aforementioned two extensions are not considered below.

E. System State and Throughput Region

Define the system state at the start of an interval j as $S[j] = \{\Gamma[j], \hat{H}[j]\}$ and let $\underline{\theta}[j]$ denote the decision (action) taken in that interval. Then, at each interval k , a scheduling policy ψ takes as input all the history up to interval k , comprising of states $\{S[j]\}_{j=0}^k$ and all decisions $\{\underline{\theta}[j]\}_{j=0}^{k-1}$, to output a decision $\underline{\theta}[k]$. Under a particular policy ψ , the throughput of the n^{th} user is denoted as

$$r_n^\psi = \lim_{J \rightarrow \infty} \frac{1}{J} \sum_{t=0}^{J-1} E[R_n^\psi[t]] \quad \forall n, \quad (13)$$

where $R_n^\psi[t] = R_n[t] \mathbf{1}(n \in \underline{\theta}[t])$ and the expectation is over the initial state and the evolution of the states and

decisions in the subsequent intervals. Note that in (13) for simplicity we have assumed that the limit exists for the selected policy. In case the limit does not exist, we can consider any sub-sequence for which the limit exists. Let Ψ be the set of all policies. The throughput region that is of interest to us is defined as the closure of the convex hull of the throughput vectors achievable under all policies in Ψ , i.e.,

$$\Lambda = CH\{\mathbf{r} : \exists \psi \in \Psi \text{ s.t.}, \mathbf{r} = \mathbf{r}^\psi\},$$

where $CH\{\cdot\}$ denotes closure of the convex hull. For each throughput vector \mathbf{r} , we obtain a utility value $U(\mathbf{r})$, where $U(\cdot)$ is the non-negative component-wise non-decreasing and concave utility function. For convenience, we also assume that the utility is continuous (and hence uniformly continuous) in the closed hypercube $[0, b]^N$ for each finite $b \in \mathbb{R}_+$. The objective then is to maximize the network utility within the throughput region, i.e., $\max_{\mathbf{r}: \mathbf{r} \in \Lambda} U(\mathbf{r})$.

III. OPTIMAL FRAME-BASED SCHEDULING POLICY

In this section, we propose a frame based policy that achieves a utility arbitrarily close to the optimal. In this policy, the time intervals are further grouped into separate frames, where each frame consists of T consecutive intervals. The scheduling decisions in each frame are based on a set of virtual queues that guide the achieved system utility towards optimal, as specified next.

A. Virtual Queue and Virtual Arrival Process

To control the achieved utilities of different users, a virtual queue is maintained for each user, denoted as $Q_n[k], k = 0, 1, \dots$ & $n = 1, \dots, N$. At the beginning of the τ^{th} frame comprising of intervals $\{\tau T, \dots, (\tau+1)T - 1\}$, where $\tau \in \{0, 1, 2, \dots\}$, the following optimization problem is solved at the scheduler

$$\max_{\mathbf{r}: \mathbf{0} \leq \mathbf{r} \leq \mathbf{r}_{\max}} V \cdot U(\mathbf{r}) - \sum_{n=1}^N Q_n[\tau T] r_n, \quad (14)$$

where r_{\max}, V are positive constants that can be freely chosen and whose role will be revealed later. We let $\mathbf{r}^*[\tau]$ be the optimal solution to the above problem. Then, the virtual arrival rate for user n is set as $r_n^*[\tau]$ in each interval in the τ^{th} frame. A scheduling policy, $\psi_{Q[\tau T]}^*$, is determined and implemented based on the virtual queue length $Q[\tau T]$ obtained at the beginning of that frame. Letting $R_n^{\Psi_{Q[\tau T]}}[k]$ denote the service rate of user n in each interval k in the τ^{th} frame under this policy, the virtual queue is then updated as

$$Q_n[k+1] = \left(Q_n[k] - R_n^{\Psi_{Q[\tau T]}}[k] \right)^+ + r_n^*[\tau], \quad (15)$$

for all $\tau T \leq k \leq (\tau+1)T - 1$ and each user n and where $(x)^+ = \max\{0, x\}$ with $Q_n[0] = 0$ for all n .

