Selection of the Coordination Strategy in the Network MIMO Downlink

Sebastian Stern and Robert F.H. Fischer
Institut für Nachrichtentechnik, Universität Ulm, Ulm, Germany
Email: {sebastian.stern, robert.fischer}@uni-ulm.de

***Abstract*— Network MIMO or distributed MIMO is one promising strategy to increase the data rates in future cellular communication systems. However, well-known schemes like linear preequalization or Tomlinson-Harashima precoding for handling the multiuser interference assume full coordination (via some centralized processing) between the base stations. This approach causes significant traffic in the (wired) backhaul. In this paper, decentralized strategies with partial coordination are considered. A selection algorithm is presented, which, given the actual MIMO channel, automatically chooses the amount of coordination which is necessary to achieve a predefined performance. The trade-off between the achievable reduction in backhaul signaling and the required signal-to-noise ratio to not exceed a tolerated bit error rate is discussed. The proposed strategies are covered by means of numerical simulations.**

I. Introduction

Over the last years, the principle of network multiple-input/multiple-output (network MIMO) has gained significant popularity in multiuser communication, e.g., [7]. Employing network MIMO, several user equipments (UEs) located within a service area are jointly supplied by a group of base stations (BSs). In our paper, we concentrate on the network MIMO downlink [9], i.e., uncoordinated UEs are served by coordinated BSs (where the coordination is realized via wired backhaul).

Meanwhile, a variety of strategies to handle multiuser interference in network MIMO has been proposed in the literature. In particular, *linear preequalization* (LPE) or the better-performing *Tomlinson-Harashima precoding* (THP), e.g., [3], [4], [6], [13], are of interest. Recently, these techniques have been extended to incorporate additional optimization criteria, for instance per-antenna/BS power constraints or a performance balancing among the UEs, e.g., [12], [14], [11], [5].

However, for all of the abovementioned schemes it is assumed that precoding takes place at some central instance. If the data to be transmitted to each UE initially arrives at its dedicated BS—as usual in uncoordinated wireless communication—the wired backhaul is burdened with a large amount of coordination traffic. To relieve this situation, we have recently proposed a strategy for a decentralized network MIMO processing [8]. After an initial, centralized computation of required precoding matrices, the encoding procedure is completely performed by direct signaling between BSs. Since this signaling can be conducted in a hierarchical way, the amount of backhaul traffic can be significantly reduced when tolerating only a small decrease in error rate performance.

Yet, still the question remains for which actual channel situations hierarchical coordination is sufficient to ensure a desired performance. If the current MIMO channel is very bad, a fully-coordinated precoding (or even an exclusion from transmission, i.e., declaring an outage) is advisable. In contrast, if good conditions are present, there is the chance to guarantee performance requirements without any or only little coordination effort. To this end, in this paper, we propose a selection algorithm which automatically switches between a fully- or hierarchically-coordinated scheme, or even deactivates coordination dependent on the actual MIMO channel. Thereby, a required performance is guaranteed while reducing the backhaul signaling as much as possible. Numerical simulations cover the effectiveness of the derived approach.

The paper is organized as follows: in Sec. II, a geometrical channel model for network MIMO scenarios is briefly reviewed. Sec. III discusses THP-type precoding schemes with different levels of coordination. An automated selection algorithm for these schemes is derived in Sec. IV. A short summary and conclusion can be found in Sec. V.

II. System Model

The geometrical network MIMO model shown in Fig. 1 is the basis for the considerations throughout the paper. All BSs are placed on a hexagonal grid with inter-site distance $r_{\mathrm{BS}} = 500$ m, where $N_{\mathrm{B}} = 3$ BSs are combined to supply $N_{\mathrm{U}} = 3$ single-antenna UEs jointly. The UEs are uniformly distributed within the service area (dark gray shaded in Fig. 1). UEs located outside this region are served by another triple of BSs leading to an identical scenario. The model is taken from [9], [5], [8], however, UEs located close to BSs are not omitted ($r_{\min} = 0$).

A. Channel Model

For the abovementioned geometrical model, the respective discrete-time (symbol interval T) complex baseband

This work was supported by Deutsche Forschungsgemeinschaft (DFG) within the framework COIN under grant Fi 982/4-2.

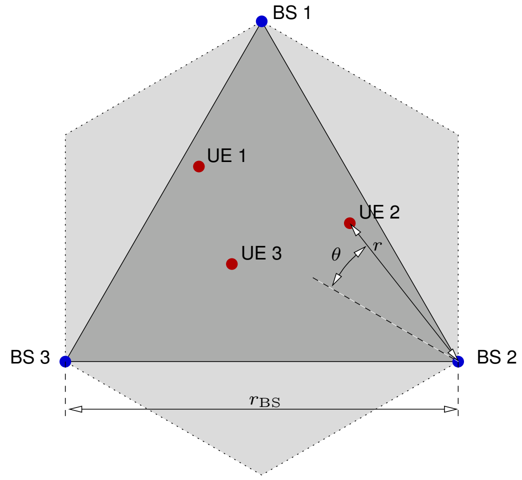

Fig. 1. Geometrical network MIMO model: a set of $N_{\mathrm{B}} = 3$ BSs serves $N_{\mathrm{U}} = 3$ UEs within the gray shaded triangular area.

MIMO system model can be defined.[1] The MIMO system model from BSs to UEs in vector/matrix notation for one time step then reads[2]

$$\boldsymbol{y} = \boldsymbol{H}\boldsymbol{x} + \boldsymbol{n}\,, \tag{1}$$

where $\boldsymbol{x} = [x_1, \ldots, x_{N_{\mathrm{B}}}]^{\mathsf{T}}$ denotes the vector of transmit symbols which are radiated from the BSs and $\boldsymbol{y} = [y_1, \ldots, y_{N_{\mathrm{U}}}]^{\mathsf{T}}$ the vector of receive symbols at the UEs. The vector $\boldsymbol{n}$ represents complex-valued zero-mean mutually uncorrelated white Gaussian noise at the receivers, i.e., $\boldsymbol{n} = [n_1, \ldots, n_{N_{\mathrm{U}}}]^{\mathsf{T}}$ and $\mathrm{E}\{\boldsymbol{n}\boldsymbol{n}^{\mathsf{H}}\} = \sigma_n^2 \boldsymbol{I}$. The noise variance of each component is given as $\sigma_n^2 \overset{\text{def}}{=} N_0/T$, with N_0 as the one-sided noise power spectral density.

The channel matrix

$$\boldsymbol{H} = \left[\, h_{u,b} \,\right]_{\substack{u=1,\ldots,N_{\mathrm{U}}\\ b=1,\ldots,N_{\mathrm{B}}}} \tag{2}$$

is assumed to be constant over some transmission burst (block-fading channel). The complex channel coefficient $h_{u,b}$ models the connection between BS b and UE u and is given as[3]

$$h_{u,b} = 10^{-L_{\mathrm{dB}}(r_{u,b})/20} \cdot 10^{A_{\mathrm{dB}}(\theta_{u,b})/20} \cdot h_{u,b,\mathrm{i.i.d.}}\,. \tag{3}$$

For the details see [1], [2], [5].

