Low-complexity Distributed MIMO Receiver and its Implementation on the OpenAirInterface Platform

Florian Kaltenberger, Rizwan Ghaffar, Raymond Knopp
Eurecom, Sophia Antipolis, France
Email (first name.last name)@eurecom.fr

***Abstract*—In order to maximize the system throughput, future wireless communication systems will employ a very tight frequency reuse. This leads to interference limited systems where the interference is high at the cell edges. The key ingredient to such networks are thus receivers that are able to exploit the structure of this interference instead of nulling or attenuating it. In this paper we apply such a receiver structure to a distributed multiple-input multiple-output (MIMO) scenario to decode two independent data streams from two synchronized base stations. Further we show how the distributed MIMO receiver is implemented in real-time on the OpenAirInterface platform and provide results from field trials and compare them to the simulation results. Applications of the distributed MIMO receiver include single-frequency cellular as well as mesh networks. OpenAirInterface is an experimental open-source real-time hardware and software platform for experimentation in wireless communications and signal processing. Its current implementation provides a full software modem comprising physical and link layer functionalities for cellular and mesh network topologies.**

I. Introduction

One major challenge in future wireless networks is interference which is caused by a very tight frequency reuse in order to increase the network throughput. Interference is especially strong for users at the cell edge severely limiting the user's throughput. Most state-of-the-art wireless systems deal with the interference either by orthogonalizing the communication links in time or frequency [1] or by allowing the communication links to share the same degrees of freedom but modelling the interference as additive Gaussian random process [2]. Both of these approaches may be suboptimal as the first approach entails an *a priori* loss of degrees of freedom in both links, no matter how weak the potential interference is while the second approach treats interference as pure noise while it actually carries information and has structure that can be potentially exploited in mitigating its effect.

In [3] a receiver structure that exploits the structure of the interference in the detection process has been proposed. It was shown that exploiting the fact that the interference comes from a finite constellation alphabet, the information rate of the desired stream can be increased by up to one bit/sec/Hz. The same authors also derive a low-complexity max-log maximum a posteriori (MAP) receiver for these scenarios [4]. In this paper we apply this kind of receiver to a distributed multiple-input multiple-output (MIMO) scenario, where one receiver is able to decode independent data streams from two different base stations.

The receiver was tested on the Eurecom OpenAirInterface platform, which is an experimental real-time open-source hardware and software platform for future wireless networks [5]. In particular, the receiver is used in the mesh network topology of the platform, where a mesh router (MR) is connected to two clusterheads (CH) simultaneously. The receiver is fully implemented in software and integrates seamlessly in the protocol stack. Synchronization between CHs is achieved by the distributed synchronization algorithm described in [5, 6].

II. OpenAirInterface Overview

In this section give a brief overview of the relevant features of OpenAirInterface needed for this paper, namely the hardware (Subsection II-A) and the physical layer (Subsection II-B). For a more detailed overview see [5].

A. Hardware Components

OpenAirInterface can be run on different hardware modules. All the experiments described in this paper are based on CardBus MIMO 1 (CBMIMO1) cards[1]. The CBMIMO1 board comprises two time-division duplex (TDD) radio frequency (RF) chains operating at 1.900-1.920 GHz with 5 MHz channels and 21dBM transmit power per antenna for an orthogonal frequency division modulated (OFDM) waveform. As the name suggests it communicates with a host PC over the CardBus/PCMCIA interface. The cards house a medium-scale field programmable gate array (FPGA) (Xilinx X2CV3000) which is mainly responsible for interfacing with RF frontend as well as framing of the transmit and receive signals. All the PHY layer signal processing is usually run in real-time on the host PC under the control of the real-time application interface (RTAI), which is an extension to the Linux operating system.

B. Physical Layer

The physical (PHY) layer of the OpenAirInterface platform targets WiMax and UMTS LTE like networks and thus uses orthogonal frequency division multiple access (OFDMA) together with multiple-input multiple-output (MIMO).