B. State-action frequency approach

We now determine the policy $\Psi_{\mathbf{Q}[\tau T]}^*$ employed in the τ^{th} frame. Notice that while the definition of the system state adopted thus far allows us to compactly describe any policy, one associated drawback is that the number of states becomes countably infinite. Fortunately, there is one aspect that we can exploit. Note that the average rates obtained upon scheduling a pair of users i, j on any interval k depends only on the corresponding coarse and fine channel estimates in interval $\kappa_{i,j}$ (which we recall denotes the prior recent-most interval over which that pair was scheduled) and the coarse channel estimates in interval k but not on those interval indices. Then, to analyze the average rates offered by any policy, it suffices to define a finite set of states, $\underline{\mathcal{S}}$, as follows. A state $\underline{s} \in \underline{\mathcal{S}}$ is defined as a particular choice $\mathbf{h}_i^{p,\text{fine}}, \mathbf{h}_j^{p,\text{fine}}, \mathbf{h}_i^{p,\text{coarse}}, \mathbf{h}_j^{p,\text{coarse}}, \mathbf{h}_i^{c,\text{coarse}}, \mathbf{h}_j^{c,\text{coarse}}$ of coarse and fine channel estimates for each pair i, j , where the superscripts p, c denote past and current estimates, respectively. Consequently there are $|\underline{\mathcal{S}}| = (|\mathcal{C}_{\text{fine}}|^2 |\mathcal{C}_{\text{coarse}}|^2)^{\frac{N(N-1)}{2}} |\mathcal{C}_{\text{coarse}}|^N$ number of states. Note that a state $S[k]$ in the previous definition would map to state $\underline{s} \in \underline{\mathcal{S}}$ which has the choice $\hat{\mathbf{h}}_i[\kappa_{i,j}], \hat{\mathbf{h}}_j[\kappa_{i,j}], \hat{\mathbf{h}}_i[\kappa_{i,j}], \hat{\mathbf{h}}_j[\kappa_{i,j}], \hat{\mathbf{h}}_i[k], \hat{\mathbf{h}}_j[k]$ for each pair i, j . A finite set of actions, $\underline{\mathcal{A}}$, is defined next to be the collection of all possible user pairs so that any $\underline{a} \in \underline{\mathcal{A}}$ uniquely identifies a user pair. Let $P(\underline{s}|\underline{s}', \underline{a})$ denote the transition probability, which we note can be determined using (1) and the facts that the finer past estimates of pairs not in \underline{a} do not change and the current coarse estimates are i.i.d. across intervals. Letting $\underline{\mathcal{P}}(\underline{\mathcal{A}})$ define the set of all probability distributions on $\underline{\mathcal{A}}$, any policy can be defined as a mapping which at each interval k takes as input all the history up-to interval k , comprising of states $\{\underline{s}[j]\}_{j=0}^k$ and all actions $\{\underline{a}[j]\}_{j=0}^{k-1}$, to output a distribution in $\underline{\mathcal{P}}(\underline{\mathcal{A}})$ from which the action $\underline{a}[k]$ can be generated. A stationary policy is one which at any interval k considers only the state $\underline{s}[k]$ to output a distribution in $\underline{\mathcal{P}}(\underline{\mathcal{A}})$ and where the output distribution depends only on the state $\underline{s}[k]$ but not on the interval index k . Under any stationary policy the sequence $\{\underline{s}[k]\}_{k=0}^\infty$ is a Markov Chain.

With these definitions in hand, we let $R_n(\underline{s}, \underline{a})$ denote the achieved transmission rate for user n when action \underline{a} is taken and the system state is \underline{s} . Denote the state action frequencies by $\{x(\underline{s}, \underline{a})\}_{\underline{s} \in \underline{\mathcal{S}}, \underline{a} \in \underline{\mathcal{A}}}$, where we note that each $x(\underline{s}, \underline{a})$ lies in the unit interval $[0, 1]$ and represents the frequency that the system state is at \underline{s} and action \underline{a} is taken. The state action frequencies need to satisfy the normalization equation

$$\sum_{\underline{s}, \underline{a}} x(\underline{s}, \underline{a}) = 1,$$

and the balance equation

$$\sum_{\underline{a}} x(\underline{s}, \underline{a}) = \sum_{\underline{s}', \underline{a}} P(\underline{s}|\underline{s}', \underline{a}) x(\underline{s}', \underline{a}).$$