[1]The channel model contains pulse-shaping at the transmitter, the radio-frequency channel, as well as matched filtering and sampling at the receiving UE. For instance, this model is suited for per carrier considerations in multicarrier transmission.

[2]Notation: $\mathrm{E}\{\cdot\}$: expectation. $\boldsymbol{A}^{\mathsf{T}}$: transpose of matrix $\boldsymbol{A}$. $\boldsymbol{A}^{\mathsf{H}}$: Hermitian of matrix $\boldsymbol{A}$. $\boldsymbol{I}$: identity matrix. $\mathbf{diag}(\cdot)$: diagonal matrix with given elements. $\mod(\cdot)$: symmetrical modulo operation; interval depends on context.

[3]Here, the factor $L_{\mathrm{dB}}(r_{u,b})$ is composed of $L_{\mathrm{dB}}(r_{u,b}) = D_{\mathrm{dB}} + 37.6\log_{10}\left(r_{u,b}/[\mathrm{km}]\right) + S_{\mathrm{dB}}$, where D_{dB} represents the basic channel attenuation, the second quantity the path loss ($r_{u,b}$ is the distance between UE u and BS b, cf. Fig 1), and the Gaussian zero-mean random variable S_{dB} (in dB; 8 dB standard deviation) shadowing effects. The factor $h_{u,b,\mathrm{i.i.d.}}$ is a zero-mean unit variance random variable to model fast fading. Finally, $A_{\mathrm{dB}}(\theta_{u,b}) = \min\left\{12\left(\theta_{u,b}/70°\right)^2,\ 20\right\}$ is assumed as antenna pattern, with $\theta_{u,b}$ denoting the angle between UE u and BS b (cf. Fig 1).

B. Precoding for Network MIMO

Since variants of THP are considered in this paper, we briefly review the respective processing/transmission model, depicted in Fig 2.

Data symbols for UEs $u = 1, \ldots, N_{\mathrm{U}}$, in vector notation $\boldsymbol{a} = [a_1, \ldots, a_{N_{\mathrm{U}}}]^{\mathsf{T}}$, are successively encoded into a vector $\tilde{\boldsymbol{x}}$ of precoded symbols $\tilde{x}_u$ via (symmetrical) modulo operation according to the given constellation [4], [8]. The feedback matrix $\boldsymbol{B} = [b_{u,\nu}]$ with lower triangular structure and unit main diagonal serves for the cancellation of interference caused by already encoded symbols. Following this, the transmit symbols x_b, $b = 1, \ldots, N_{\mathrm{B}}$, to be radiated from the BSs, are calculated via the feedforward matrix $\boldsymbol{F} = [\boldsymbol{f}_1 \cdots \boldsymbol{f}_{N_{\mathrm{B}}}]^{\mathsf{T}}$.

Depending on the actual precoding/coordination strategy (see below), both the encoding order of the data symbols (due to the successive procedure) and the assignment of transmit symbols to BSs (in case of non-equivalent BSs) might affect the performance. Both orderings can be modeled via permutation matrices $\boldsymbol{P}_{\mathrm{U}}$ (sorting of UEs) before and $\boldsymbol{P}_{\mathrm{B}}$ (assignment of BSs) after the precoding feedback loop. Equivalently, a permuted channel matrix $\boldsymbol{P}_{\mathrm{U}}\boldsymbol{H}\boldsymbol{P}_{\mathrm{B}}$, with row permutation matrix $\boldsymbol{P}_{\mathrm{U}}$ and column permutation matrix $\boldsymbol{P}_{\mathrm{B}}$, is present.

Throughout the paper, the data symbols are assumed to be drawn from a conventional zero-mean QAM constellation with variance $\sigma_a^2 \overset{\text{def}}{=} \mathrm{E}\{|a_u|^2\}$, $\forall u$. The transmit power at each antenna is given by $\sigma_{x_b}^2 \overset{\text{def}}{=} \mathrm{E}\{|x_b|^2\}$, $\forall b$. We employ the *sum power constraint* (SPC)

$$\sum\nolimits_{b=1}^{N_{\mathrm{B}}} \sigma_{x_b}^2 = \mathrm{E}\{\boldsymbol{x}^{\mathsf{H}}\boldsymbol{x}\} \le P_{\mathrm{sum}} = N_{\mathrm{U}}\sigma_a^2\,. \tag{4}$$

In network MIMO systems, it is additionally reasonable to impose a *per-antenna power constraint* (PPC)

$$\mathrm{E}\{|x_b|^2\} \le P_{\mathrm{per}}, \quad b = 1, \ldots, N_{\mathrm{B}}\,. \tag{5}$$

We choose $P_{\mathrm{per}} = 1.5\,\sigma_a^2$, as this value is a good compromise between bit-error rate (BER) and power restrictions [5], [8].

Following the noisy MIMO channel, an individual receiver-side scaling, represented by the matrix $\boldsymbol{G} \overset{\text{def}}{=} \mathbf{diag}(g_1, \ldots, g_{N_{\mathrm{U}}})$, is used to scale the signals suitably, adjusted to the subsequent threshold device (assuming uncoded transmission). The produced estimated symbols are denoted as $\hat{a}_u$, $u = 1, \ldots, N_{\mathrm{U}}$.

As in [5], [8], we again include the basic attenuation of the radio-frequency channels into the definition of the signal-to-noise ratio (SNR), which is expressed as *transmitted energy per information bit over the noise power spectral density*. It reads

$$\frac{E_{\mathrm{b}}}{N_0'} = \frac{\sigma_a^2}{\sigma_n^2 \log_2(M)}\,, \tag{6}$$

with $N_0' = N_0 \cdot 10^{D_{\mathrm{dB}}/10}$.

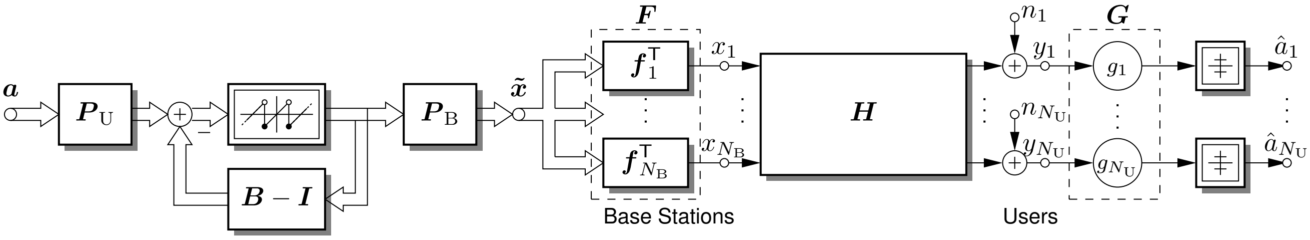

Fig. 2. Block diagram of THP in network MIMO systems.

III. Coordination Strategies for Decentralized Processing

In the literature on network MIMO usually a central processing unit (CU) fulfilling two different tasks is expected: i) assuming perfect channel knowledge and bursty transmission, the CU initially computes the precoding matrices for the specific scheme which are valid for one transmission frame; ii) subsequently, in each time step of this burst, the CU is responsible for calculating the transmit symbols $\boldsymbol{x}$ from the actual data symbols $\boldsymbol{a}$.