OpenAirInterface makes use of punctured binary codes (64-state rate 1/2 convolutional or 8-state rate 1/3 3GPP/LTE

Eurecom's Research is supported in part by its industrial partners: Swisscom, Thales, SFR, Orange France, Hitachi Europe, STMicroelectronics, Bouygues Telecom, Sharp, Cisco, BMW Group. This work is also funded in part by the European Commission through project CHORIST (ICT/FP6), the French National Research Agency (ANR) through projects AIRNET and HNPS, and by the Wiener Wissenschafts-, Forschungs-und Technologiefonds (WWTF) through the project PUCCO.

[1]A successor to CBMIMO1 called Express MIMO will be available soon.

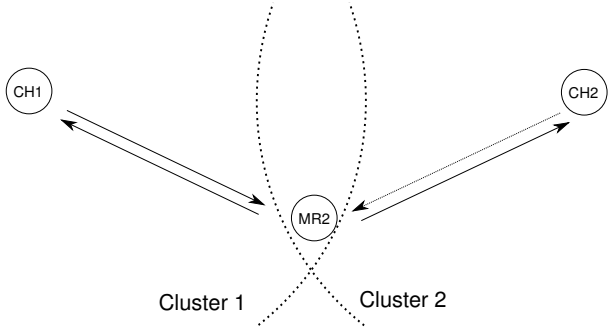

Fig. 1. The mesh network topology is organized in clusters. Each cluster is controlled by a cluster head (CH). Other nodes in the network are called mesh routers (MR) since they can be used to relay information between CHs.

Turbo code). Puncturing can use either 3GPP rate matching or random puncturing in order to fine tune the coding rate to adapt to configurable transport block sizes delivered to PHY by the MAC. Each transmitted block is punctured and then passed to a bit-interleaver and modulation mapper. This technique is referred to as bit interleaved coded modulation (BICM). OpenAirMesh supports QPSK, 16-QAM and 64-QAM modulation although currently only QPSK is used. The transmitted transport blocks can be split into two spatial streams in the case of point-to-point MIMO transmission.

The modulated symbols are multiplied by an adjustable amplitude and passed to the space-time-frequency (STF) parser. The STF parser multiplexes the pilot symbols and the data symbols into OFDM symbols. In the case of one available spatial stream, the STF parser also performs fast antenna cycling, i.e., every subcarrier is transmitted from a different antenna. This way each stream sees all the degrees of freedom of the channel. In the case of two spatial streams, the STF parser guarantees that both streams use different antennas in the same time/frequency dimension. This is a form of superposition coding since the two streams are combined additively in the air.

This design allows using the same transmitter and receiver structure both for point-to-point MIMO as well as distributed MIMO transmission. In the latter case one spatial stream is used at each source and the second stream originates in another part of the network, either in the same cluster or an adjacent cluster. A particular user can decode both streams or simply select the one it requires. In Section IV we derive a low-complexity max-log MAP receiver for this design.

III. Signal Model

Consider the scenario depicted in Fig. 1 with two transmitters (called clusterheads, CH) and one receiver (called mesh router, MR). We assume the network is fully synchronized. This can be achieved by using CH1 as the reference clock and letting the MR relay the synchronization symbol to CH2 [5, 6]. We assume that each CH has n_t antennas and MR has n_r antennas. Let $\mathbf{x}_{m,q}^{(j)}$ denote $n_t \times 1$ vector of the transmit symbols for subcarrier q of OFDM symbol m of CH j, $j = 1, 2$. We assume that the transmit symbols are taken from a signal set $\chi_j \subseteq \mathbb{C}$ of size $|\chi_j| = M_j$ with a Gray labeling map $\mu_j : \{0,1\}^{\log|M_j|} \to \chi_j$ while $j \in [1, 2]$. [2]

Cascading the IFFT at the MR and the FFT at the CHs with cyclic prefix (CP) extension, the received signal at MR2 at q-th frequency tone and the m-th OFDM symbol can be expressed as:

$$\mathbf{y}_{m,q} = \mathbf{H}_q^{(1)}\mathbf{x}_{m,q}^{(1)} + \mathbf{H}_q^{(2)}\mathbf{x}_{m,q}^{(2)} + \mathbf{z}_{m,q} \quad (1)$$

where $\mathbf{H}_q^{(1)}$ and $\mathbf{H}_q^{(2)}$ denote the $n_r \times n_t$ MIMO channel between CH1 and MR2 and between CH2 and MR2. The channel is assumed to be frequency selective (i.e., it varies with subcarrier index q) and block fading (i.e., constant over the OFDM symbols of a frame). $\mathbf{z}_{m,q} \in \mathbb{C}^{n_r}$ is the vector of circularly symmetric complex white Gaussian noise of variance σ^2.