The above two equations form a state-action polytope $\underline{\mathcal{X}}$ and let \mathbf{x} denote any vector of state action frequencies lying in $\underline{\mathcal{X}}$. We next define a rate region as

$$\tilde{\Lambda} = \{\mathbf{R} : R_n = \sum_{\underline{s}} \sum_{\underline{a}} R_n(\underline{s}, \underline{a}) x(\underline{s}, \underline{a}), \forall n \ \& \ \mathbf{x} \in \underline{\mathcal{X}}\}. \quad (16)$$

Then, given the virtual queue length $\mathbf{q} = \mathbf{Q}[\tau T]$ we consider the following linear program (LP),

$$\begin{aligned} \max_{\mathbf{x}} \quad & \sum_{\underline{s}, \underline{a}} \mathbf{q}^T \mathbf{R}(\underline{s}, \underline{a}) x(\underline{s}, \underline{a}) \\ \text{s.t.} \quad & \mathbf{x} \in \underline{\mathcal{X}}. \end{aligned} \quad (17)$$

We use \mathbf{x}^* to denote an optimal solution to the linear program and define $\mathbf{R}^* = [R_1^*, \dots, R_N^*]^T$, where

$$R_n^* = \sum_{\underline{s}} \sum_{\underline{a}} R_n(\underline{s}, \underline{a}) x^*(\underline{s}, \underline{a}), \forall n. \quad (18)$$

Using the Bayesian rule, we can identify the corresponding stationary policy $\Psi_{\mathbf{Q}[\tau T]}^*$, which at any interval k in the τ^{th} frame first maps the state $S[k]$ to its counterpart $\underline{s} \in \underline{\mathcal{S}}$. Then, if $\sum_{\underline{a}'} x^*(\underline{s}, \underline{a}') > 0$, it chooses action \underline{a} using the probabilistic rule

$$P(\text{pick } \underline{a} \text{ at state } \underline{s}) = \frac{x^*(\underline{s}, \underline{a})}{\sum_{\underline{a}'} x^*(\underline{s}, \underline{a}')}, \forall \underline{a} \in \underline{\mathcal{A}}.$$

On the other hand, if $\sum_{\underline{a}'} x^*(\underline{s}, \underline{a}') = 0$, it chooses action \underline{a} arbitrarily. Let $\mathbf{R}^{\text{frame}[k]}, \tau T \leq k \leq (\tau+1)T-1$, denote the service rate vectors obtained under this policy for the intervals in the τ^{th} frame.

We list the following results which can be obtained using those that have been derived before for weakly communicating Markov Decision Processes [10],[11].

Lemma 1. *The region Λ defined in (13) is identical to the region $\tilde{\Lambda}$ defined in (16). Further, for each frame τ and any given $\mathbf{Q}[\tau T]$, an optimal solution to the LP in (17) can be found for which the corresponding policy $\Psi_{\mathbf{Q}[\tau T]}^*$ is also deterministic.*

Henceforth, we assume $\Psi_{\mathbf{Q}[\tau T]}^*$ to be also deterministic.

Lemma 2. *For arbitrarily fixed $\delta > 0$ there exists a large enough frame length T_o and constants γ, β such that for each frame length $T \geq T_o$ and all $\mathbf{Q}[\tau T]$*

$$\begin{aligned} \Pr \left(\left\| \frac{1}{T} \left(\sum_{j=0}^{T-1} \mathbf{R}^{\text{frame}[\tau T + j]} \right) - \mathbf{R}^* \right\| > \delta \mid \mathbf{Q}[\tau T] \right) \\ \leq \gamma \exp(-\beta T). \end{aligned} \quad (19)$$

C. Optimality of the frame-based policy

Define Lyapunov function $L(\mathbf{Q}[\tau T]) = \frac{1}{2} \sum_{n=1}^N Q_n^2[\tau T]$. Then the T -step average Lyapunov drift is expressed as

$$\Delta_T(\mathbf{Q}[\tau T]) = \frac{1}{T} E [L(\mathbf{Q}[(\tau + 1)T]) - L(\mathbf{Q}[\tau T]) \mid \mathbf{Q}[\tau T]],$$

where the expectation is over the initial states at interval τT induced by the policies adopted in the previous frames and the evolution of the states and decisions in the τ^{th} frame under the policy $\Psi_{\mathbf{Q}[\tau T]}^*$. Our first result is the following, the proof of which is included in [13].