For the case that each data symbol is present at one (previously assigned) BS—a situation relevant in practice—significant data traffic in the wired backhaul is caused. All data symbols have to be transmitted to the CU in order to perform the precoding; the CU then returns the transmit symbols. If we measure the coordination effort by the number (denoted as β) of complex symbols to be communicated via the backhaul in each discrete time step, for classical CU processing, we have $\beta = 6$.

In contrast, a strategy for decentralized processing has been presented in [8], where the CU is just responsible for the initial computation of the precoding matrices. The calculation of the transmit symbols takes place separately in each BS, hence, the encoding is realized via a direct communication between the BSs without any CU. In the following, THP-type precoding for decentralized encoding employing different levels of coordination (no, hierarchical and full coordination) is assessed, as illustrated in Fig. 3.

A. No Coordination

If all UEs possess good channel conditions (usually when all UEs are located near their assigned BS), each BS may individually supply its UE. In that case, interference caused by other BSs does not significantly lower the performance. Consequently, there is a chance to avoid backhaul traffic completely ($\beta = 0$; cf. Fig. 3 left), and the data symbols directly constitute the transmit symbols,[4] i.e., $\boldsymbol{x} = \boldsymbol{a}$.

The performance largely depends on the assignment from data symbols to BSs. For non-coordinated (NC) transmission, there are $N_{\mathrm{B}}! = 6$ different assignments, i.e., six column permutations $\boldsymbol{H}\boldsymbol{P}_{\mathrm{B}}$ of the channel matrix. The end-to-end cascade from data to receive symbols thus reads $\boldsymbol{C}^{(\mathsf{NC})} = [c_{u,b}^{(\mathsf{NC})}] = \boldsymbol{H}\boldsymbol{P}_{\mathrm{B}}$. For uncoordinated transmission, the *signal-to-interference-plus-noise ratio* (SINR)—the relevant quantity for the BER performance—is then given as ($\zeta \overset{\text{def}}{=} \sigma_n^2/\sigma_a^2$ is the inverse SNR)

$$\mathrm{SINR}_u^{(\mathsf{NC})} \overset{\text{def}}{=} \frac{|c_{u,u}^{(\mathsf{NC})}|^2}{\zeta + \sum\limits_{\substack{l=1,\ldots,N_{\mathrm{U}} \\ l \neq u}} |c_{u,l}^{(\mathsf{NC})}|^2} \,. \qquad (7)$$

Since the overall BER performance is dominated by the worst-case user, it is convenient to maximize the minimum SINR over all users, like proposed in [12], [14], [11], [5] (min SINR criterion). Hence, the optimization task reads

$$\mathrm{minSINR}^{(\mathsf{NC})} \overset{\text{def}}{=} \min_{u=1,\ldots,N_{\mathrm{U}}} \mathrm{SINR}_u^{(\mathsf{NC})} \overset{\boldsymbol{P}_{\mathrm{B}}}{\longrightarrow} \max \,. \qquad (8)$$

This strategy requires only a very low initial computation effort, as the SINRs are directly obtained from the permutations of $\boldsymbol{H}$. Noteworthy, the power constraints (4) and (5) are immediately fulfilled.

B. Hierarchical Coordination

In [8], we have proposed a THP-type precoding scheme for a hierarchical distribution of knowledge with successive encoding; this strategy is shown in Fig. 3 (middle). BS 1 only knows its data symbol a_1 and calculates $\tilde{x}_1 = a_1$ and $x_1 = f_{1,1}\tilde{x}_1$. The precoded symbol $\tilde{x}_1$ is communicated to BS 2, which, together with its data symbol a_2 calculates $\tilde{x}_2 = \mod(a_2 - b_{2,1}\tilde{x}_1)$ and $x_2 = f_{2,1}\tilde{x}_1 + f_{2,2}\tilde{x}_2$. Finally, both precoded symbols, $\tilde{x}_1$ and $\tilde{x}_2$, are sent to BS 3, which calculates $\tilde{x}_3 = \mod(a_3 - b_{3,2}\tilde{x}_2 - b_{3,1}\tilde{x}_1)$ and $x_3 = f_{3,1}\tilde{x}_1 + f_{3,2}\tilde{x}_2 + f_{3,3}\tilde{x}_3$. This directly corresponds to a lower triangular feedforward matrix $\boldsymbol{F}$. In summary, only $\beta = 3$ symbols have to be communicated via the backhaul.

Again, it is of importance which data symbol is communicated to which BS (permutation matrix $\boldsymbol{P}_{\mathrm{B}}$). However, in addition, here the ordering of the successive encoding also has to be optimized. As in conventional THP, this ordering can be described by a row permutation matrix $\boldsymbol{P}_{\mathrm{U}}$. In summary, $N_{\mathrm{U}}!N_{\mathrm{B}}! = 36$ permutations $\boldsymbol{P}_{\mathrm{U}}\boldsymbol{H}\boldsymbol{P}_{\mathrm{B}}$ are possible in case of hierarchical coordination (HC). The end-to-end cascade is now given as $\boldsymbol{C}^{(\mathsf{HC})} = [c_{u,b}^{(\mathsf{HC})}] = \boldsymbol{P}_{\mathrm{U}}\boldsymbol{H}\boldsymbol{P}_{\mathrm{B}}\boldsymbol{F}$, and the optimization task according to the min SINR criterion reads

$$\mathrm{minSINR}^{(\mathsf{HC})} \overset{\text{def}}{=} \min_{u=1,\ldots,N_{\mathrm{U}}} \mathrm{SINR}_u^{(\mathsf{HC})} \overset{(\boldsymbol{P}_{\mathrm{U}},\boldsymbol{P}_{\mathrm{B}})}{\longrightarrow} \max \,, \qquad (9)$$

[4]Regarding Fig. 2, this can be seen as THP with feedforward and feedback matrix $\boldsymbol{F} = \boldsymbol{I}$, $\boldsymbol{B} = \boldsymbol{I}$. At the receiver scaling via $\boldsymbol{G} = \mathbf{diag}((c_{1,1}^{(\mathsf{NC})})^{-1}, \ldots, (c_{3,3}^{(\mathsf{NC})})^{-1})$ is performed.