Since each clusterhead transmits only one spatial stream and antenna cycling is used, only one element of $\mathbf{x}_{m,q}^{(j)},\ j = 1, 2$ is non-zero for every m and q. We identify this non-zero element with $x_{m,q}^{(j)},\ j = 1, 2$ and can rewrite (1) equivalently as

$$\mathbf{y}_{m,q} = \mathbf{h}_q^{(1)}x_{q,m}^{(1)} + \mathbf{h}_q^{(2)}x_{m,q}^{(2)} + \mathbf{z}_{m,q}, \quad (2)$$

where $\mathbf{h}_q^{(1)}$ and $\mathbf{h}_q^{(2)}$ are the equivalent channel vectors for the non-zero elements. The complex symbols $x_{m,q}^{(1)}, x_{m,q}^{(2)}$ of the 2 streams are assumed to be independent and of variances σ_1^2 and σ_2^2 respectively. Assuming that the first stream is the desired stream, the signal to noise ratio (SNR) is given by $\frac{\sigma_1^2}{\sigma^2}$ and the signal to interference ratio (SIR) by $\frac{\sigma_1^2}{\sigma_2^2}$.

For notational convenience, we can drop the frequency and OFDM symbol indexes and rewrite the system model as

$$\mathbf{y} = \mathbf{h}_1 x_1 + \mathbf{h}_2 x_2 + \mathbf{z}$$
$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{z}$$

where $\mathbf{H} = [\mathbf{h}_1\ \mathbf{h}_2]$ and $\mathbf{x} = [x_1, x_2]^T$.

IV. Dual-stream Interference Cancellation Receiver Architecture

The key ingredient for allowing spatial reuse one in wireless networks is interference mitigation. In the mesh architecture described in the previous section, MRs that are between two CHs must be able to communicate concurrently with both CHs (one stream for each CH on transmit and receive) on the same time-frequency resources since the scheduling between clusters is not coordinated. In this section we describe two different dual-stream multi-antenna receiver structures that can be used for this purpose, namely a minimum mean squared error (MMSE) receiver (see Subsection IV-A) and a low-complexity max-log MAP receiver (see Subsection IV-B) [4]. Both receivers were first implemented in Matlab and later on the OpenAirInterface platform.

[2] *Notation.* Let $\mathbb{C}$ denote the set of complex numbers. Scalars are denoted by x. Column vectors and matrices are denoted by $\mathbf{a}$ and $\mathbf{A}$ and their elements are denoted by a_i and $A_{i,j}$ respectively. Transpose and Hermitian transpose are denoted by $\cdot^T$ and $\cdot^H$. $\mathbf{I}_M$ is the identity matrix of size M and $\mathbf{0}_M$ is an M-dimensional vector of zeros. The Euclidean (ℓ_2) norm of a vector $\mathbf{a}$ is denoted by $\|\mathbf{a}\|$ and the Frobenius norm of a matrix $\mathbf{A}$ is denoted by $\|\mathbf{A}\|_F$. Subscripts $(.)_R$ and $(.)_I$ indicate the real and imaginary parts.

