ACCEPTED FROM OPEN CALL

CELLULAR SYSTEM PHYSICAL LAYER THROUGHPUT: HOW FAR OFF ARE WE FROM THE SHANNON BOUND?

CHRISTIAN MEHLFÜHRER, SEBASTIAN CABAN, AND MARKUS RUPP,
VIENNA UNIVERSITY OF TECHNOLOGY

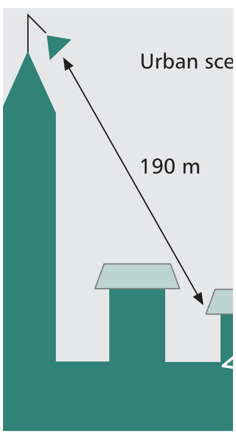

The authors take a closer look at WiMAX and HSDPA, two successful cellular systems that are being operated currently in many countries and check their truly achievable performance.

ABSTRACT

Cellular wireless systems have made impressive progress over the past two decades. They currently connect more than five billion people worldwide. With the advent of turbo decoders in the 1990s, the design of optimal decoders very close to the Shannon bound became possible in AWGN channels. Nowadays many researchers believe that this is true for entire wireless systems and that therefore there is not much left to investigate. In this contribution, we take a closer look at WiMAX and HSDPA, two successful cellular systems being operated currently in many countries, and check their truly achievable performance. We have measured the physical layer throughput of WiMAX and HSDPA in various realistic environments (urban and mountainous) with high-quality equipment and different antenna configurations. We furthermore compared the throughput measured to the Shannon bound. Based on our measurements, we analyze the losses in design and implementation (e.g., pilots, guard carriers, coding, equalization, and channel estimation) and report our findings. Surprisingly, we are currently only utilizing 40 percent of the available channel capacity; or roughly equivalently, we are up to 10 dB off the Shannon bound at typical operational points, thus providing a lot of potential improvement for future 4G systems. In advanced four transmit antenna configurations of HSDPA, the losses are even more pronounced, showing that our current standards are not well suited to take advantage of the much higher capacity provided.

This work has been funded by the Christian Doppler Laboratory for Wireless Technologies for Sustainable Mobility, the Institute of Communications and Radio Frequency Engineering, and KATHREIN-Werke KG as well as A1 Telekom Austria AG. The financial support of the Federal Ministry of Economy, Family and Youth and the National Foundation for Research, Technology and Development is gratefully acknowledged. We further thank Gottfried Lechner for designing the LDPC code for WiMAX and Christoph Mecklenbräuker for his profound leadership.

INTRODUCTION

In 1948, Claude Shannon provided us with a clear channel capacity bound for point-to-point connections [1]. This bound was extended independently in the mid-1990s by Gerry Foschini and Mike Gans [2] as well as by Emre Telatar [3] to the case of multiple transmit and receive antennas (multiple-input multiple-output, MIMO). It is this bound measured in small frequency bins over the utilized bandwidth (Eq. 1 ahead) to which we refer and compare our throughput results. It is this bound that coding experts compare to their turbo codes as well as their iterative receivers. Note, however, that this bound is a single user bound, even if it supports several transmit and receive antennas. Considering multiple users and/or interference is much more difficult. In many situations, the definition of capacity alone is still under discussion. Furthermore, the Shannon bound describes aspects of the physical layer system, not the scheduling techniques, network routing issues, or application layer facets. It is thus far from describing the entire third-generation (3G) system performance, but merely an important piece: the channel.

Nowadays, mathematical models are very common for predicting the performance of systems designed by either mankind or nature. We apply mathematical models to predict tomorrow's weather, our climate 100 years from now, and the performance of next-generation wireless systems. Clearly, models reflect the behavior of the system under investigation. Depending on how precisely we know such a modeled system or subsystem, our prediction models can be more or less accurate. Higher accuracy typically comes with the price of parameter and complexity increase, yet the beauty of a model lies in its simplicity. Once they contain hundreds of parameters and require many hours of simulation time for each new parameter setup, models become less and less treatable. Unfortunately, wireless 3G systems are very complex systems with hundreds of parameters. The channel, a significant part of wireless systems, is still a hot topic for