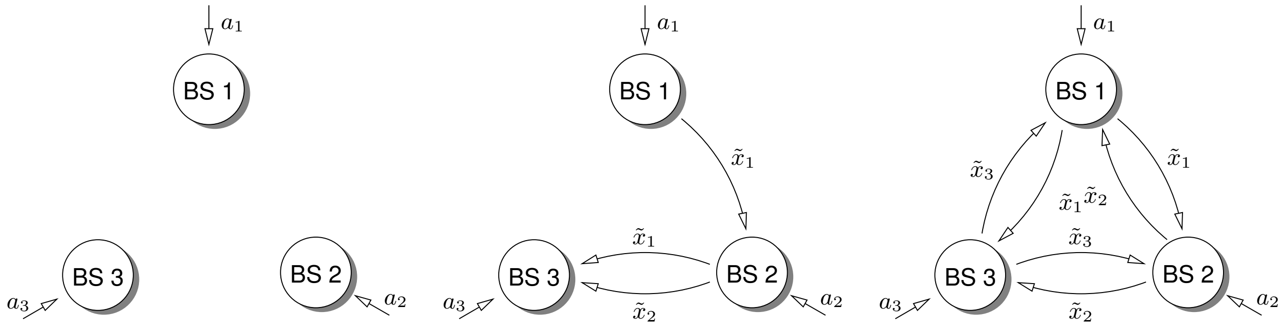

Fig. 3. Precoding strategies for decentralized network MIMO processing with different amount of coordination between BSs. Left: no coordination. Middle: hierarchical coordination. Right: full coordination.

where the SINR calculates to [5], [8]

$$\mathrm{SINR}_u^{(\mathsf{HC})} \overset{\text{def}}{=} \frac{|c_{u,u}^{(\mathsf{HC})}|^2}{\zeta + \sum\limits_{\substack{l=1,\dots,N_\mathrm{U} \\ l>u}} |c_{u,l}^{(\mathsf{HC})}|^2} \,. \tag{10}$$

The computation of $\boldsymbol{F}$ and $\boldsymbol{B}$ obeying the power constraints (4) and (5) can be performed using a *second-order cone program*, cf. [11, Algorithm 2] (numerically implementable via [10]). Regrettably, the precoding matrices have to be computed for all 36 possible permutations in order to find the best solution. To lower the initial calculation effort, in [8], a heuristic strategy for restricting the set of permutation pairs $(\boldsymbol{P}_\mathrm{U}, \boldsymbol{P}_\mathrm{B})$, for which the optimization is carried out, has been presented. In the present paper, we employ this preselection strategy[5] and only test 6 permutation pairs.

C. Full Coordination

When one or several UEs suffer from bad channel conditions, the employment of fully-coordinated precoding is advisable. In this case of classical THP, all precoded symbols $\tilde{x}_1$, $\tilde{x}_2$, and $\tilde{x}_3$ have to be known at all BSs; the respective coordination effort amounts to $\beta = 6$, cf. Fig. 3 (right). As known from conventional THP, here only the precoding order is of importance; the $N_\mathrm{U}! = 6$ row permutations $\boldsymbol{P}_\mathrm{U}\boldsymbol{H}$ lead to different SINRs. The end-to-end cascade for full coordination (FC) precoding reads $\boldsymbol{C}^{(\mathsf{FC})} = [c_{u,b}^{(\mathsf{FC})}] = \boldsymbol{P}_\mathrm{U}\boldsymbol{H}\boldsymbol{F}$ and the optimization task is given as

$$\mathrm{minSINR}^{(\mathsf{FC})} \overset{\text{def}}{=} \min_{u=1,\dots,N_\mathrm{U}} \mathrm{SINR}_u^{(\mathsf{FC})} \xrightarrow{\boldsymbol{P}_\mathrm{U}} \max \,, \tag{11}$$

where $\mathrm{SINR}_u^{(\mathsf{FC})}$ is defined in (10), replacing the superscript "HC" by "FC".

For FC precoding, the BLAST algorithm yields the (nearly) optimum encoding order [13]. Using the associated permutation $\boldsymbol{P}_\mathrm{U}$ the precoding matrices are again computed via second-order cone program obeying (4) and (5) [5].

[5] Although the channel model in this paper has slightly changed compared to [8] (no exclusion of UEs located near to BSs), the empirically obtained probabilities, based on which the permutations are chosen, are almost the same as in [8].

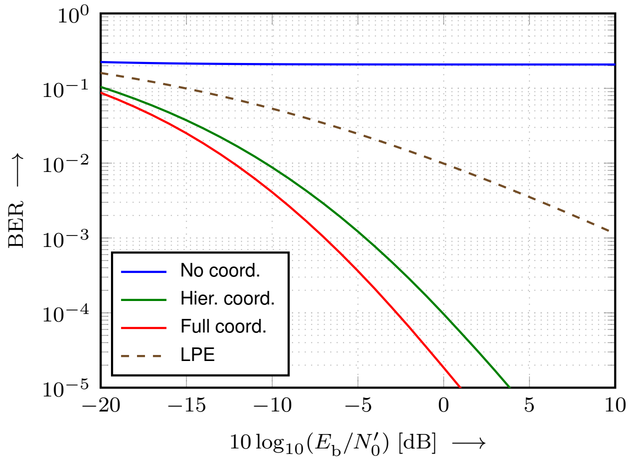

Fig. 4. BER over the SNR for the network MIMO scenario at hand. 16QAM signaling; per-antenna power constraint $P_\mathrm{per} = 1.5\,\sigma_a^2$. Different levels of coordination; for comparison, the result for the optimized version of LPE [5] is given.

D. BER Performance and Comparison

For a comparison of the achievable performance of the abovementioned coordination strategies, Fig. 4 shows the BER curves over the SNR. In the Monte Carlo simulations 50000 channel realizations were considered; each one for a burst of 50000 data symbols.[6] The overall BER performance is dominated by the UEs possessing the lowest SINRs (among all possible channel instances and users). Noteworthy, the computation of precoding matrices according to the min SINR criterion leads to a nearly balanced SINR among all users [5]. In contrast, the SINR values for uncoordinated transmission may extremely differ due to the lack of interference handling.

Apparently, transmission without any coordination is not reasonable as almost all channel realizations cause significant interference. In contrast, between hierarchically- and fully-coordinated THP, merely a gap of about 2.5 dB is present in the high-SNR/low-BER regime. This means that for many channel realizations, the application of partly-coordinated network MIMO is sufficient. This fact is further examined and discussed in the following section.

[6] The curves slightly differ from that in [5], [8], as $r_\mathrm{min} = 0$ is chosen.

Noteworthy, even hierarchical precoding significantly outperforms fully-coordinated ($\beta = 6$) linear preequalization.

IV. Automated Selection of the Coordination Strategy

In this section, the possibility of an automated selection of the coordination effort is discussed. Depending on the current channel, our objective will be the reduction of the backhaul traffic, ensuring a predefined transmission quality (BER demand).

A. SINR Statistics

As already discussed in the last section, regardless of the specific precoding strategy, the minimum SINR is determining the BER performance. Consequently, we are interested in the distribution of this quantity. Fig. 5 depicts the empirically determined *probability density function* (pdf) of $\min\mathrm{SINR}^{(\cdot)}$ in dB for all considered precoding strategies. For comparison, Gaussian distributions with the same mean value and the same standard deviation are shown (black lines). Additionally, Fig. 6 illustrates the resulting mean values μ_{dB} and standard deviations σ_{dB} over the SNR (each in dB).[7]

As can be seen, for uncoordinated transmission the SNR only has a minor impact on the pdf (interference-limited regime). This behavior explains the hardly decreasing BER curve in Fig. 4. Just in a few cases, the uncoordinated scheme results in an acceptable transmission quality. In contrast to that, the statistics of hierarchical coordination do not differ a lot from the ones for full coordination. Especially in the high-SNR regime full coordination shows its advantage; in double-logarithmic scale, expectation and standard deviation linearly increase over the SNR[8] for FC, whereas some residual interferences are still present in HC.