A. MMSE Receiver

Detection based on MMSE equalization [7] involves applying a spatial filter $\mathbf{M}$ to the received signal vector $\mathbf{y}$, i.e., $\tilde{\mathbf{x}} = \mathbf{M}\mathbf{y}$ where $\tilde{\mathbf{x}}$ is the biased estimate of $\mathbf{x}$. Frequency domain MMSE filter $\mathbf{M}$ is given as

$$\mathbf{M} = \left(\sigma^2\mathbf{P}^{-1} + \mathbf{H}^H\mathbf{H}\right)^{-1}\mathbf{H}^H$$

where $\mathbf{P}$ is the diagonal power distribution matrix with the diagonal as $\left[\sigma_1^2, \sigma_2^2\right]$. It is followed by an unbiasing operation i.e. $\hat{\mathbf{x}} = \Gamma^{-1}\tilde{\mathbf{x}}$ where $\Gamma = \text{diag}\left(\mathbf{MH}\right)$. Post detection interference is assumed to be Gaussian which on one hand reduces the computational complexity but on the other adds to the suboptimality of MMSE detection. MMSE preprocessing decouples the spatial streams and the bit metric for the i-th bit for bit value b of the symbol x_k on k-th stream is given as

$$\lambda_k^i\left(\mathbf{y}, b\right) \approx \max_{x_k \in \chi_{k,b}^i}\left[-\frac{\gamma_k^2}{N_0}\left|\hat{x_k} - x_k\right|^2\right] \tag{3}$$

for $k = 1, 2$ where γ_k is the i-th diagonal element of Γ. $\chi_{k,b}^i$ denotes the subset of the signal set $x_k \in \chi_k$ whose labels have the value $b \in \{0, 1\}$ in the position i. Based on these bit metrics, bit log likelihood ratios (LLRs) are calculated which after de-interleaving are passed to the channel decoder.

Linear spatial MMSE filters are able to minimize the level of interference and are being considered as likely candidates for future wireless systems [8]. However, it is well known that MMSE is suboptimal for non Gaussian alphabets in low dimensional systems (low number of interferers) [9] and does not exploit the interference structure. Further, the implementation of the MMSE receiver on a fixed point processor is not trivial due to the necessary matrix inversion. The reason for this is high dynamic range of the determinant, especially in frequency selective channels. See [5] for more details.

B. Low-Complexity max-log MAP Detector

This detector is a low complexity version of the max-log MAP detector and is based on the matched filter outputs [4]. Its low complexity is based on the reduction of one complex dimension. This detector instead of attenuating the interference exploits its structure in mitigating the effect of interference. Without loss of generality, let's consider the first stream being the desired stream. The max-log MAP bit metric for the bit b of the desired stream x_1 is given as [10]

$$\begin{aligned}\lambda_1^i\left(\mathbf{y}, b\right) &\approx \min_{x_1 \in \chi_{1,b}^i, x_2 \in \chi_2}\left\|\mathbf{y} - \mathbf{h}_1 x_1 - \mathbf{h}_2 x_2\right\|^2 \\ &= \min_{x_1 \in \chi_{1,b}^i, x_2 \in \chi_2}\Big\{\left\|\mathbf{y}\right\|^2 + \left\|\mathbf{h}_1 x_1\right\|^2 - \left(2y_1 x_1^*\right)_R \\ &\quad + \left(2p_{12}x_1^* x_2\right)_R - \left(2y_2 x_2^*\right)_R + \left\|\mathbf{h}_2 x_2\right\|^2\Big\}\end{aligned} \tag{4}$$

where $y_1 = \mathbf{h}_1^H\mathbf{y}$ and $y_2 = \mathbf{h}_2^H\mathbf{y}$ are the matched filter outputs for the first and the second stream resp. and $p_{12} = \mathbf{h}_1^H\mathbf{h}_2$ is the cross correlation between the first and the second channel. Writing terms in their real and imaginary parts, we have

$$\begin{aligned}\lambda_1^i\left(\mathbf{y}, b\right) &\approx \min_{x_1 \in \chi_{1,b}^i, x_2 \in \chi_2}\Big\{\left\|\mathbf{h}_1 x_1\right\|^2 - \left(2y_1 x_1^*\right)_R \\ &+ \left(2\left(p_{12,R}x_{1,R} + p_{12,I}x_{1,I}\right) - 2y_{2,R}\right)x_{2,R} + \left\|\mathbf{h}_2\right\|^2 x_{2,R}^2 \\ &+ \left(2\left(p_{12,R}x_{1,I} - p_{12,I}x_{1,R}\right) - 2y_{2,I}\right)x_{2,I} + \left\|\mathbf{h}_2\right\|^2 x_{2,I}^2\Big\}.\end{aligned} \tag{5}$$