 1536-1284/11/$25.00 © 2011 IEEE

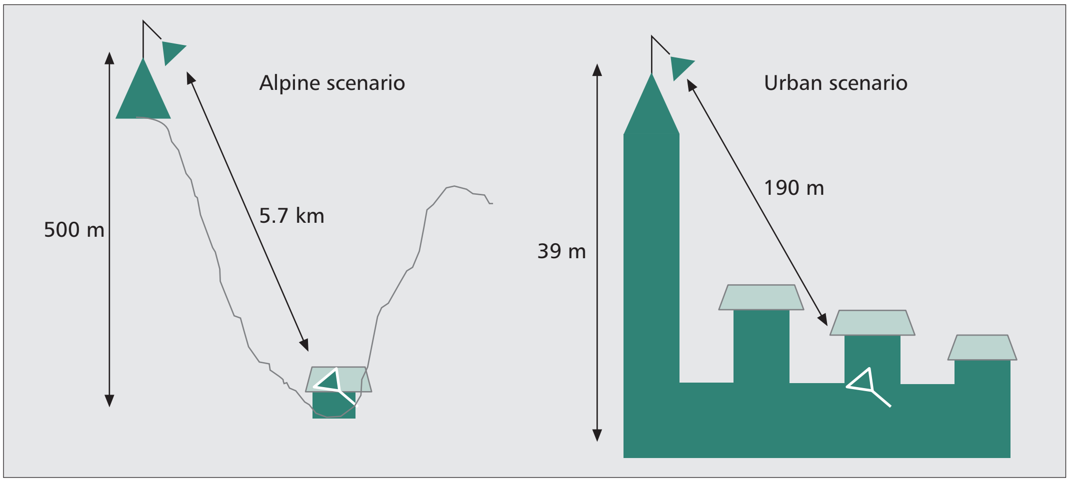

Figure 1. *Schematic view of the two measurement scenarios, left: alpine, right: urban.*

Eventually, measurements are more honest as they exclude assumptions in general. If you do not know, for example, the channel, the frequency offset, or the noise variance, this is natural for a measurement, but not so for a simulation.

investigation. A recent channel model (e.g., Winner Phase 2+) may contain more than 100 parameters to support a multitude of situations with extreme flexibility. If not an expert in modeling, such complex channel models can easily be mistreated and wrong results generated. We therefore based our investigations on measurements.

Compared to simulations, measurements are tedious. On the other hand, each measurement scenario is a true, realistic, and physically correct scenario. With even more effort, measurements can be made repeatable, although not easily reproducible.[1] In order to minimize the effect of our measurement hardware on the results, we employed rather expensive self-built measurement equipment with high linearity, a large range of operation, and a very low noise figure. In the digital domain, we employed receiver structures with the highest performance still feasible for real-time implementation. Our linear minimum mean square error (LMMSE) equalizers are computed in 64-bit floating point MATLAB with a length of 60 taps, certainly in quality much higher than what we can expect in commercial low-cost fixed-point equipment. In high-speed downlink packet access (HSDPA) we also applied successive interference cancellation of the synchronization and pilot signals as well as eight iterations of the turbo channel decoder in order to obtain the best results. Due to cost limitations, however, such optimized receivers are unlikely to be employed in commercial 3G systems. Therefore, the measured implementation loss can be interpreted as a *truly lower bound for commercial systems*.

Eventually, measurements are more honest as they exclude assumptions in general. If you do not know, for example, the channel, the frequency offset, or the noise variance, this is natural for a measurement, but not so for a simulation. More severe, unknown impairments (of unknown sources) only occur in measurements, never in simulations.

This article is organized as follows. We briefly explain our measurement setup and methodology; the interested reader can find more details in the references provided. We provide a very brief introduction to MIMO WiMAX and HSDPA. We explain the transmit precoding as well as the receiver processing as it is important for understanding and interpreting the results. The major part of the article in which we present the performance losses observed for WiMAX and HSDPA in single-input single-output (SISO) as well as 2 × 2 MIMO scenarios. For HSDPA we also present some advanced studies employing up to four antennas. We conclude the article with some final comments.

Measurement Methodology and Setup

Several scenarios have been measured by us over the past years; two of them are compared in this article: an alpine scenario and an urban one (Fig. 1). We picked those two scenarios as they both reflect existing cells in a Carinthian Alpine valley and in downtown Vienna, Austria.