Particularly in the case of fully-coordinated THP (optimized according to the min SINR criterion), the pdfs of the minimum SINR are very well approximated by a *log-normal* one, i.e.,

$$\mathrm{pdf}(\min\mathrm{SNR}^{(\mathsf{FC})}) = \frac{1}{\sqrt{2\pi\sigma_{\mathrm{dB}}^2}} \cdot \exp\left(-\frac{(10\log_{10}(\min\mathrm{SINR}^{(\mathsf{FC})}) - \mu_{\mathrm{dB}})^2}{2\,\sigma_{\mathrm{dB}}^2}\right) . \quad (12)$$

This fact is subsequently used to develop an adequate selection model.

B. Selection Algorithm

Since fully-coordinated THP is the best-performing precoding strategy (among that considered in this paper), for an automated selection, its average BER curve (cf. Fig. 4) can be seen as a reference. In turn, to each SNR value E_{b}/N_0', a corresponding SINR value, which is minimally required, is associated. If, for the *actual* channel, uncoordinated transmission or hierarchical precoding are able to achieve a $\min\mathrm{SINR}^{(\cdot)}$ at least as large as the required one, these schemes are sufficient to guarantee the desired performance but with lowered backhaul traffic.

In addition, performance can be increased by discarding "poor" channel realizations, hence causing an outage.[9]

For an automated selection, still the question has to be answered of how to specifically choose the required SINR (for a given SNR) for which uncoordinated or hierarchical transmission, respectively, are accepted. To this end, we consider that value of $\min\mathrm{SINR}^{(\mathsf{FC})}$, de-

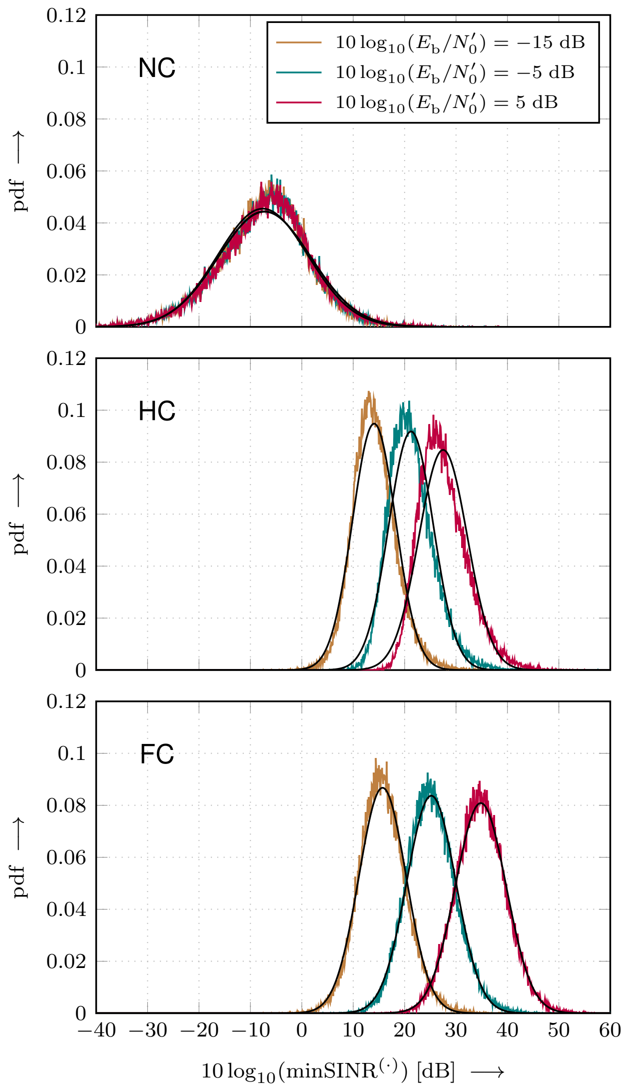

Fig. 5. Probability density function of the minimum SINR (in dB) for different SNRs (empirically evaluated via 50000 channel realizations). Top: no coordination (NC); Middle: hierarchical coordination (HC); Bottom: full coordination (FC). Black: Gaussian distributions with same mean value and same standard deviation. 16QAM signaling; per-antenna power constraint $P_{\mathrm{per}} = 1.5\,\sigma_a^2$.

[7] In dependency of $\zeta = \sigma_n^2/\sigma_a^2$, these parameters are well approximated by $\mu_{\mathrm{dB}} = 9.47\log_{10}(\zeta^{-1}) - 4.15$ [dB] and $\sigma_{\mathrm{dB}} = 0.16\log_{10}(\zeta^{-1}) + 4.28$ [dB] for fully-coordinated precoding.

[8] Since interference is handled, $\mathrm{SINR}_u^{(\mathsf{FC})} \approx |c_{u,u}^{(\mathsf{FC})}|^2 \cdot \frac{\sigma_a^2}{\sigma_n^2}$.

[9] In such cases, alternative orthogonal multiplexing strategies such as time-division or frequency-division multiplexing may be used.

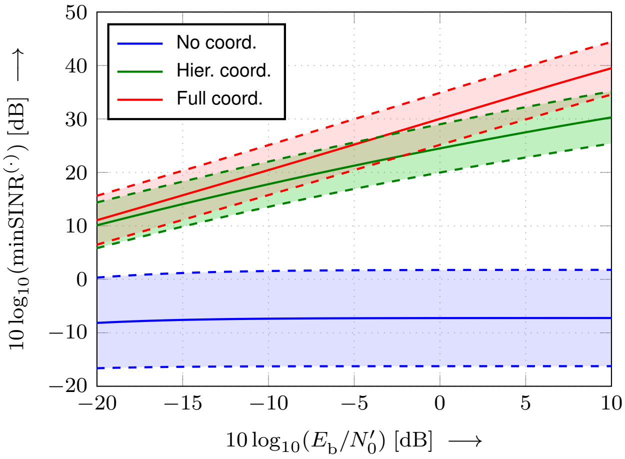

Fig. 6. Expectation μ_dB (solid) and standard deviation σ_dB (dashed; $\mu_\mathrm{dB} \pm \sigma_\mathrm{dB}$ shown) for different coordination strategies over the SNR (in dB; empirically evaluated via 50000 channel realizations). 16QAM signaling; per-antenna power constraint $P_\mathrm{per} = 1.5\,\sigma_a^2$.