For x_2 belonging to the equal energy alphabets, the values of $x_{2,R}$ and $x_{2,I}$ which minimize (5) need to be in the opposite directions of $\left(2\left(p_{12,R}x_{1,R} + p_{12,I}x_{1,I}\right) - 2y_{2,R}\right)$ and $\left(2\left(p_{12,R}x_{1,I} - p_{12,I}x_{1,R}\right) - 2y_{2,I}\right)$ resp. thereby evading search on alphabets of x_2 and reducing one complex dimension of the system. The bit metric is therefore written as

$$\begin{aligned}\lambda_1^i\left(\mathbf{y}, b\right) &\approx \min_{x_1 \in \chi_{1,b}^i, x_2 \in \chi_2}\Big\{\left\|\mathbf{h}_1 x_1\right\|^2 - \left(2y_1 x_1^*\right)_R \\ &- \left|\left(2\left(p_{12,R}x_{1,R} + p_{12,I}x_{1,I}\right) - 2y_{2,R}\right)\right|\left|x_{2,R}\right| \\ &- \left|\left(2\left(p_{12,R}x_{1,I} - p_{12,I}x_{1,R}\right) - 2y_{2,I}\right)\right|\left|x_{2,I}\right|\Big\}\end{aligned} \tag{6}$$

For non equal energy alphabets, it is the minimization problem of a quadratic function again trimming one complex dimension of the system. In that case, the real and imaginary parts of x_2 which minimize (4) are given as

$$\begin{aligned}x_{2,R} &\rightarrow -\frac{\left(p_{12,R}x_{1,R} + p_{12,I}x_{1,I}\right) - y_{2,R}}{\left\|\mathbf{h}_2\right\|^2} \\ x_{2,I} &\rightarrow -\frac{\left(p_{12,R}x_{1,I} - p_{12,I}x_{1,R}\right) - y_{2,I}}{\left\|\mathbf{h}_2\right\|^2}\end{aligned} \tag{7}$$

where $\rightarrow$ indicates the quantization process in which amongst the finite available points, the point closest to the calculated continuous value is selected.

The reduced complexity max-log MAP detector requires also less operations than the MMSE receiver [3]. Furthermore it can be implemented without any division and therefore it is numerically more stable on a fixed point processor.

V. Experiments and Applications

In this section we investigate the performance of the two dual-stream receiver structures described in the previous section. Firstly, in Subsection V-A, we perform computer simulations using a simple synthetic channel model. Secondly, in Subsection V-B we present performance results from the real-time implementation on the OpenAirInterface platform. All performance comparisons (both for simulation and lab tests) were done using the broadcast channel (BCH) of the primary clusterhead (CH1) with interference from the BCH from the secondary clusterhead (CH2). The BCH uses QPSK modulation and rate 1/2 convolutional code with a block length of 1056 bits which corresponds to 8 OFDM symbols with 132 data subcarriers each.

A. Computer Simulations

1) Channel Model: For the simulations, the 2×2 MIMO channel matrices $\mathbf{H}_q^{(1)}$ and $\mathbf{H}_q^{(2)}$ are modeled as spatially white and independent. The channel is assumed to be constant during a block and varies independently between blocks. We use both a frequency flat fading model as well as a frequency

