

Design & Experimental Validation of MIMO Multiuser Detection for Downlink Packet Data

Dragan Samardzija, Angel Lozano and Constantinos Papadias*

November 14, 2005

Abstract

In single-user MIMO communication, the first-order throughput scaling is determined by the smallest of the number of transmit and receive antennas. This typically renders terminals the constraining bottleneck. In a multiuser downlink, this bottleneck can be bypassed by having the base station communicate with multiple terminals simultaneously, in which case the receive antennas at those terminals are effectively pooled in terms of the capacity scaling. This, however, requires that the base have instantaneous channel information. Without such information, the structure and statistics of the channel can be exploited to form multiple simultaneous beams towards the various users, but these beams are in general mutually interfering. This paper proposes the use of multiuser detection to discriminate the signals conveyed over interfering beams. This approach is formulated and experimentally evaluated on an HSDPA MIMO testbed that involves a commercial base station, multi-antenna terminals, and custom ASICs.

*The authors are with Bell Laboratories (Lucent Technologies), Holmdel, NJ07733, USA.

I Introduction

MIMO (multi-input multi-output) schemes utilizing multiple transmit and receive antennas are posed to be a major ingredient in the evolutionary process of wireless communication. Widely recognized features associated with MIMO are: spatial diversity, signal enhancements, interference mitigation and spatial multiplexing. The latter, in particular, has driven a lot of the research over the last decade, ever since it was shown in [1, 2] that—in adequate channel conditions—the ergodic capacity (in bits/s/Hz) of a MIMO link as function of the average SNR (signal-to-noise ratio) behaves as

$$C(\text{SNR}) = \min(n_T, n_R) \log_2 \text{SNR} + O(1) \quad (1)$$

where n_T and n_R denote the number of transmit and receive antennas, respectively. This linear scaling with the number of antennas is a powerful means to achieve high spectral utilization provided that antenna arrays can be effectively deployed.

In actual wireless systems, of course, links do not operate in isolation: each base station must actively communicate with a plurality of users and thus a number of MIMO links have to coexist. The behavior expressed by (1) can be immediately translated onto a multiuser environment by partitioning either time or frequency onto orthogonal sets, each of which is assigned to a particular user link. Focusing on the downlink, where n_T indicates the number of transmit antennas at the base station while n_R represents the number of receive antennas at the terminal, such orthogonal multiplexing incurs only a small loss in capacity if $n_R \gg n_T$. Unfortunately, the small size and cost sensitivity of mobile terminals precludes the deployment of a large number of antennas thereon and hence the most likely scenario for mobile systems corresponds with $n_T \geq n_R$. (In some cases, n_R may be tightly restricted to equal 1.) In these conditions, orthogonal multiplexing severely constrains the capacity.

Without the constraint of time/frequency multiplexing, a downlink with n_T antennas at the base and n_R antennas at each of K users can yield a sum capacity that behaves as

$$C(\text{SNR}) = \min(n_T, Kn_R) \log_2 \text{SNR} + O(1) \quad (2)$$

whereby the tight restrictions on n_R become immaterial and the burden of limiting the capacity shifts to the base station, where n_T can be more easily scaled. Unfortunately, achieving (2) requires that the base station have accurate and instantaneous information about the state of the fading channel between each of its antennas and those at each of the mobiles [3]–[6]. This represents a total of $n_T \times K \times n_R$ time-varying complex coefficients, whose instantaneous tracking by the base station is not a viable option in frequency-duplexed systems¹.

¹In time-duplexed systems, the reciprocity in the channel propagation characteristics makes it feasible

Without instantaneous channel state information at the transmitter, simultaneous transmission to multiple users becomes challenging. In these conditions, interestingly, antenna correlation—often detrimental in MIMO—facilitates the formation of beams that can be directed to individual users providing partial isolation between their simultaneous transmissions. Moreover, this approach (already recognized and incorporated into the UMTS evolutionary process [7, 8]) results in simple receiver structures. The degree of user isolation that can be attained through the composition of beams, however, is directly determined by the location of the users in the cell and by the characteristics of their propagation channels. Unless $K \ll n_T$, every beam will often illuminate users other than the intended one resulting in significant levels of multiuser interference.

In this paper, we formulate and experimentally evaluate a scheme that provides resilience to strong multiuser interference when multiple beams are simultaneously active. The cornerstone of this scheme is the recognition that well-known MUD (multiuser techniques) [9], developed originally for CDMA (code-division multiple access), can be applied to the mitigation and removal of spatial interference. This represents, to some extent, an abandonment of the idea of simple and basically passive terminals and an embracing of the concept of smart terminals that actively participate in the task of discriminating transmissions to the different users. This conceptual shift is grounded on the rapid improvement in processing power that stems from Moore's law.

In the remainder of the paper, we describe the MIMO-MUD scheme and we quantify its benefits through a series of experiments executed on a testbed that involves a commercial base station equipped with multiple antennas, terminals also equipped with multiple antennas, and specially designed MIMO ASICs (application specific integrated circuits). In order to render the experiments specific, the testbed is set up to comply with the HSDPA (high-speed downlink packet access) channel, which is foreseen to become the main vehicle for the provision of packet data in UMTS. To the best of our knowledge, these are the first such reported experiments.

The paper is organized as follows. Having justified the interest in the simultaneous transmission to multiple users through parallel beams, in section II we review several well-known MUD approaches and discuss their applicability to the problem of mitigating the effects of multiuser interference across beams. In Section III, in turn, we briefly describe the key features of UMTS HSDPA and describe in detail its implementation on the MIMO testbed platform. Finally, Section IV lays down a number of experimental results that validate the applicability of the chosen MIMO-MUD approach.

to track these coefficients instantaneously as long as the Doppler spread is small enough. Note, however, that reciprocity does not necessarily apply to the transceivers and thus careful calibration may be required. Note also that the wide-area communication marketplace is currently dominated by frequency-duplexed systems.

II MIMO-MUD: Formulation

Although not a requirement when MUD is used, we shall limit the number of active users to be $K \leq n_T$, which allows for the generation of beams that are orthogonal in origin [10]. Larger numbers of users can of course be accommodated via time/frequency-multiplexing.