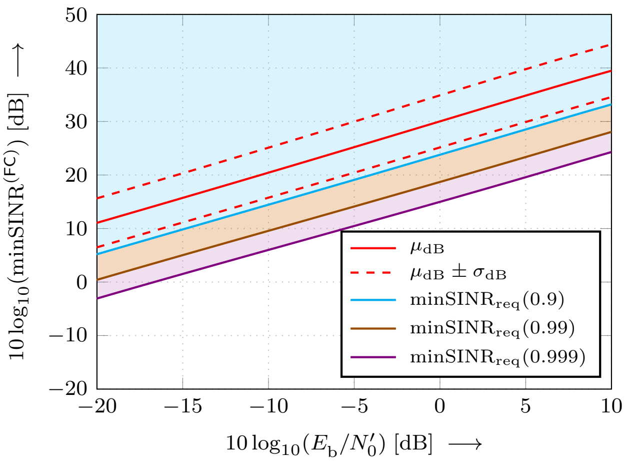

Fig. 7. Quantiles $\mathrm{minSINR}_\mathrm{req}(p_\mathrm{e})$ of the log-normal model for the $\mathrm{minSINR}^{(\mathsf{FC})}$ distribution over the SNR. Empirically determined expectations μ_dB and standard deviations σ_dB taken from Fig. 6. 16QAM signaling; per-antenna power constraint $P_\mathrm{per} = 1.5\,\sigma_a^2$.

noted as $\mathrm{minSINR}_\mathrm{req}(p_\mathrm{e})$, which (relying on the log-normal model) is exceeded with a given probability $p_\mathrm{e} \in [0,1]$. Mathematically, $\mathrm{minSINR}_\mathrm{req}(p_\mathrm{e})$ corresponds to the $(1-p_\mathrm{e})$-*quantile* of the $\mathrm{minSINR}^{(\mathsf{FC})}$ distribution, i.e., the inverse of the *cumulative distribution function* (cdf) at probability $1-p_\mathrm{e}$. For the log-normal distribution at hand, these quantiles are determined as

$$\mathrm{minSINR}_\mathrm{req}(p_\mathrm{e}) = 10^{(\mu_\mathrm{dB}+\sigma_\mathrm{dB}\sqrt{2}\,\mathrm{erf}^{-1}(1-2\,p_\mathrm{e}))/10}\,, \tag{13}$$

where $\mathrm{erf}^{-1}(\cdot)$ is the inverse of the error function. Fig. 7 exemplarily depicts the (SNR-dependent) quantiles $\mathrm{minSINR}_\mathrm{req}(p_\mathrm{e})$ for a choice of $p_\mathrm{e} = 0.9,\, 0.99,\, 0.999$.

In the same manner, we are able to define a threshold $\mathrm{minSINR}_\mathrm{out}(p_\mathrm{no})$, below which we declare an outage. From the above mathematics, $1-p_\mathrm{no}$ is then the outage probability (p_no is the no-outage probability) for a fully-coordinated scheme. Obviously, $p_\mathrm{e} \leq p_\mathrm{no}$ has to

Algorithm 1 Automated selection of coordination strategy

Require: channel matrix $\boldsymbol{H}$, inverse SNR ζ, exceedance probability p_e and no-outage probability p_no

$\mathrm{minSINR}_\mathrm{req} := 10^{(\mu_\mathrm{dB}+\sigma_\mathrm{dB}\sqrt{2}\,\mathrm{erf}^{-1}(1-2\,p_\mathrm{e}))/10}$;
$\mathrm{minSINR}_\mathrm{out} := 10^{(\mu_\mathrm{dB}+\sigma_\mathrm{dB}\sqrt{2}\,\mathrm{erf}^{-1}(1-2\,p_\mathrm{no}))/10}$;
// Check if non-coordinated transmission sufficient
$\boldsymbol{P}_\mathrm{B} := \mathrm{optimizeNC}(\boldsymbol{H},\zeta)$; *// cf.* (8)
if $\mathrm{minSINR}^{(\mathsf{NC})}(\boldsymbol{H}\boldsymbol{P}_\mathrm{B},\zeta) \geq \mathrm{minSINR}_\mathrm{req}$ **then**
  $\boldsymbol{P}_\mathrm{U} := \boldsymbol{I}$; $\boldsymbol{F} := \boldsymbol{I}$; $\boldsymbol{B} := \boldsymbol{I}$;
  return $\boldsymbol{P}_\mathrm{U}$, $\boldsymbol{P}_\mathrm{B}$, $\boldsymbol{F}$, $\boldsymbol{B}$; *// Return if successful*
end if
// Check if hierarchical coordination sufficient
$[\boldsymbol{F},\boldsymbol{B},\boldsymbol{P}_\mathrm{U},\boldsymbol{P}_\mathrm{B}] := \mathrm{optimizeHC}(\boldsymbol{H},\zeta)$; *// cf.* (9)
if $\mathrm{minSINR}^{(\mathsf{HC})}(\boldsymbol{P}_\mathrm{U}\boldsymbol{H}\boldsymbol{P}_\mathrm{B}\boldsymbol{F},\zeta) \geq \mathrm{minSINR}_\mathrm{req}$ **then**
  return $\boldsymbol{P}_\mathrm{U}$, $\boldsymbol{P}_\mathrm{B}$, $\boldsymbol{F}$, $\boldsymbol{B}$; *// Return if successful*
end if
// Check if full coordination sufficient
$[\boldsymbol{F},\boldsymbol{B},\boldsymbol{P}_\mathrm{U}] := \mathrm{optimizeFC}(\boldsymbol{H},\zeta)$; *// cf.* (11)
if $\mathrm{minSINR}^{(\mathsf{FC})}(\boldsymbol{P}_\mathrm{U}\boldsymbol{H}\boldsymbol{F},\zeta) \geq \mathrm{minSINR}_\mathrm{out}$ **then**
  $\boldsymbol{P}_\mathrm{B} := \boldsymbol{I}$;
  return $\boldsymbol{P}_\mathrm{U}$, $\boldsymbol{P}_\mathrm{B}$, $\boldsymbol{F}$, $\boldsymbol{B}$; *// Return if successful*
else
  return FAIL; *// Declare outage*
end if

hold. Subsequently, a suited choice of the probabilities p_e and p_no, respectively, or corresponding thresholds $\mathrm{minSINR}_\mathrm{req}(p_\mathrm{e})$ and $\mathrm{minSINR}_\mathrm{out}(p_\mathrm{no})$, is discussed in detail.

From the above theoretical considerations, we are able to state an automated selection algorithm for reducing the backhaul traffic while guaranteeing the desired performance, see Algorithm 1. As input quantities, the algorithm demands for the actual channel matrix $\boldsymbol{H}$, the inverse SNR ζ, and a proper choice of the exceedance probability p_e and the no-outage probability p_no. First, the thresholds $\mathrm{minSINR}_\mathrm{req}$ and $\mathrm{minSINR}_\mathrm{out}$ are calculated from p_e and p_no, respectively (cf. (13)). Then, it is tested whether non-coordinated transmission suffices, i.e., if $\mathrm{minSINR}^{(\mathrm{NC})}$ is above the threshold. If this is the case, the algorithm can already be terminated. Otherwise, it is tested whether hierarchical THP suffices, i.e., if $\mathrm{minSINR}^{(\mathrm{HC})}$ is above the threshold. Fully-coordinated precoding is the fall-back option. Alternatively, if outage is allowed ($p_\mathrm{no} < 1$) it is tested whether fully-coordinated precoding fulfills the set demand.

C. Numerical Results

For the numerical evaluation of the selection algorithm, we again examine 50000 different channel realizations, where each channel matrix is constant over a burst of 50000 symbols. First, we restrict to the case when the exclusion of channels is turned off ($p_\mathrm{no} = 1$); then, we study the effect of allowing outage.