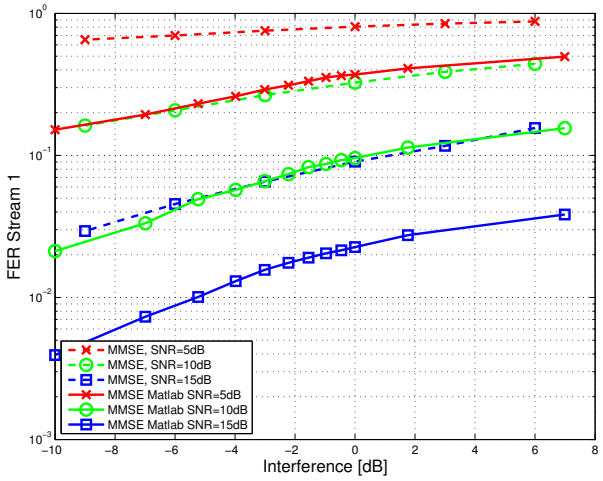

Fig. 2. FER of the first stream of the MMSE receiver (C and Matlab implementation) for a frequency flat Rayleigh fading channel. Results are plotted for different SNR levels of the first stream. The x axis denotes the interference of the second stream w.r.t the first stream.

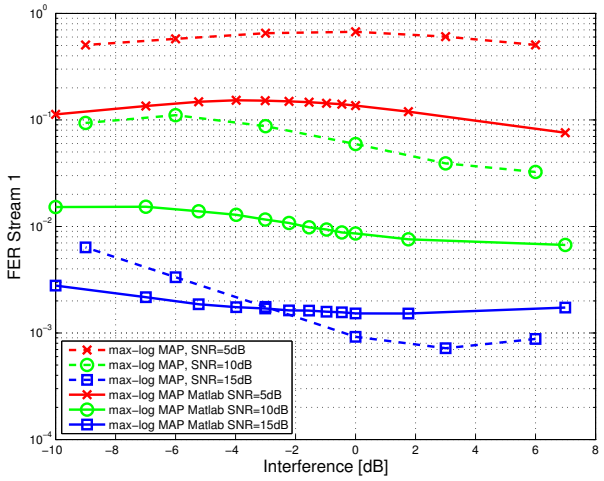

Fig. 3. FER of the first stream of the max-log MAP receiver (C and Matlab implementation) for a frequency flat Rayleigh fading channel. Results are plotted for different SNR levels of the first stream. The x axis denotes the interference of the second stream w.r.t the first stream.

selective model. In the frequency flat case the channel matrices stay constant over all subcarriers q with channel coefficients drawn from a Rayleigh distribution with unit variance. In the frequency selective case we model the channel as a tapped delay line with 8 sample-spaced taps with an exponential power delay profile. Each tap undergoes Rayleigh fading.

2) Simulation Results: The simulation model described above as well as the two receiver structures described in Section IV were implemented in Matlab as well as in fixed-point C. The C simulator however includes the FFT and CP insertion at the transmitter and the corresponding IFFT and CP removal at the receiver. The channel is thus simulated in the time domain rather than in the frequency domain. Also the C implementation does a full channel estimation while the Matlab implementation assumes perfect channel knowledge. We perform Monte Carlo simulations with both implementations and compare the frame error rates (FERs) of the first stream w.r.t. the interference from the second stream.

Fig. 2 shows the performance comparison of the two implementations of the MMSE receiver for different SNR levels. It can be seen that the fixed point C implementation looses approximately 5dB compared to the floating point implementation in Matlab. This underperformance is due to the loss of accuracy of the fixed point implementation and the channel estimation errors. Fig. 3 shows the performance comparison of the two implementations of the max-log MAP receiver for different SNR levels. It can be seen that in general the Matlab implementation performs better for low SNR levels (5 and 10dB). For an SNR level of 15dB the two implementations perform very similar. Compared to the MMSE implementation, the max-log MAP implementation is clearly preferable. Also note that the performance of the max-log MAP receiver actually gets better when the interference gets stronger. This improvement is attributed to the ability of max-log MAP receiver to exploit the interference structure in decoding the desired stream. Max-log MAP receiver has a coding gain as interference gets stronger in contrary to the MMSE receiver which has a coding loss [11].