The baseband complex linear model describing the communication between the base station and the k -th terminal, $k \in \{1, \dots, K\}$, is

$$\mathbf{y}_k = \mathbf{H}_k \mathbf{x} + \mathbf{n}_k \quad (3)$$

where \mathbf{y}_k is the $(n_R \times 1)$ vector received by terminal k , \mathbf{n}_k is the corresponding additive white Gaussian noise vector with one-sided spectral density

$$N_0 = \frac{E[\|\mathbf{n}\|^2]}{n_R}$$

and \mathbf{H}_k is the $(n_R \times n_T)$ channel random matrix whose (i, j) -th entry represents the transfer coefficient² between the j -th base transmit antenna and the i -th receive antenna at terminal k . In turn, \mathbf{x} is the $(n_T \times 1)$ transmit vector, common to all users and structured as

$$\mathbf{x} = \sum_{\ell=1}^K \mathbf{w}_\ell s_\ell \quad (4)$$

where s_ℓ is the information-bearing signal intended for terminal ℓ while the vector \mathbf{w}_ℓ contains the set of deterministic coefficients that, applied to each of the transmit antennas, generate the corresponding beam. Without loss of generality, the \mathbf{w}_ℓ 's are chosen such that $\|\mathbf{w}_\ell\|=1$, $\ell \in \{1, \dots, K\}$, and the power radiated for user ℓ is then $P_\ell = E[|s_\ell|^2]$. It is important to point out that, without instantaneous channel information at the transmitter, the coefficients in the set of vectors \mathbf{w}_ℓ , $\ell \in \{1, \dots, K\}$, cannot depend on the random matrices \mathbf{H}_ℓ , $\ell \in \{1, \dots, K\}$, but only on their distributions.

From the standpoint of user k , we can conveniently rearrange (3) and (4) as

$$\mathbf{y}_k = \underbrace{\mathbf{H}_k \mathbf{w}_k s_k}_{\substack{\text{signal intended} \\ \text{for terminal } k}} + \underbrace{\sum_{\ell \neq k} \mathbf{H}_k \mathbf{w}_\ell s_\ell}_{\text{interference}} + \underbrace{\mathbf{n}_k}_{\text{noise}} \quad (5)$$

where the interference corresponds to the signals that are being beamed towards terminals other than k , orthogonal in origin but—in general—not upon reception because of

²In order to focus on the spatial processing aspects, the channel fading is modelled as frequency-flat. The formulation, nonetheless, can be extended to frequency-selective fading.

the random matrix \mathbf{H}_k whose realization is unknown to the transmitter. The realization of \mathbf{H}_k , in contrast, is considered known to the receiver, which may estimate it provided, for example, that each individual beam is associated with a unique pilot. Multiple secondary pilots are already supported in UMTS [8] and are expected to be equally available in future system designs. More specifically, this enables receiver k to estimate the effective channels $\mathbf{H}_k \mathbf{w}_\ell$ for $\ell \in \{1, \dots, K\}$.

There are several manners in which the presence of the interference can be addressed:

- The simplest approach is to ignore the interference by matching the receiver at terminal k to the effective channel for its desired signal generating the decision statistic $(\mathbf{w}_k^\dagger \mathbf{H}_k^\dagger \mathbf{y}_k)$, which exhibits an average signal-to-interference and noise ratio [9]

$$\text{SINR}_k = E \left[\frac{P_k (\mathbf{w}_k^\dagger \mathbf{H}_k^\dagger \mathbf{H}_k \mathbf{w}_k)^2}{N_0 \mathbf{w}_k^\dagger \mathbf{H}_k^\dagger \mathbf{H}_k \mathbf{w}_k + \mathbf{w}_k^\dagger \mathbf{H}_k^\dagger \mathbf{H}_k \left(\sum_{\ell \neq k} P_\ell \mathbf{w}_\ell \mathbf{w}_\ell^\dagger \right) \mathbf{H}_k^\dagger \mathbf{H}_k \mathbf{w}_k} \right]$$

which depends strongly on the structure of \mathbf{H}_k . We shall use this SUMF (single-user matched filter) receiver as a baseline for later comparisons.

- A more robust approach consists of mitigating the interference through MMSE (minimum mean-square error) linear processing, which exploits the information provided by the conditional interference covariance

$$\Phi = N_0 \mathbf{I} + \sum_{\ell \neq k} P_\ell \mathbf{H}_k \mathbf{w}_\ell \mathbf{w}_\ell^\dagger \mathbf{H}_k^\dagger \quad (6)$$

The resulting average SINR at terminal k is [9]

$$\text{SINR}_k = P_k \mathbf{w}_k^\dagger E \left[\mathbf{H}_k^\dagger \Phi^{-1} \mathbf{H}_k \right] \mathbf{w}_k$$

which must lie within

$$\frac{P_k}{\sum_{\ell \neq k} P_\ell} \leq \text{SINR}_k \leq \frac{P_k \mathbf{w}_k^\dagger E[\mathbf{H}_k^\dagger \mathbf{H}_k] \mathbf{w}_k}{N_0} \quad (7)$$

The lower bound in (7) corresponds to an interference-limited situation with \mathbf{H}_k having independent entries, in which case the use of beams provides no significant advantage over time/frequency multiplexing. The upper bound, on the other hand, corresponds to a highly structured channel allowing for the formation of beams that remain essentially orthogonal regardless of the realization of \mathbf{H}_k , in which case terminal k receives no interference from any of the beams directed to other users.

- The most ambitious approach, and the one embraced in the remainder of the paper, is based on the joint detection of the signals transmitted on all beams, of which only the intended one is decoded and passed on to the higher layers while the remaining ones are simply discarded. In this case, the average SNR per receive antenna at terminal k is simply

$$\begin{aligned} \text{SNR}_k &= \frac{E[|\mathbf{H}_k \mathbf{x}|^2]}{E[|\mathbf{n}_k|^2]} \\ &= \frac{\sum_{\ell} \mathbf{w}_{\ell}^{\dagger} E[\mathbf{H}_k^{\dagger} \mathbf{H}_k] \mathbf{w}_{\ell}}{N_0 n_R} \end{aligned}$$

More specifically, the MIMO MUD solution that we propose relies on terminal k using its knowledge of $\mathbf{H}_k \mathbf{w}_{\ell} \forall \ell$ to perform ML (maximum likelihood) detection as

$$\{\hat{s}_1, \dots, \hat{s}_K\} = \arg \min \|\mathbf{y}_k - \mathbf{H}_k \sum_{\ell} \mathbf{w}_{\ell} \hat{s}_{\ell}\|^2 \quad (8)$$

where \hat{s}_k is the estimate of the signal s_k , retained and processed, while $\hat{s}_{\ell}, \ell \neq k$, are the signals intended for other users, discarded after detection.