In the alpine environment, the transmitter was placed 5.7 km from the receiver with an essentially strong line-of-sight (LOS) component and moderate scattering. The root mean square (RMS) delay spread of the channels was 260 ns (corresponding to roughly one chip duration of HSDPA). The other scenario was an urban measurement with only 190 m distance but without an LOS component. The scattering was much richer, the RMS delay spread 1.1 μs. Other scenarios, not presented in this article but measured for reference purposes, had a characteristic performance in between the alpine and urban scenarios.

Although many more situations were investigated, we present here only the results for SISO and 2 × 2 MIMO with cross-polarized transmit antennas. In addition, we present results of a four-transmit-antenna measurement of advanced HSDPA to provide a flavor of what more antennas can offer. We employed commercially available cross-polarized base station antennas (Kathrein 80010543), as they support the highest diversity at the smallest size, and are thus of great interest to providers. Measurements with equally polarized antennas typically provided much smaller capacity and correspondingly smaller throughput. Our receiver has a noise figure of 1.9 dB,

[1] For simulations this seems to be a very simple task; still, simulation code is rarely provided in the signal processing community, and reproducibility by others is hampered [4].

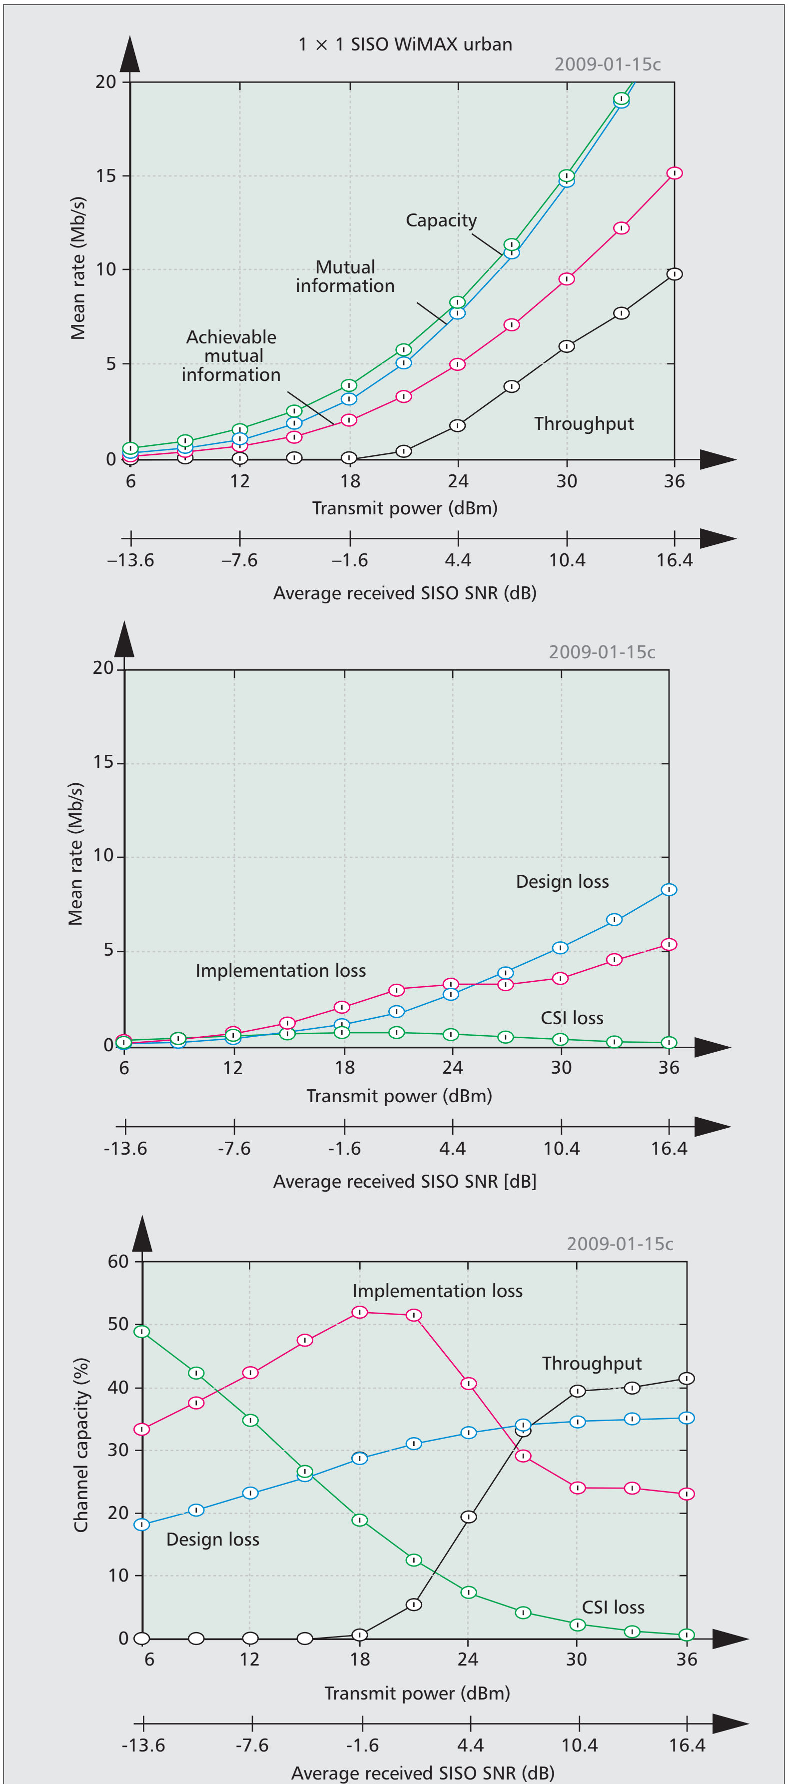

Figure 2. *Example of measured performance metrics (left) and derived throughput losses, absolute (middle) and relative (right).*

which increases by the attenuation of the feeder cable (1 dB/m cable length) to a total of approximately 4 dB. We operate our 5 MHz transmit signal in 6.25 MHz available spectrum to facilitate sharp spectral filtering. Our power amplifiers are capable of transmitting at a power of 30 W (per antenna) of which we use less than 15 percent to ensure high-linearity transmissions.