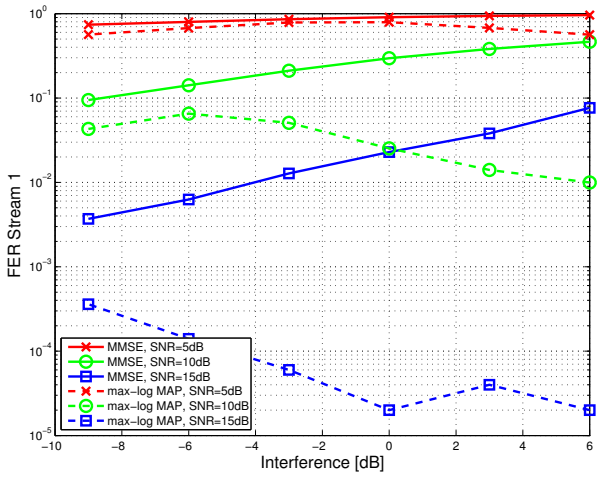

Fig. 4. FER of the first stream of the of the MMSE and max-log MAP receiver (C implementation) for a frequency selective Rayleigh fading channel. Results are plotted for different SNR levels of the first stream. The x axis denotes the interference of the second stream w.r.t. the first stream.

Fig. 4 shows the performance comparison of the C implementation of the MMSE as well as the max-log MAP receiver for different SNR levels in a frequency selective channel. It can be seen that the max-log MAP receiver profits most from the the additional frequency diversity of the channel, while the performance of the MMSE receiver hardly improves. The max-log MAP receiver has full diversity gain while MMSE receiver loses one order of diversity [11].

B. Lab Tests

1) Test Setup: The dual stream receiver was fully integrated in the software modem of the OpenAirInterface platform[3]. In order to evaluate its performance, it was also integrated in the Eurecom MIMO OpenAir Sounder (EMOS) [12]. The EMOS

[3]Note that the real-time system uses the same fixed point code for the receiver as the simulator.

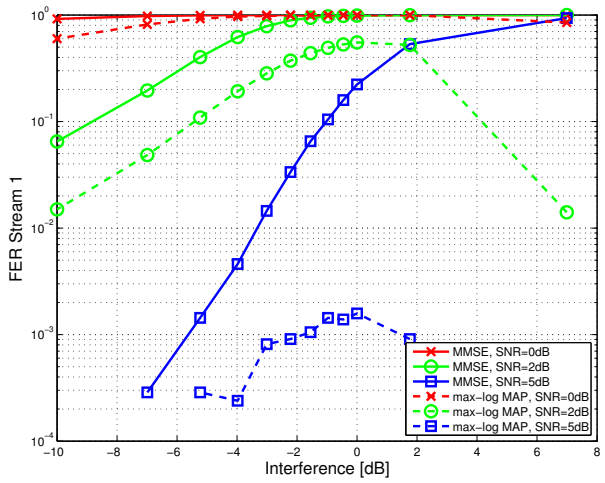

Fig. 5. Measured FER of the first stream using real-time MMSE and max-log MAP receiver in the Eurecom lab. Results are plotted for different power levels of the two clusterheads.

can be seen as a stand-alone version of the physical layer of the OpenAirInterface. Only the synchronization symbols (CHSCH, MRSCH) and the broadcast channels (CHBCH, MRBCH) are transmitted. Instead of the scheduled access channels, additional pilot symbols are transmitted for channel sounding purposes. At the receiver we record the FERs (based on the CRC check) of the CHBCH as well as the MIMO channel matrices from the channel estimation. The latter also allow a post performance evaluation that is usefull because the SNR and SIR can be better controlled.

2) Results: For the experiments we set up three nodes (CH1, CH2 and MR2) in in three different rooms (with a distance of about 10m). We carried out several experiments with different transmit powers of the two CHs to simulate different SNR and SIR levels at the receiver. The FERs that were recorded with the EMOS showed a very high volatility. While for some power settings the FER was 0%, for other settings it was very high. We believe that this is mostly due to the sensitivity of the receiver to the wireless channel, which cannot be controlled as nicely as in the simulations.

Therefore, for this paper we decided to show results using the recorded channel matrices. Here, the SNRs and the interference levels can be controlled just the same way as in the computer simulations of Section V-A. Fig. 5 shows the FER of the second stream for different SNR values for the max-log MAP and the MMSE respectively.