1) Selection without Outage: Fig. 8 depicts the average BER over all users and channel realizations over the SNR

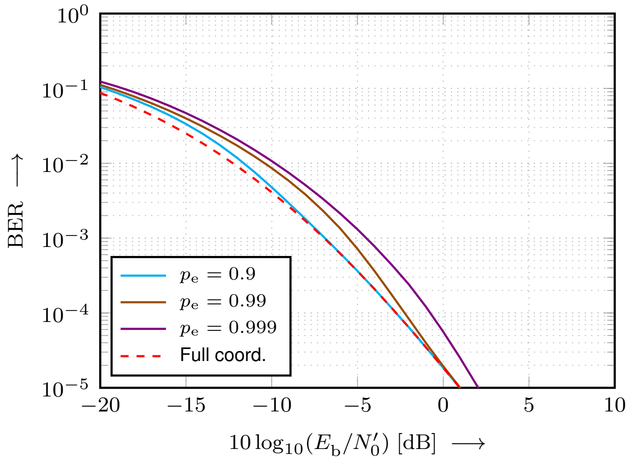

Fig. 8. BER over the SNR when employing the automated selection algorithm without outage. 16QAM signaling; per-antenna power constraint $P_\mathrm{per} = 1.5\,\sigma_a^2$. Parameter: exceedance probability p_e. Fully-coordinated precoding (cf. Fig. 4) is shown as reference.

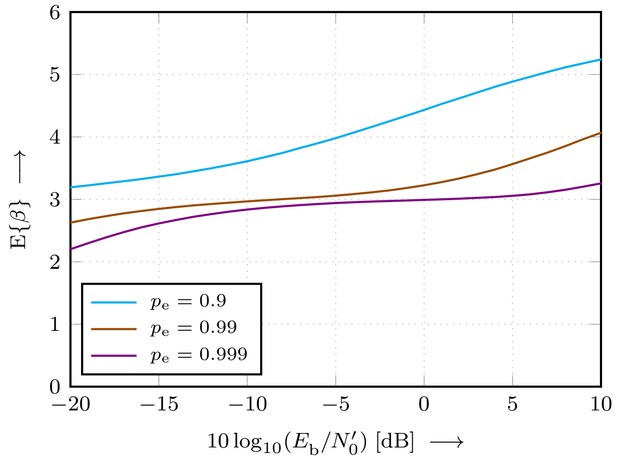

Fig. 10. Average backhaul traffic $\mathrm{E}\{\beta\}$ (in complex symbols per discrete time step) over the SNR when employing Algorithm 1. 16QAM signaling; per-antenna power constraint $P_\mathrm{per} = 1.5\,\sigma_a^2$. Parameter: exceedance probability p_e; no outage ($p_\mathrm{no} = 1$).

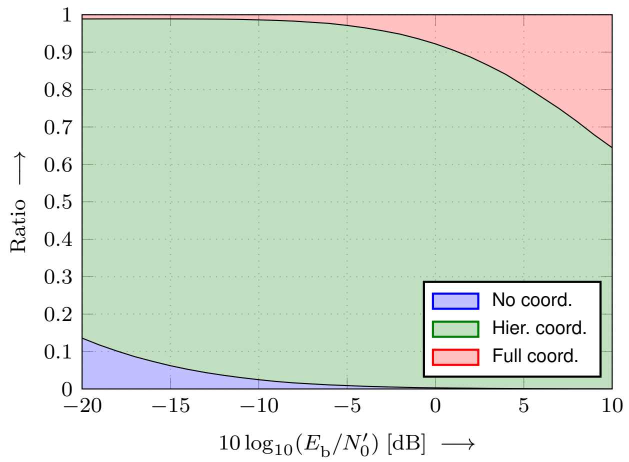

Fig. 9. Ratio of network MIMO coordination strategies in dependency of the SNR when employing Algorithm 1. 16QAM signaling; per-antenna power constraint $P_\mathrm{per} = 1.5\,\sigma_a^2$. Exceedance probability $p_\mathrm{e} = 0.99$; no outage ($p_\mathrm{no} = 1$).

for different choices of p_e. For comparison, the reference BER curve of fully-coordinated precoding is given. If we consider the case $p_\mathrm{e} = 0.9$, which means that we use a high threshold for the acceptance of the NC or HC schemes (cf. Fig. 7), there is hardly any difference from the results using the selection to that when using only FC THP. Assuming $p_\mathrm{e} = 0.99$, there is a noticeable difference for low SNRs. However, in the low-BER/high-SNR regime, that choice for the parameter is reasonable since the BER performance is nearly equivalent to fully-coordinated precoding. In contrast, when selecting $p_\mathrm{e} = 0.999$, very often NC or HC is chosen, hence a poorer performance (loss of 1 to 2.5 dB) is obtained.

Obviously, both p_e and the SNR at which the MIMO system operates determine the ratio of the coordination strategies selected by the algorithm. Fig. 9 illustrates each strategy's percentage of usage for $p_\mathrm{e} = 0.99$. Since the fully-coordinated scheme is the reference, increasing the SNR lowers the tolerated BER and hence increases the required SINR. Non-coordinated transmission is selected only for very low SNRs, which coincides with the conclusions obtained from Fig. 5. In the low- to mid-SNR range, HC precoding is sufficient for most of the channel realizations. Only in the high-SNR regime (low BER requirements), full coordination becomes important as there the superiority of FC over HC is more pronounced (cf. Fig. 5 or 6).

Since the ratio of each coordination strategy directly determines the (average) traffic in the backhaul, both the SNR and p_e influence the number β of complex symbols to be transmitted per discrete time step. In Fig. 10, $\mathrm{E}\{\beta\}$ is visualized in dependency of these two quantities. Again, increasing the SNR lowers the tolerated BER and hence FC is more and more preferred. Consequently, an increase in average coordination effort is visible over the SNR. This effect also holds for lowering p_e, as the acceptance ratio of NC and HC is decreased, too.

If we select $p_\mathrm{e} = 0.9$ (which is a good choice in the low-SNR regime, cf. Fig. 8), the traffic can be lowered from $\beta = 6$ for FC to $\mathrm{E}\{\beta\} \approx 3\ldots4$. The gains of the selection become even more pronounced in the high SNR region. Here, $p_\mathrm{e} = 0.99$ is reasonable and the average coordination effort $\mathrm{E}\{\beta\}$ can be kept below 3.5. In summary, using the proposed selection strategy, without noticeable loss in BER performance, the backhaul traffic can almost be halved.