III High-Speed Downlink Packet Access MIMO Testbed

A High-Speed Downlink Packet Access

For delay-tolerant data traffic, upcoming releases of UMTS will allocate a fraction of the power and code space to HSDPA, whose main features are:

- *Time multiplexing.* Users are time-multiplexed in short frames.
- *Multicode signaling.* The entire HSDPA code space is assigned to the active user. Thus, the transmit signal consists of a superposition of orthogonal codes.
- *No power control.* Power control is disabled.
- *Link adaptation.* The transmit rate is adapted based on feedback from the terminals.
- *Hybrid ARQ.* The link-layer automatic repeat request (ARQ) mechanism is combined with the physical-layer forward error correction [11].

With the incorporation of MIMO, the possibility of having active users on separate beams is enabled and, correspondingly, the use of MIMO-MUD becomes alluring. In the remaining, we validate this idea using a 5-MHz MIMO testbed that operates at 2.1 GHz and supports $n_T=4$ antennas at the base and $n_R=4$ antennas at each terminal [12]. The testbed is currently compliant with the above-described HSDPA features.

In order to test the MIMO-MUD concept under the harshest conditions, trivial beams are employed: each w_ℓ is identically zero except for the ℓ -th entry, which is set to 1. With that, the beams give rise to severe interference as no attempt is made to isolate the transmission to different users.

B Transmitter Implementation

At the base, omnidirectional vertically polarized $1/4$ -wavelength antennas are set 4 wavelengths apart along a line at a height of about 3 m. As shown in Fig. 1, the transmitter is mounted on a prototype Lucent base station (OneBTSTM prototype). A prototype mezzanine board is used to implement the physical and MAC (medium access control) layers. The rest of the base station, including RF (radio frequency) front-end, backplane and network interface, is also used. The RF front-end, in particular, meets the EVM (error vector magnitude) requirements set by Release 5 of the UMTS specifications.

A FPGA (field programmable gate array) is used to implement the multi-antenna physical layer transmitter. The corresponding functional block scheme is depicted in Fig. 2 with each functional block being HSDPA compliant. Up to 4 independent data streams d_ℓ , $\ell \in \{1, 2, 3, 4\}$, are passed down from the MAC layer, each intended for a distinct user. After being independently processed, every stream is radiated out of one of the antennas with a 24-bit CRC (cyclic redundancy code) word appended to each data block. These data blocks are encoded using a rate- $1/3$ turbo code and the desired transmission data rate is realized via a rate matching procedure that performs either puncturing or repetition of the encoder outputs. Binary words are then mapped to a particular QAM constellation (both QPSK and 16-QAM are supported by HSDPA) and then assigned to specific length-16 orthogonal channelization (i.e., spreading) codes. In addition, a unique pilot drawn from a set of secondary UMTS pilots [8] is assigned to each transmit antenna. The pilots are mutually orthogonal and orthogonal to the data-carrying spreading codes. The pilot power is set to 10% of the total radiated power [13]. The same scrambling code is used at every transmit antenna and the primary and secondary synchronization channels are also transmitted allowing mobile terminals to achieve chip-level, slot-level and frame-level synchronization and to perform cell search procedures. The above functional blocks (in Fig. 2) are implemented on FPGA Xilinx Virtex II 6000, with the clock rate of 61.44 MHz using approximately 15% of the available logic (i.e., logic slices), for each user (i.e.,

transmitted stream). Furthermore, approximately 100 KByte of memory is used per user.

The rate controller in Fig. 2 is closely coupled with the multiuser scheduling that is executed at the MAC layer. Specifically, the rate controller is responsible for setting, for each 2-msec time transmission interval, the rate matching parameters, modulation (QPSK or 16-QAM) and number of active spreading codes. Effectively, it optimizes the transmission data rates for a given channel and data traffic conditions. For the experimental results presented in Section IV, QPSK modulation was used with rate-1/2 coding and 10 length-16 active spreading codes.

To support multiple users, the MAC layer is implemented on a processor platform. Specifically, the multiuser scheduling and Hybrid-ARQ are implemented on a digital signal processor (Texas Instruments DSP 6701), while interfacing to an IP (internet protocol) network is implemented on an embedded processor (Motorola PowerPC 8260). The standard HSDPA specifications are retained at the MAC layer and thus only the physical layer is aware of the presence of MIMO.

C Receiver Implementation

As shown in Fig. 3, the terminal antennas are low-profile bow-tie printed dipoles with alternating 45° polarizations occupying vertexes of an $8.2'' \times 5.2''$ rectangle with the entire array fitting on the back of a laptop. Note that the fifth antenna, which is placed in the center of the rectangle (in Fig. 3), is vertically polarized and is used for the uplink transmission (a conventional single transmit antenna uplink is used). Physically different downlink and uplink antennas are used to simplify the design by avoiding implementation of an analog antenna coupler (which is otherwise needed when the same antenna is used both for the uplink and downlink, simultaneously).

The functional block scheme of the multi-antenna physical layer receiver is illustrated in Fig. 4, where \hat{d}_k is the estimate of the transmitted data for terminal k . After the AD (analog-to-digital) conversion, the received signal is sent to the MIMO-MUD ASIC which, in turn, outputs LLRs (log-likelihood ratios) that are then fed to the rate dematcher. After dematching, the turbo-decoder ASIC performs iterative decoding. The physical layer is implemented on 3 interconnected printed circuit cards that are next described in more detail.

Card 1 implements the analog RF front-end outputting up to 4 complex baseband signals, each corresponding to a receive antenna. A heterodyne receiver with a 140-MHz IF and noise figure under 8 dB is utilized, after which 10-bit AD conversion takes place.