It can be seen that the basic trends are the same as in the computer simulations, in the sense that the MMSE receiver has in general a worse performance than the max-log MAP receiver and that its performance degrades as the interference increases. The max-log MAP receiver on the other hand can exploit the interference and the FER decreases again. However, the "turning point", where the FER decreases is further shifted to the right (at around 0 dB SIR compared to -6 dB). We attribute this phenomenon to the channel correlation, which was very strong in this measurement. We will investigate this issue further in future publications.

VI. Conclusions

In this paper we have shown the feasibility of distributed MIMO on the real-time open-source OpenAirInterface platform. It was shown that the reduced complexity max-log MAP detector has several advantages over the linear MMSE receiver. First of all its performance is much better (both diversity and coding gain) [11], especially when the interference level is high. Further it can be implemented without any divisions which is very advantageous on a fixed point processor. The implementation of the MMSE receiver on the other hand requires a matrix inversion, which is not trivial using fixed point arithmetic and looses a lot of accuracy. The receiver was also evaluated on the OpenAirInterface real-time testbed. It is difficult to produce FER curves for this setting since the distributed MIMO receiver is very sensitive to channel conditions. Even small displacements of the antennas can make a big difference in FER. In field trials we have seen that the best performance is achieved if the two transmitters have a line of sight to the receiver, but are well separated in space. Significant differences in the received powers from the two sources can also improve the performance (in case of the max-log MAP receiver).

References

[1] D. Gesbert, S. Kiani, A. Gjendemsj, and G. Øien, "Adaptation, coordination, and distributed resource allocation in interference-limited wireless networks," *Proceedings of the IEEE*, vol. 95, no. 12, pp. 2393–2409, Dec. 2007.

[2] M. Russell and G. Stuber, "Interchannel interference analysis of OFDM in a mobile environment," in *Proc. IEEE 45th Vehicular Technology Conference*, vol. 2, 1995, pp. 820–824.

[3] R. Ghaffar and R. Knopp, "Dual-antenna BICM reception with applications to MIMO broadcast and single frequency cellular system," in *Proc. IEEE Intl. Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, Cannes, France, Sep. 2008.

[4] ——, "Interference suppression for next generation wireless systems," in *Proc. IEEE Vehicular Technology Conference VTC-Spring*, Barcelona, Spain, Apr. 2009.

[5] F. Kaltenberger, R. Ghaffar, R. Knopp, H. Anouar, and C. Bonnet, "Openairinterface—an experimental platform for next generation wireless networks," *EURASIP Journal on Communications and Networking*, 2009, submitted.

[6] H. Anouar, C. Bonnet, F. Kaltenberger, and R. Knopp, "OpenAirMesh—an experimental platform for cooperative mesh networks," in *Proc. 1st COST2100 Workshop on MIMO and Cooperative Communications*, Trondheim, Norway, Jun. 2008.

[7] I. Medvedev, B. Bjerke, R. Walton, J. Ketchum, M. Wallace, and S. Howard, "A comparison of MIMO receiver structures for 802.11n WLAN - performance and complexity," in *Proc. IEEE 17th International Symposium on Personal, Indoor and Mobile Radio Communications*, Sept. 2006.

[8] E. Dahlman, H. Ekstrom, A. Furuskar, Y. Jading, J. Karlsson, M. Lundevall, and S. Parkvall, "The 3G long-term evolution - radio interface concepts and performance evaluation," in *Proc. IEEE 63rd Vehicular Technology Conference VTC-Spring*, vol. 1, May 2006, pp. 137–141.

[9] H. Poor and S. Verdu, "Probability of error in MMSE multiuser detection," *IEEE Trans. Inf. Theory*, vol. 43, no. 3, pp. 858–871, May 1997.

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

[11] R. Ghaffar and R. Knopp, "Spatial interference cancellation and pairwise error probability analysis," in *Proc. IEEE International Conference on Communications (ICC)*, Dresden, Germany, Jun. 2009.

[12] F. Kaltenberger, M. Kountouris, D. Gesbert, and R. Knopp, "On the trade-off between feedback and capacity in measured MU-MIMO channels," *IEEE Trans. Wireless Commun.*, 2009, in press.