Typically, the performance of wireless systems is evaluated by drive tests (i.e., measuring the throughput in a vehicle continuously while driving along a road). Unfortunately, the same test repeated at another day can result in entirely different outcomes. In order to make our measurements repeatable we apply a measurement technique employing XY positioning tables on which we automatically change the antenna positions and measure many locations in an area of 3λ × 3λ. At each antenna position in the same scenario, we perform repeated measurements of all transmission schemes of interest. By such a setup we ensure that all schemes are measured under exactly the same conditions. The averaged *measurement results* over all antenna positions later reflect an estimated mean value for the throughput at a given transmit power P_{Tx}. Similarly, this is also the method by which the channel coefficients were evaluated. We measure a set of channels at highest transmit power P_{Tx} (e.g., at 36 dBm; Fig. 2) and evaluate the channel coefficients by the best estimator we have available (LMMSE) utilizing all data (not just the pilots). This provides us with an estimate of the highest quality, certainly an order of magnitude higher in quality than the estimator later employed based on the pilots only. Taking these channel estimates, we then calculated the channel capacity according to Eq. 1 and averaged it to obtain the average channel capacity for a given signal-to-noise ratio (SNR), which is plotted in the figures. Similarly, we proceeded for mutual information[2] according to Eq. 2 and further constrained measures such as achievable mutual information according to Eq. 3. In addition, we calculated the 99 percent confidence intervals to gauge the precision of our estimates. Taking the measurement with highest SNR and all symbols for estimating the channel ensured that the level of uncertainty observed is small compared to the results obtained.

We compute for each measurement point the 99 percent confidence interval by bootstrapping methods and plot the point augmented on a small scale (e.g., Figs. 2 and 3) according to the size of the achieved confidence interval. Most of these only appear as points, indicating extremely high precision. See [8] for a more detailed description of the measurement procedure employed.