Proposition 1. *For any given $\epsilon > 0$, there exists a frame length T_o such that for all frame lengths $T \geq T_o$ the T -step average Lyapunov drift can be bounded as*

$$\Delta_T(\mathbf{Q}[\tau T]) \leq BT - \sum_{n=1}^N Q_n[\tau T] R_n + \sum_{n=1}^N Q_n[\tau T] r_n^*[\tau], \quad (20)$$

where B is a constant and $\mathbf{R} = [R_1, \dots, R_N]^T$ is any vector such that $\mathbf{R} + \epsilon \mathbf{1} \in \Lambda$.

Consider the ϵ -interior of Λ , i.e., $\Lambda_\epsilon = \{\mathbf{R} : \mathbf{R} + \epsilon \mathbf{1} \in \Lambda\}$. Denote $\mathbf{r}_\epsilon^{\text{opt}}$ as the optimal value of the following optimization problem.

$$\begin{aligned} & \max U(\mathbf{r}) \\ & \text{s.t. } \mathbf{r} \in \Lambda_\epsilon; \mathbf{r} \preceq r_{\max} \mathbf{1}. \end{aligned}$$

Our main result is the following which is proved in [13].

Theorem 1. *For any given $\epsilon > 0$, there exists a T_o such that for all frame lengths $T \geq T_o$*

$$\liminf_{J \rightarrow \infty} U \left(\frac{1}{J} \sum_{t=0}^{J-1} E [\mathbf{R}^{\text{frame}}[t]] \right) \geq U(\mathbf{r}_\epsilon^{\text{opt}}) - BT/V.$$

Thus, by choosing ϵ , framelength T and parameters V, r_{\max} appropriately, our frame based policy can be made arbitrarily close to optimal.

For comparison we will use the conventional MU-MIMO scheduling described in Section II-A. In addition, we also use the following *myopic* policy. This policy operates in a manner similar to the frame based policy but with the following important differences. Firstly, the frame-length is set as $T = 1$ so that the arrival rates are computed at the start of each interval and the virtual queues are updated at the end of that interval. Then, at each interval k the current state $S[k]$ is mapped to its image $\underline{s} \in \underline{\mathcal{S}}$. Considering the queue length $\mathbf{q} = \mathbf{Q}[k]$, the action $\hat{\underline{a}} = \arg \max_{\underline{a} \in \underline{\mathcal{A}}} \mathbf{q}^T \mathbf{R}(\underline{s}, \underline{a})$ is selected. Clearly, this policy does not consider the transition probabilities (and the possible future evolutions) at all while deciding an action. Nevertheless, as seen in the following section, this policy indeed offers a competitive performance.

IV. SIMULATION RESULTS

We consider a narrowband downlink with four single-antenna users that are served by a BS equipped with four transmit antennas. All users are assumed to experience an identical (large scale fading) pathloss factor δ and thus see an identical average SNR, which models the physical scenario in which all users are equidistant from the BS. Further, we model the small-scale fading seen by each user as Rayleigh fading so the channel response vector of each user is assumed to have i.i.d. $\mathcal{CN}(0, \delta^2)$ elements. Consequently the normalized channel response vector (i.e., channel direction) is isotropically distributed in $\mathbb{C}^{4 \times 1}$. Moreover, the channel response vectors evolve independently across intervals and are independent across users. In the following simulations, each user quantizes its channel norm and channel direction separately. In particular, the channel norm is quantized using a scalar quantizer which for simplicity we assume to be identical for both fine and coarse estimates. On the other hand, to quantize the channel direction, in order to obtain the finer estimate, the quantization codebook used comprises of a set independently generated instances of isotropic vectors in $\mathbb{C}^{4 \times 1}$ (a.k.a. random vector codebook), where we note that for large codebook sizes random vector codebooks have been shown to be a good choice for both SU-MIMO and conventional MU-MIMO. The quantization of the channel direction to obtain the coarser estimate is accomplished using Grassmannian codebooks.