Card 2 contains the basic processing elements of the multi-antenna receiver: (i) MIMO-

MUD ASIC ³ [14] and (ii) turbo decoder ASIC [15]. A block scheme of the MIMO detector is given in Fig. 5. The detector is based on a bank of despreaders matched to the data-carrying spreading codes, whose outputs are fed into the ML detection that corresponds with (8). In Fig. 5, \hat{s}_{lk} , $l \in \{1, \dots, L\}$, is the estimate of the transmitted symbol corresponding to code l for user k . Furthermore, each LLR corresponds to 1 channel bit with an 8-bit resolution.

An estimate of the MIMO channel, essential to the detection process, is obtained from an on-chip estimator. This is illustrated in Fig. 6, where \hat{h}_{ij} denotes the estimate of the (i, j) -th entry of the MIMO channel matrix \mathbf{H}_k . The on-chip estimator is based on a bank of despreaders corresponding to each of the length-512 pilot codes. In the case of frequency-flat fading, the presented estimator results in an ML channel estimate (see [13] and references therein). To lower the estimation noise, an optional integrator with forgetting factor α is available. For the experimental results in Section IV, $\alpha=0$.

Card 3, finally, holds the FPGA that acts as interconnect matrix between ADs, MIMO-MUD ASIC and turbo decoder ASIC. Furthermore, it executes (i) synchronization, (ii) frequency offset compensation, (iii) physical channel desegmentation (iv) rate dematching, (v) CRC check, and numerous auxiliary functions. To all of these functions, the use of MIMO is immaterial. The above functions are implemented on FPGA Xilinx Virtex II 6000, with the clock rate of 61.44 MHz using approximately 25% of the available logic and 70 KByte of memory.

IV MIMO-MUD Experimental Results

Indoor over-the-air measurements, mostly in static conditions, were carried out in a laboratory/office environment. The receiver was placed at various locations in the room. QPSK modulation was used with rate-1/2 coding and 10 length-16 orthogonal codes. The measurements include thermal as well as quantization noise.

We measured FERs (frame error rates) with the 2-msec time transmission interval specified for HSDPA. Based on the FER and on the 3.84-MHz chip rate, the throughput T is obtained as

$$T = 3.84 n_T \frac{10}{16} (1 - \text{FER}) \text{ [Mbps]}$$

This corresponds to a system with ARQ where the frames in error are discarded.

Fig. 7 presents the measured CDF (cumulative distribution) of T for a transmit power of 0 dBm (1 mW) over 30 locations. We show, for $n_T=4$ and $K=4$ terminals, each with $n_R=1$,

³The MIMO-MUD ASIC is manufactured using 0.18-micron CMOS technology, with 438000 gates, 300 mW core power, and size of 3.7 mm \times 3.7 mm.

a comparison of SUMF and MIMO-MUD receivers. Also depicted is the throughput for $K=1$ and $n_R=1$, for which the SUMF is optimal. Notice the large gains arising from the use of multiple antenna transmission. It is also worth noticing the value of multi-user detection alone, which for example leads to a throughput increase of more than 0.5 Mbps (at the 50% percentile point) compared to the corresponding single-user optimal transceiver. Fig. 8 presents the average throughput for different transmit power levels. MIMO-MUD results, at high transmit powers, in an almost 4-fold increase in average throughput. It should be noted that a higher order constellation could be used to combat the flooring effect shown in the figure for the single user transceivers. However, multi-user detection would still offer some gains, as evidenced by the fact that it has superior performance even before flooring starts to occur (e.g. at 0 dBm transmit power). Fig. 9 presents the average throughput for $n_R=1, 2, 3, 4$ with $n_T=4$ and with 0 dBm (solid line) and 10 dBm (dashed line) transmit powers. Fig. 10 presents corresponding results for $n_T=2$. In both Fig. 9 and Fig. 10 we see a sizeable improvement in throughput associated with the use of MIMO-MUD, especially when n_T is larger than or comparable to n_R . Although, for higher n_R , the SUMF approaches the MIMO-MUD throughput, this is in part an artifact of the fact that only QPSK is used. With 16-QAM available, we expect the MIMO-MUD advantage to be largely sustained.

In order to further demonstrate the capabilities of our HSDPA MIMO prototype, we also implemented a video-streaming application (using the Real-Time Streaming Protocol). Video streaming rates of up to 2 Mbps were achieved over the air (the higher layer ARQ introduced only a slight reduction in the overall throughput). In terms of interference mitigation performance, when using real-time video as each user's signal, MUD at the receiver performed very closely to the predicted behaviour and managed to separate the interfering video signals without any perceived degradation of performance as compared to each user's video stream transmitted alone. More information about these experiments can be found in [16].

V Conclusions

Multuser detection is a natural approach to signal detection in multuser environments. Although much of the developments in this area have been motivated by CDMA, multuser techniques are equally well suited to the spatial processing that arises with the use of MIMO, where the role of the CDMA spreading sequences is played by the fading coefficients between the various transmit and receive antennas.

In this paper, we have applied MUD to the detection of mutually interfering downlink beam transmissions aimed at different terminals. Without instantaneous channel state information at the base, these beams cannot be rendered orthogonal at the terminal re-

ceivers. Rather than simply enduring their mutual interference, we have proposed to jointly detect the signals transmitted on the intended and unintended beams.

Besides formulating such MIMO-MUD reception, we have experimentally validated the approach using a testbed that includes a commercial multi-antenna base station, multi-antenna terminals and custom MIMO ASICs. The results confirm the power of MUD, especially when the number of receive antennas at each terminal does not exceed the number of transmit antennas at the base.

Besides the application that has constituted the focus of the paper, MIMO-MUD schemes carry over to other multiuser MIMO settings. If, instead of parallel beams, time/frequency multiplexing is utilized, MIMO-MUD can be applied to mitigate the impact of interference from neighboring co-channel base stations. Although, in this case, individualizing the channel estimate for each interfering base station may not always be feasible, joint detection of desired and undesired transmissions can be applied to a few dominant neighbors. Furthermore, simpler linear MMSE processing can be applied if only the aggregate interference covariance in (6) can be estimated. Theoretical assessments of the advantage associated with knowledge of such covariance in MIMO communication can be found in [17]–[19]. Actually, even before the advent of MIMO systems, earlier pioneering contributions had already demonstrated the interference suppression capability of multiple receive antennas [10, 20, 21].