CELLULAR SYSTEMS INVESTIGATED: WIMAX AND HSDPA

WIMAX AND HSDPA

We implemented, measured, and analyzed two standardized 3G cellular systems:

- The WiMAX physical layer, as defined in IEEE 802.16-2004, Section 8.3: This standard

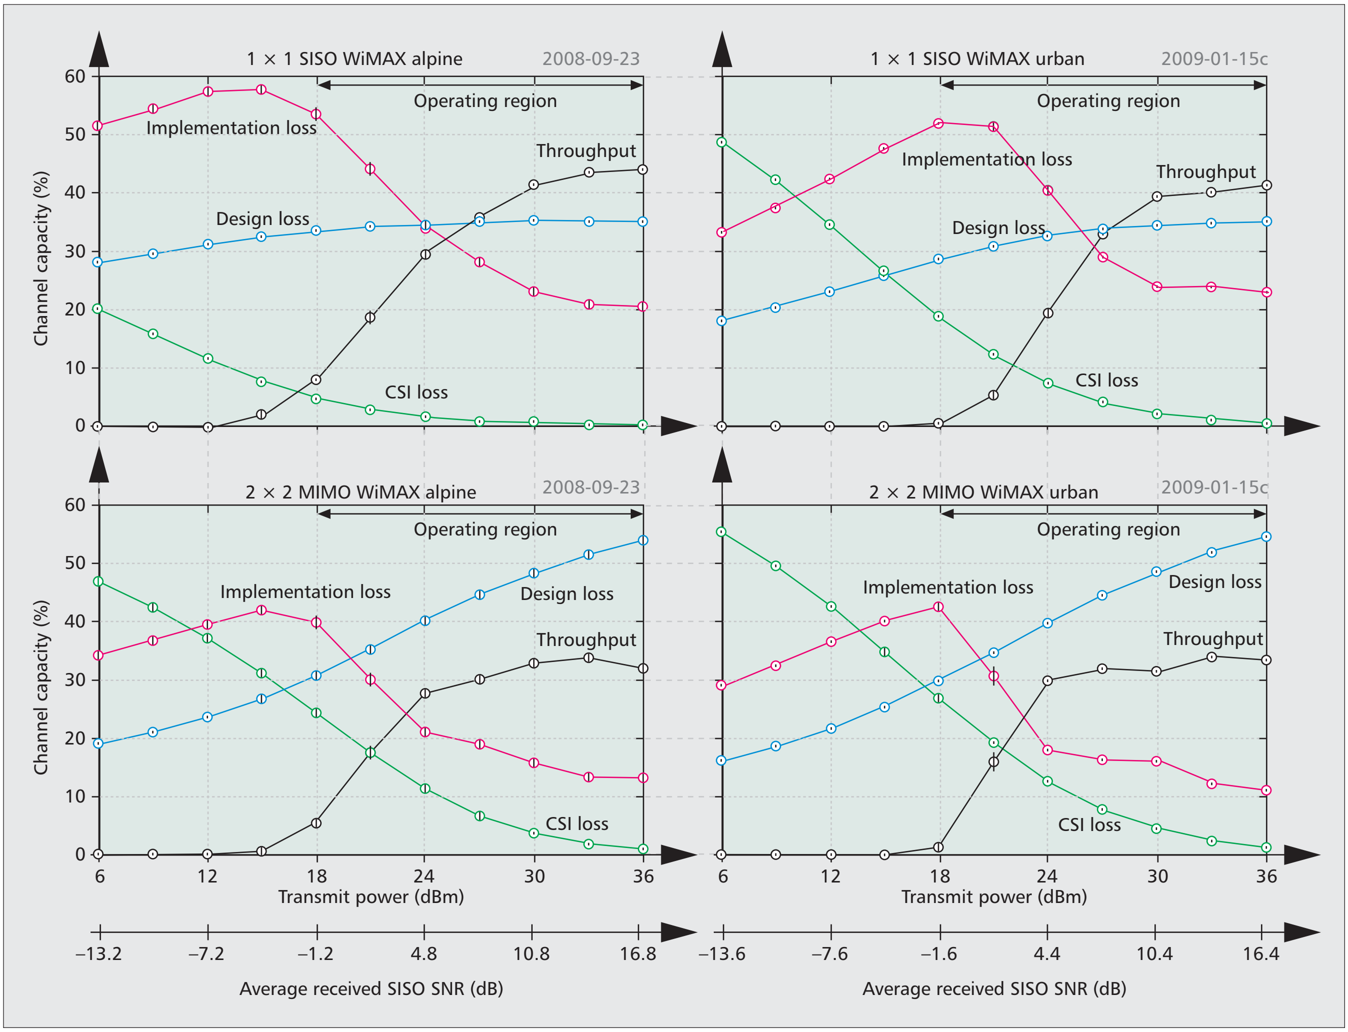

Figure 3. *Relative losses in WiMAX 1 × 1 and 2 × 2 transmissions: left, alpine; right, urban; upper, SISO; lower, MIMO.*

was developed to provide wireless Internet access to stationary and low-mobility users [9]. In our measurements, we employed an orthogonal frequency-division multiplexing (OFDM) physical layer with 256 narrowband subcarriers. By choosing one out of seven adaptive modulation and coding (AMC) schemes at the transmitter, the data rate is adjusted to the current channel conditions, thereby maximizing the data throughput. The standard defines various channel codes out of which we selected for our evaluations the mandatory Reed-Solomon convolutional code (RS-CC) and the optional convolutional turbo code (CTC). We furthermore implemented and measured a regular low-density parity check (LDPC) channel code, which is not defined in the standard. Unless stated otherwise, all our results refer to the CTC. In order to utilize transmit diversity, the standard furthermore foresees simultaneous transmission on two transmit antennas by Alamouti space-time coding.

•The HSDPA mode of the Universal Mobile Telecommunications System (UMTS) [10]: The first version of HSDPA was introduced in Release 5 of UMTS to provide high data rates to mobile users. This is achieved by several techniques such as fast link adaptation, fast hybrid automated repeat request, and fast scheduling. In contrast to the pure transmit power adaptation performed in UMTS, fast link adaptation in HSDPA adjusts the data rate and the number of spreading codes depending on so-called channel quality indicator (CQI) feedback. MIMO HSDPA, standardized in Release 7 of UMTS, further increases the maximum downlink data rate by spatially multiplexing two independently coded and modulated data streams. Additionally, channel-adaptive spatial precoding is implemented at the base station. This is achieved by a standardized set of precoding vectors from which one vector is chosen based on so-called precoding control indication (PCI) feedback obtained from the user equipment.