Before offering our results, we consider an interval k and decision $\underline{\theta}$ and describe the mapping rules alluded to in Section II-C. We determine a good direction (i.e., unit-norm beamforming vector) for multicasting using the alternating optimization based multicast beamforming design algorithm [12] that takes only the coarse estimates $\hat{\mathbf{h}}_{u_1}[k]$ and $\hat{\mathbf{h}}_{u_2}[k]$ as inputs and set $\frac{\mathbf{z}[k,2]}{\|\mathbf{z}[k,2]\|}$ and $\frac{\mathbf{z}[k,3]}{\|\mathbf{z}[k,3]\|}$ to be equal to this direction. The precoding matrix $\mathbf{W}_{u_1}[\kappa]$ is obtained by extending the naive zero-forcing design of conventional MU-MIMO to the model in (8). In particular at interval κ the BS naively assumes that coarse estimates $\hat{\mathbf{h}}_{u_1}[\kappa], \hat{\mathbf{h}}_{u_2}[\kappa]$ it has are indeed equal to their respective exact channels (and hence their respective finer estimates). Then, at any future interval k (the knowledge of k is not assumed during interval κ) when pair (u_1, u_2) is next scheduled, under the naive assumption (8) would reduce to

$$\begin{aligned} y_{u_1}[\kappa, 1] - \frac{y_{u_1}[k, 2]}{\mathbf{h}_{u_1}[k] \mathbf{z}[k, 2]} &= \hat{\mathbf{h}}_{u_1}[\kappa] \mathbf{x}_{u_1}[\kappa] + n_{u_1}[\kappa, 1] \\ &\quad - \frac{n_{u_1}[k, 2]}{(\mathbf{h}_{u_1}[k] \mathbf{z}[k, 2])}, \\ \frac{y_{u_1}[k, 3]}{(\mathbf{h}_{u_1}[k] \mathbf{z}[k, 3])} &= \hat{\mathbf{h}}_{u_2}[\kappa] \mathbf{x}_{u_1}[\kappa] + \frac{n_{u_1}[k, 3]}{(\mathbf{h}_{u_1}[k] \mathbf{z}[k, 3])}. \quad (21) \end{aligned}$$

To remove dependence on k , all noise covariances are averaged so that (21) reduces to a point-to-point MIMO channel with channel




Fig. 2. Comparison with conventional MU-MIMO

matrix $[\hat{\mathbf{h}}_{u_1}[\kappa]; \hat{\mathbf{h}}_{u_2}[\kappa]]$ and noise covariance $\text{diag}\{1 + E[1/|\mathbf{h}_{u_1}[k]\mathbf{z}[k, 2]|^2], E[1/|\mathbf{h}_{u_1}[k]\mathbf{z}[k, 3]|^2]\}$. Notice however that due to the power constraints these expected values in turn depend on the choice of precoders $\mathbf{W}_{u_1}[\kappa], \mathbf{W}_{u_2}[\kappa]$. As a further simplification, we fix these expected values to be suitable scalars which are determined offline. The precoder $\mathbf{W}_{u_1}[\kappa]$ can now be obtained using the standard point-to-point MIMO precoder design algorithm [7]. The precoder $\mathbf{W}_{u_2}[\kappa]$ is computed in an analogous manner. Finally, the norms of the precoding vectors are fixed as $\|\mathbf{z}[k, 2]\| = \frac{\sqrt{P}}{\|\hat{\mathbf{h}}_{u_1}[\kappa]\mathbf{W}_{u_2}[\kappa]\|}$ and $\|\mathbf{z}[k, 3]\| = \frac{\sqrt{P}}{\|\hat{\mathbf{h}}_{u_2}[\kappa]\mathbf{W}_{u_1}[\kappa]\|}$.