Acknowledgments

The authors gratefully acknowledge the support and encouragement of many colleagues led by Theodore Sizer, Reinaldo Valenzuela and Stephen Wilkus, all from Lucent Technologies. The authors are particularly grateful to Peter Bosch, Sape Mullender, Susan Walker, Tran Cuong, Francis Mullany, Eric Beck, Arnold Siegel, Thomas Gvoth and Ilya Korisch for their support. Part of this work was done under the IST project FITNESS, sponsored and funded by the FP5 European research framework.

References

- [1] G. J. Foschini and M. J. Gans, "On the limits of wireless communications in a fading environment when using multiple antennas," *Wireless Personal Communications*, pp. 315–335, 1998.
- [2] I. E. Telatar, "Capacity of multi-antenna Gaussian channels," *Eur. Trans. Telecom*, vol. 10, pp. 585–595, Nov. 1999.
- [3] W. Yo and J. M. Cioffi, "Sum capacity of a gaussian vector broadcast channel," *IEEE Intern. Symp. on Inform. Theory*, p. 498, July 2002.
- [4] G. Caire and S. Shamai, "On achievable throughput of a multiantenna Gaussian broadcast channel," *IEEE Trans. on Inform. Theory*, vol. 49, no. 7, pp. 1691–1706, July 2003.
- [5] P. Viswanath and D. N. C. Tse, "Sum capacity of vector gaussian broadcast channel and uplink-downlink duality," *IEEE Trans. on Inform. Theory*, vol. 49, no. 8, pp. 1912–1921, Aug. 2003.
- [6] S. Vishwanath, N. Jindal, and A. Goldsmith, "On the capacity of multiple input multiple output broadcast channel," *IEEE Intern. Conf. on Communications*, vol. 3, pp. 1444–1450, May 2002.
- [7] K.I. Pedersen, P.E. Mogensen, and J. Ramiro-Moreno, "Application and performance of downlink beamforming techniques in UMTS," *IEEE Comm. Magazine*, pp. 134–143, Oct. 2003.
- [8] "Beamforming enhancements," *TR 25.887 v. 1.3.0, 3rd Generation Partnership Project*, Oct. 2002.
- [9] S. Verdú, *Multiuser Detection*, Cambridge University Press, 1998.
- [10] J. H. Winters, "On the capacity of radio communication systems with diversity in a Rayleigh fading environment," *IEEE J. Select. Areas Commun.*, vol. 5, no. 5, pp. 871–878, June 1987.
- [11] Q. Zhang and S. A. Kassam, "Hybrid ARQ with selective combining for fading channels," *IEEE J. Select. Areas Commun.*, vol. 17, pp. 867–880, May 1999.
- [12] A. Burg, E. Beck, D. Samardzija, M. Rupp, and et al., "Prototype experience for MIMO BLAST over third generation wireless system," *IEEE JSAC Special Issue on MIMO Systems and Applications*, vol. 21, no. 3, April 2003.
- [13] D. Samardzija and N. Mandayam, "Pilot assisted estimation of MIMO fading channel response and achievable data rates," *IEEE Transactions on Signal Processing, Special Issue on MIMO*, vol. 51, no. 11, pp. 2882–2890, November 2003.
- [14] D. C. Garrett, L. M. Davis, and G. K. Woodward, "19.2 Mbit/s 4 x 4 BLAST/MIMO detector with soft ML outputs," *IEE Electr. Letters*, vol. 39, no. 2, pp. 233–235, Jan. 2003.
- [15] M. Bickerstaff, L. Davis, C. Thomas, D. Garrett, and C. Nicol, "A 24 Mb/s radix-4 turbo decoder," *IEEE Intern. Solid-State Circuits Conf.*, Feb. 2003.

- [16] IST-FITNESS D3.3, "Description of UMTS MTMR reconfigurability demo," www.ist-fitness.org.
- [17] A. Lozano and A. M. Tulino, "Capacity of multiple-transmit multiple-receive antenna architectures," *IEEE Trans. on Inform. Theory*, vol. 48, Dec. 2002.
- [18] A. Lozano, A. M. Tulino, and S. Verdú, "Multiple-antenna capacity in the low-power regime," *IEEE Trans. on Inform. Theory*, vol. 49, no. 10, Oct. 2003.
- [19] A. L. Moustakas, S. H. Simon, and A. M. Sengupta, "MIMO capacity through correlated channels in the presence of correlated interferers and noise: a (not so) large N analysis," *IEEE Trans. on Inform. Theory*, vol. 49, no. 10, pp. 2545–2561, Oct. 2003.
- [20] J. H. Winters, J. Salz, and R. D. Gitlin, "The impact of antenna diversity on the capacity of wireless communication systems," *IEEE Trans. Commun.*, vol. 42, no. 2/3/4, pp. 1740–1751, Feb./March/Apr. 1994.
- [21] J. H. Winters, "Optimum combining in digital mobile radio with cochannel interference," *IEEE J. Select. Areas Commun.*, vol. 2, no. 4, pp. 528–539, July 1984.



Figure 1: Multi-antenna base station.

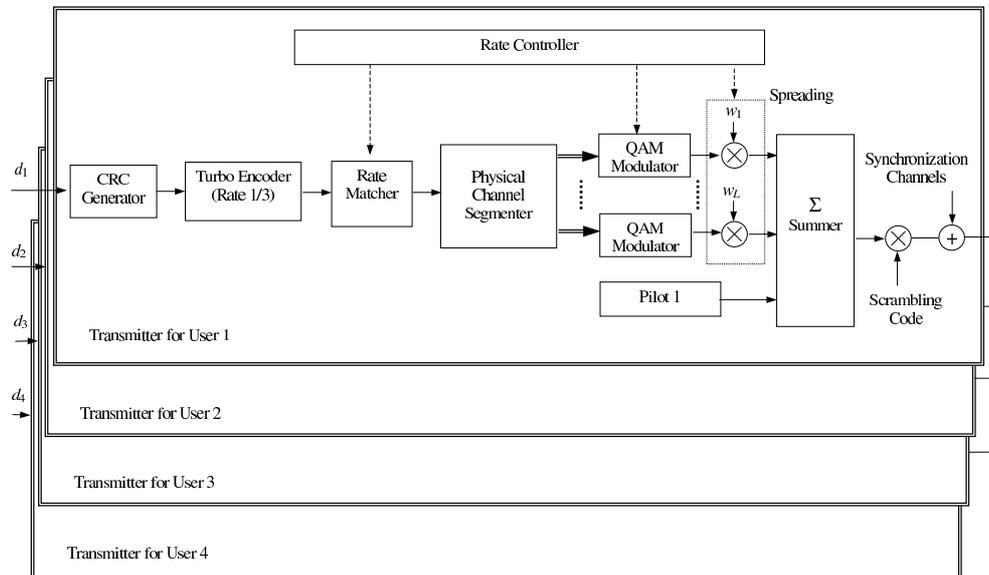


Figure 2: Functional block scheme of multi-antenna HSDPA physical layer transmitter.