The trade-off between the backhaul traffic and the SNR required for guaranteeing a desired target BER (here $\mathrm{BER} = 10^{-4}$) is depicted in Fig. 11. Again, for the moment, no outage is allowed ($p_\mathrm{no} = 1$, red curve). Obviously, the lowest required SNR in order to achieve the target BER is obtained by restricting to FC precoding (i.e, $p_\mathrm{e} = 0$ or $\min \mathrm{SINR}_\mathrm{req} = \infty$) which results in the worst-case backhaul traffic of $\beta = 6$. When enabling the automated selection by increasing p_e (lowering

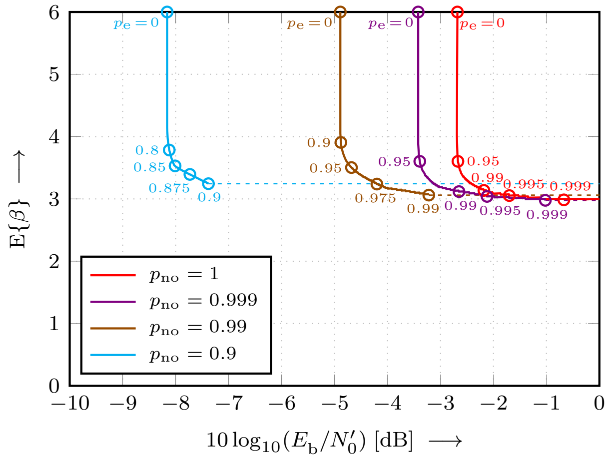

Fig. 11. Trade-off between required SNR and average backhaul traffic $E\{\beta\}$ for guaranteeing a target BER of 10^{-4}. 16QAM signaling; per-antenna power constraint $P_{per} = 1.5\,\sigma_a^2$. Parameter: no-outage probability p_{no}; variation of p_e. A dashed line is drawn if, in each case, a BER lower than the target BER is achieved.

$\min\mathrm{SINR}_{req}$), $E\{\beta\}$ decreases significantly almost without any increase in required SNR. Employing a probability p_e beyond approximately 0.99 does provide only very little additional gains in $E\{\beta\}$ but significantly increases the required SNR; the curve flattens out at $E\{\beta\} = 3$. Hence, $p_e \approx 0.99$ is a good compromise between the amount of coordination and SNR requirements.

2) Selection with Outage: Finally, we investigate the influence of allowing an outage, i.e., the effects when "bad" channels are excluded from transmission. To this end, we again consider the trade-off between backhaul traffic and SNR required for guaranteeing a desired target BER (here $\mathrm{BER} = 10^{-4}$), see Fig. 11. The no-outage probability is chosen to $p_{no} = 0.999$, 0.99, and 0.9, i.e., for the 0.1, 1, and 10 % worst performing channels according to the Gaussian model an outage is declared.

As expected, the exclusion of channels with poor conditions in general leads to lower required SNRs. The trade-off curves basically exhibit the same characteristics (L shape) as in the case of no outage. However, the specific values of p_e for the optimum trade-off depend on the actual value p_{no}. Remarkably, even the exclusion of only the 1 % worst performing channels can lower the required SNR by about 2 dB. As already explained, only $p_e \leq p_{no}$ is reasonable. Hence, the curves end at $p_e = p_{no}$. Beyond this point (visualized by the dashed line) the precoding strategy selected by the algorithm always performs better than the desired target BER.

V. Summary and Conclusion

In this paper, an approach for the automated selection of coordination strategies in network MIMO scenarios has been presented. To this end, different precoding strategies (no coordination, hierarchical precoding, conventional fully-coordinated precoding) have briefly been reviewed and assessed. A straightforward selection algorithm has been proposed, which classifies the current channel situation and activates only that coordination strategy which guarantees a desired performance with the lowest amount of backhaul signaling. For that purpose, a log-normal model for the minimum SINR over the users achievable with fully-coordinated precoding has been employed.

Numerical simulations have revealed that the selection is able to nearly halve the amount of backhaul traffic without any noticeable negative impact on the BER performance. By adjusting the free parameter (exceedance probability or minimally required SINR) a trade-off between required SNR and signaling is enabled. Finally, we have assessed the exclusion of bad-performing channels. Allowing an outage, the required SNR can be lowered significantly.

In summary, the automated selection of the degree of cooperation is a promising technique for future network MIMO scenarios. For the vast majority of channel realizations, fully-cooperated precoding (classical THP) is not required at all; very often the hierarchical scheme proposed in [8] is sufficient. Consequently, only the amount of backhaul traffic really needed should be invested.

References

[1] 3GPP. Technical Specification Group Radio Access Network, RF System Scenarios, Technical Report TR 25.996 V3.3.0, 2002.

[2] 3GPP. Technical Specification Group Radio Access Network, Physical Layer Aspects for Evolved Universal Terrestrial Radio Access (UTRA), Technical Report TR 25.814 V7.1.0, 2006.

[3] R.F.H. Fischer, C. Windpassinger, A. Lampe, J.B. Huber. MIMO Precoding for Decentralized Receivers. In *Proc. Int. Symp. on Information Theory (ISIT)*, Lausanne, Switzerland, p. 496, 2002.

[4] R.F.H. Fischer. *Precoding and Signal Shaping for Digital Transmission*. John Wiley & Sons, New York, 2002.

[5] R.F.H. Fischer, S. Heinrichs. The Network MIMO Downlink with Per-Antenna Power Constraint: Linear Preequalization vs. Tomlinson-Harashima Precoding. In *Proc. 17th Int. ITG Workshop on Smart Antennas (WSA)*, Stuttgart, Germany, Mar. 2013.

[6] M. Joham, J. Brehmer, W. Utschick. MMSE Approaches to Multiuser Spatio-Temporal Tomlinson-Harashima Precoding. In *Proc. 5th Int. ITG Conf. on Source and Channel Coding (SCC)*, pp. 387–394, Erlangen, Germany, Jan. 2004.

[7] M.K. Karakayali, G.J. Foschini, R.A. Valenzuela. Network Coordination for Spectrally Efficient Communication in Cellular Systems. *IEEE Wireless Communications*, pp. 56–61, Aug. 2006.

[8] S. Stern, R.F.H. Fischer. Hierarchical Precoding for the Network MIMO Downlink. In *Proc. 10th Int. ITG Conference on Systems, Communications and Coding (SCC)*, Hamburg, Germany, Feb. 2015.

[9] C. Stierstorfer, C. Siegl, R.F.H. Fischer, T. Wild, C. Hoek. Network MIMO Downlink Transmission. In *Proc. 15th Int. OFDM Workshop (InOWo)*, Hamburg, Germany, Sept. 2010.

[10] J. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software 11–12; Special issue on Interior Point Methods*, pp. 625–653, 1999.

[11] A. Tölli, M. Codreanu, M. Juntti. Linear Multiuser MIMO Transceiver Design With Quality of Service and Per-Antenna Power Constraints. *IEEE Transactions on Signal Processing*, pp. 3049–3055, July 2008.

[12] A. Wiesel, Y.C. Eldar, S. Shamai (Shitz). Linear Precoding via Conic Optimization for Fixed MIMO Receivers. *IEEE Transactions on Signal Processing*, pp. 161–176, Jan. 2006.

[13] C. Windpassinger. *Detection and Precoding for Multiple Input Multiple Output Channels*. PhD Thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg, June 2004.

[14] W. Yu, T. Lan. Transmitter Optimization for the Multi-Antenna Downlink With Per-Antenna Power Constraints. *IEEE Transactions on Signal Processing*, pp. 2646–2660, June 2007.