In Fig. 2 we compare the sum rate utility obtained using conventional MU-MIMO that only uses the current CSI with that obtained using the myopic scheduling that uses only the delayed CSI (EMAT with delayed) and the myopic scheduling that uses the hybrid CSI (EMAT with hybrid), where for the latter two schemes the average rates are computed assuming both the sub-optimal and the optimal filtering. In all cases the channel norms were assumed to be perfectly quantized whereas a 2-bit coarse codebook and 5-bit fine codebook were employed to quantize the channel directions, respectively. As seen from the figure, the conventional MU-MIMO gets interference limited and the policy using the finer albeit delayed CSI offers significant gains, which are further improved by utilizing the hybrid CSI. The improvement is more marked upon using optimal filtering. We also compared the sum rates obtained using our proposed policy and the myopic one, respectively, for a simpler examples having fewer number of states. We found that for well designed quantization codebooks, the myopic policy performs very close to the optimal frame based policy. This observation coupled with the fact that the complexity of the myopic policy scales much more benignly with the system size, makes it well suited to practical implementation. Additional details and results

are included in [13].

V. CONCLUSIONS

We considered the DL MU-MIMO scheduling problem with hybrid CSIT and proposed an optimal frame-based joint scheduling and feedback approach. There are two important and interesting issues that are the focus of our current research. The foremost one pertains to the exceedingly large number of states that are needed to accommodate practical system sizes which makes implementation of the frame based policy challenging even upon using commercial LP solvers. While the sparse nature of these linear programs can indeed be exploited, an efficient and significant reduction in the number states is necessary. The second issue is the choice of the precoding matrices and vectors. Recall that in this work we have assumed the choice of precoders to be pre-determined and fixed for each (state,action) pair. To fully exploit the precoding gains and the availability of “precoded pilots” in future networks, we should relax this restriction. Finally, we remark that incorporating practical considerations such as delay constraints on scheduling are other important open issues.

REFERENCES

- [1] D. Gesbert, M. Kountouris, R. W. Heath Jr, C. Chae and T. Salzer, “From single user to multiuser communications: shifting the MIMO paradigm,” *IEEE Signal Proc. Mag.*, Oct., 2007.
- [2] G. Caire, N. Jindal, M. Kobayashi and N. Ravindran, “Multiuser MIMO Achievable Rates With Downlink Training and Channel State Feedback,” *IEEE Transactions on Information Theory*, Jun., 2010.
- [3] A. Adhikary, H. C. Papadopoulos, S. A. Ramprasad and Giuseppe Caire, “Multi-User MIMO with outdated CSI: Training, Feedback and Scheduling,” *Allerton*, Oct., 2011.
- [4] H. Shirani-Mehr, G. Caire and M. Neely, “MIMO Downlink Scheduling with Non-Perfect Channel State Knowledge,” *IEEE Trans. on Comm.*, July, 2010.
- [5] M. A. Maddah-Ali and D. Tse, “Completely Stale Transmitter Channel State Information is Still Very Useful,” *IEEE Trans. on Information Theory*, July, 2012.
- [6] M. Kobayashi, S. Yang, D. Gesbert and Xiping Yi, “On the Degrees of Freedom of time correlated MISO broadcast channel with delayed CSIT,” *IEEE Trans. on Information Theory*, Jan., 2013.
- [7] D. Tse and P. Viswanath, “*Fundamentals of wireless communication*,” Cambridge University Press, 2005.
- [8] A. Wiesel, Y.C. Eldar and S. Shamai, “Zero-forcing precoding and generalized inverses,” *IEEE Trans. Signal Process.*, Sept. 2008.
- [9] 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. on Comm.*, Oct., 2011.
- [10] E. Altman, “*Constrained Markov Decision Processes*”, Chapman & Hall, 1999.
- [11] K. Jagannathan, S. Mannor, I. Menache and E. Modiano, “A state action frequency approach to throughput maximization over uncertain wireless channels,” *IEEE INFOCOM*, Shanghai, China, Apr. 2011.
- [12] H. Zhu, N. Prasad and S. Rangarajan, “Precoder Design for Physical Layer Multicasting,” *IEEE Trans. Sig. Proc.*, Nov. 2012.
- [13] W. Ouyang, N. Prasad and S. Rangarajan, “Exploiting Hybrid Channel Information for Downlink Multi-User MIMO Scheduling,” *Tech. Report*, arXiv, Mar. 2013.