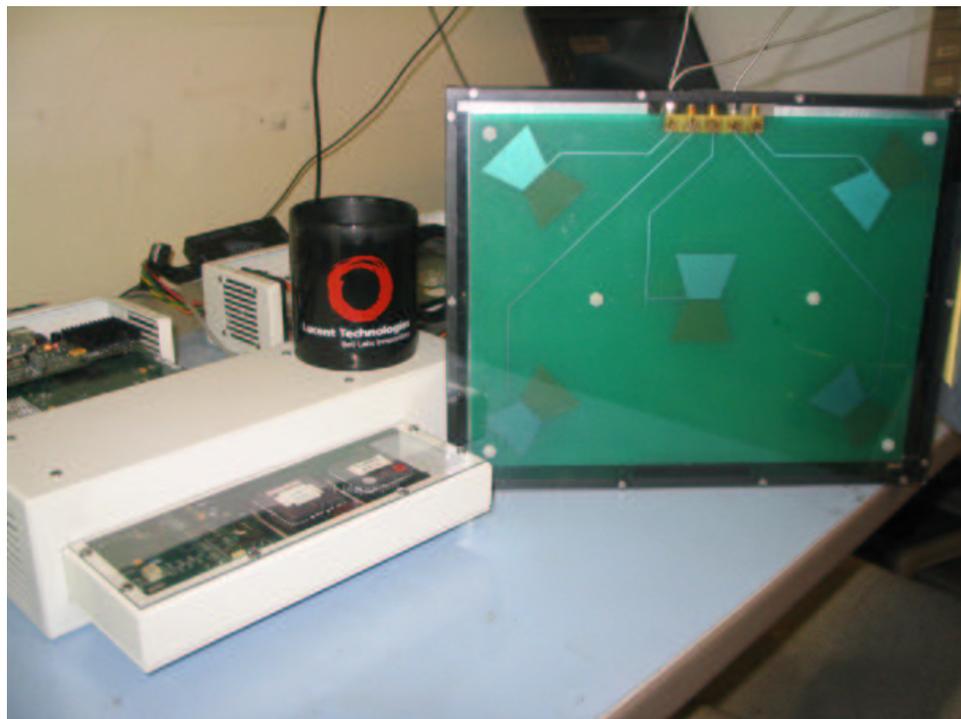


Figure 3: Terminal and receive antenna array.

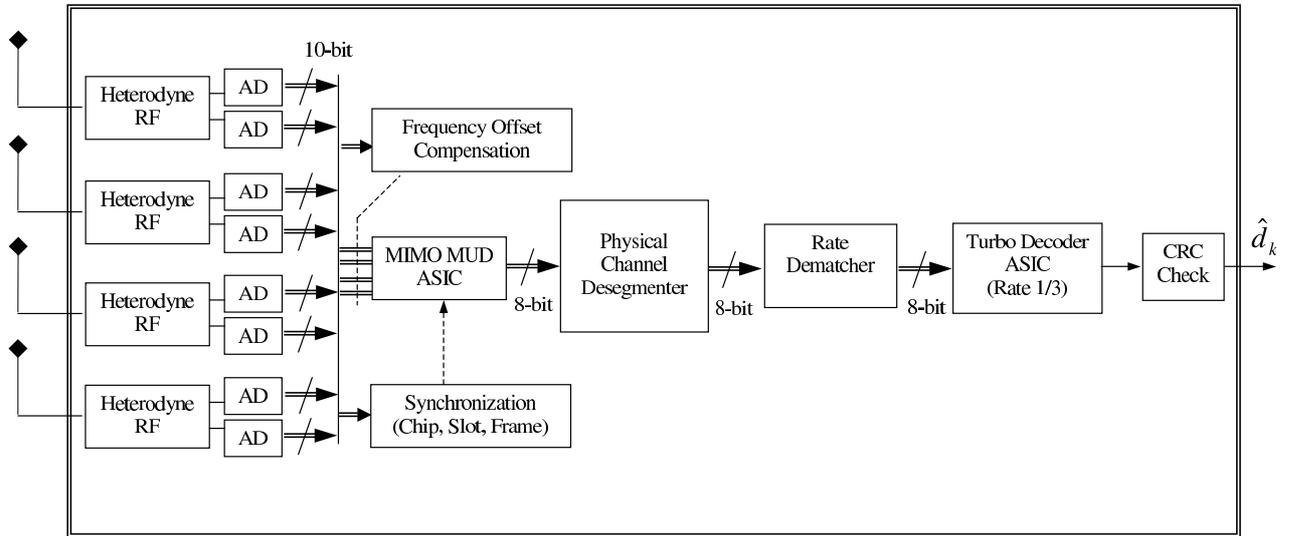


Figure 4: Functional block diagram of multi-antenna HSDPA physical layer receiver.

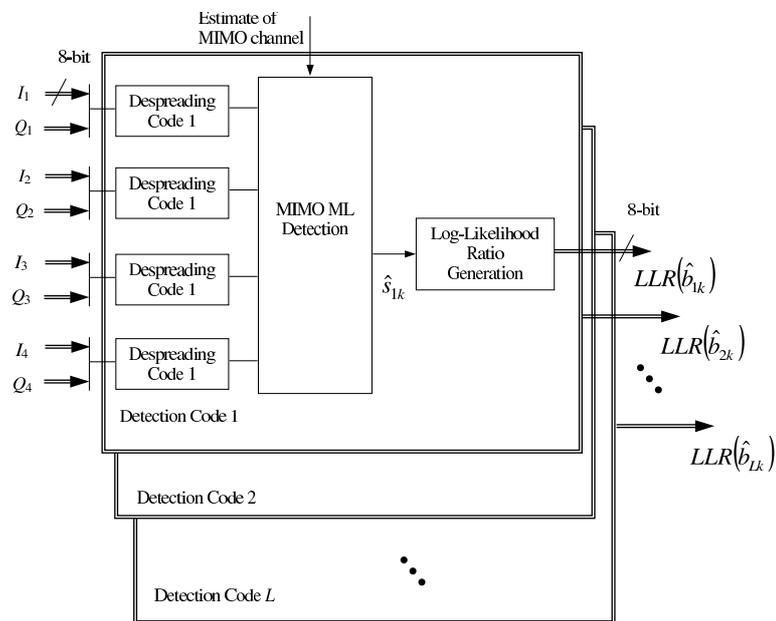


Figure 5: MIMO detection as implemented on the MIMO-MUD ASIC.

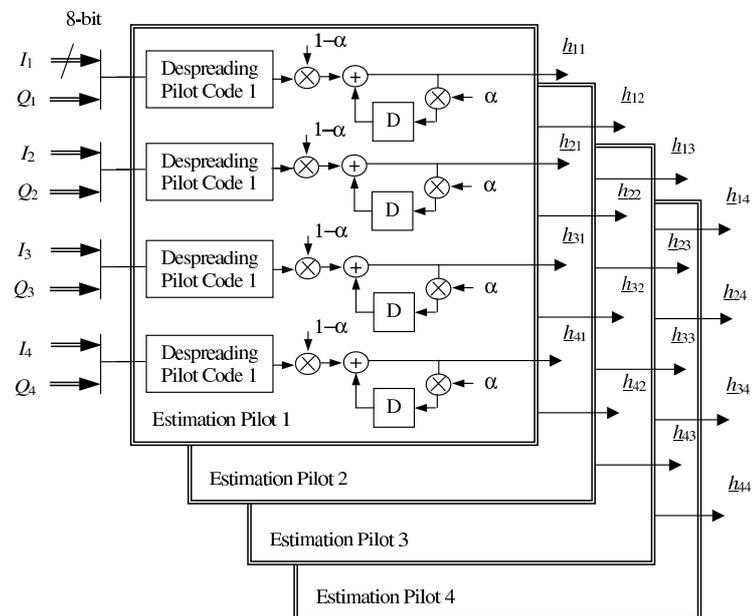


Figure 6: MIMO channel estimation as implemented on the MIMO-MUD ASIC.

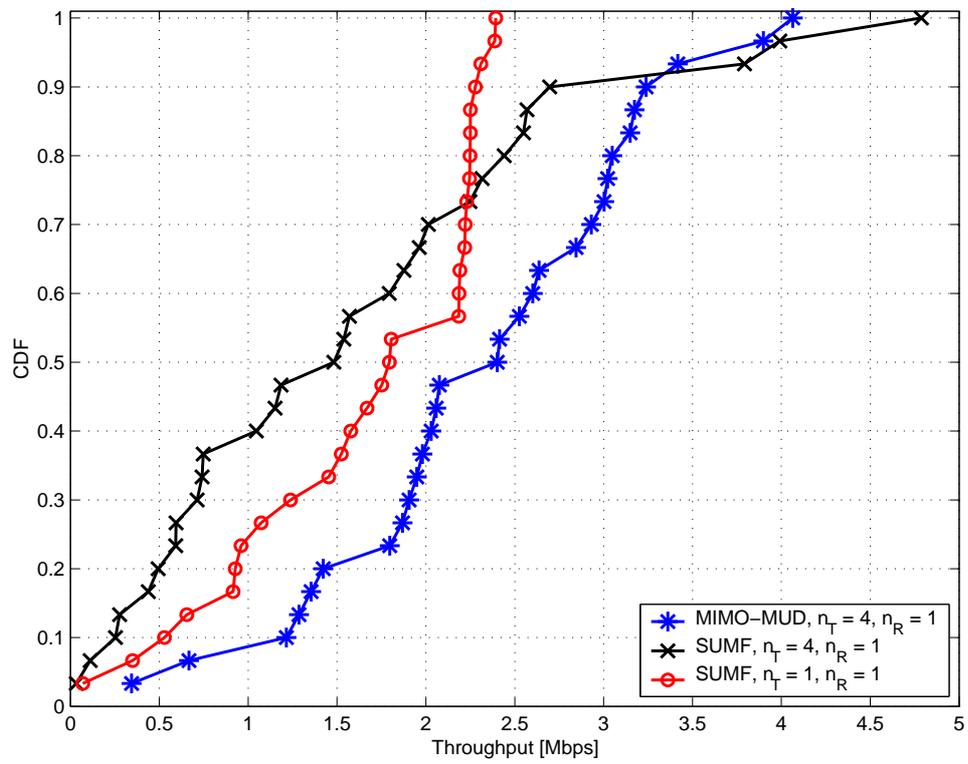


Figure 7: Measured CDF of throughput for 0 dBm over 30 locations.

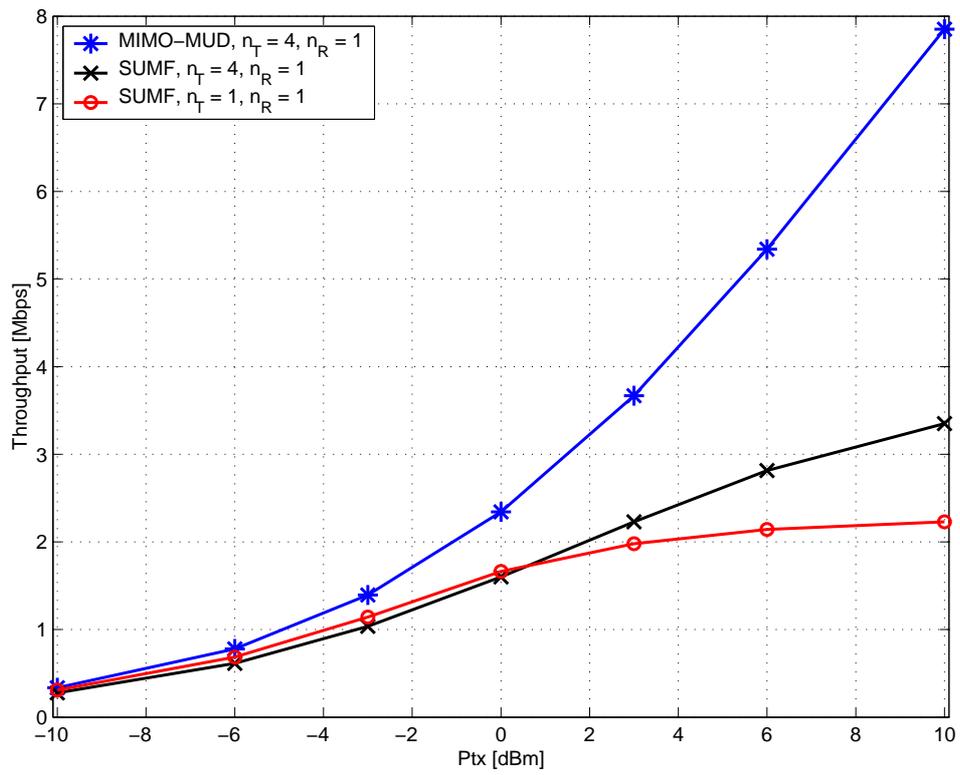


Figure 8: Measured average throughput vs. transmit power over 30 locations.

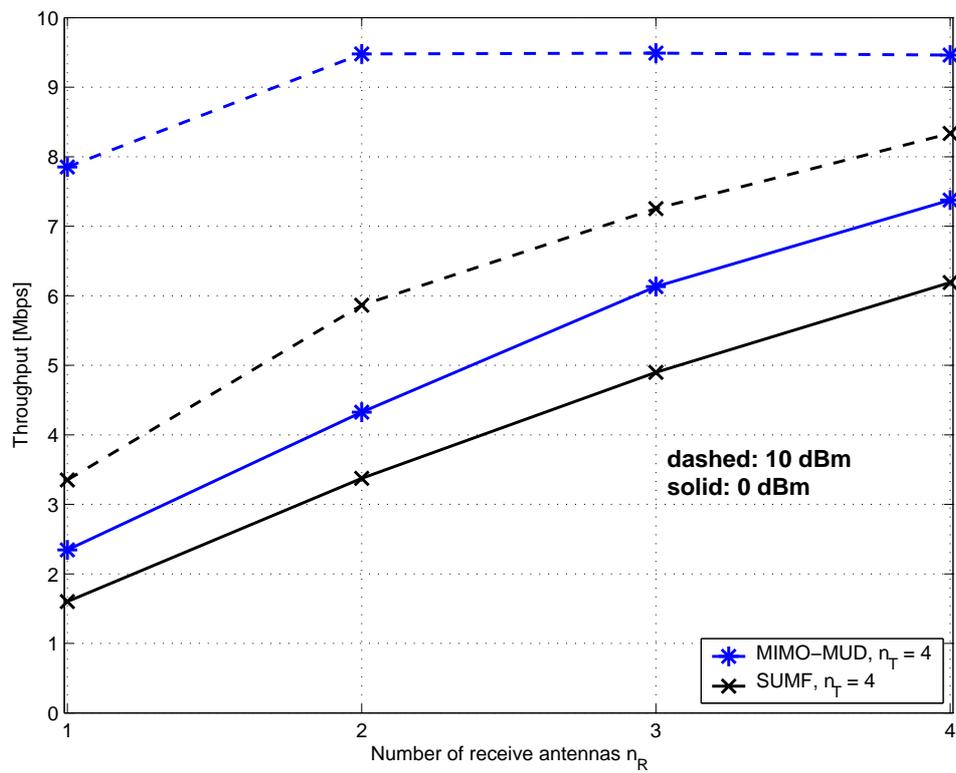


Figure 9: Measured average throughput vs. n_R at 0 dBm (solid line) and 10 dBm (dashed line) over 30 locations.

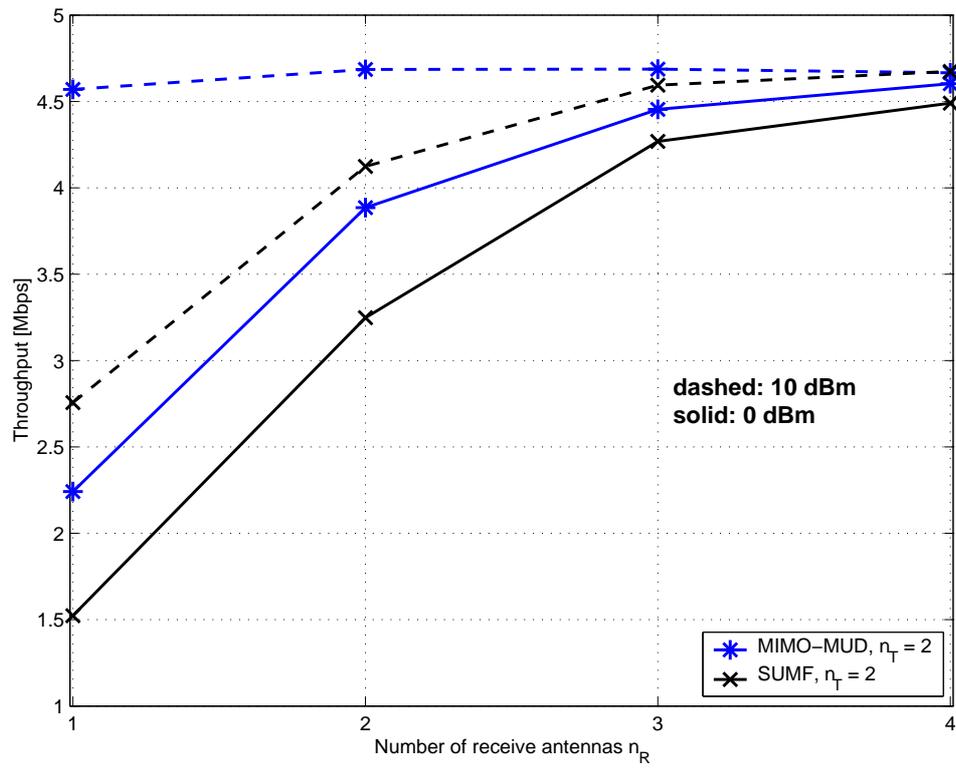


Figure 10: Measured average throughput vs. n_R at 0 dBm (solid line) and 10 dBm (dashed line) over 30 locations.