Throughput Bounds and System Losses

In this section, we define several bounds for the data throughput. The differences between the bounds can be considered as system losses.

In 1998, Foschini and Gans [2] as well as Telatar [3] extended the Shannon capacity C to MIMO systems. For a system working on discrete frequencies[3] $k = 1 \dots K$, the capacity C as a function of the transmit power P_{Tx}, the channel matrix $\mathbf{H}_k \in \mathbb{C}^{N_R \times N_T}$ at the kth frequency bin with bandwidth B/K, the entire channel band-

[2] *Note that there are many different terms in literature describing a constrained capacity, the most prominent being non-coherent capacity [5], functional capacity [6], and simply capacity without CSI at the transmitter [7].*

[3] *We show the formulas in terms of discrete frequency bins* k = 1…K, *as this is very natural for OFDM systems. In HSDPA an equivalent discrete Fourier transform (DFT) on the channel impulse responses has been applied to calculate frequency domain channel coefficients, which in turn are required to compute capacity and mutual information. As the channels have approximately the same bandwidth, the resulting channel capacities match.*

For a MIMO receiver the situation is more complicated, as one antenna may consistently receive less energy than another; we average over all individual SISO channels, or four channels for a 2 × 2 MIMO system.

width B, the receiver noise variance σ_n^2, and the number of transmit antennas N_T is given by

$$C(P_{Tx}) = \max_{\sum \mathrm{tr}\{\mathbf{R}_k\} \le K} \frac{B}{K} \sum_{k=1}^{K} \log_2 \det\left(\mathbf{I} + \frac{P_{Tx}}{\sigma_n^2 N_T} \mathbf{H}_k \mathbf{R}_k \mathbf{H}_k^H\right) \quad (1)$$

A transmission system that is designed to achieve the channel capacity $C(P_{Tx})$ has to perform frequency-selective and spatial precoding according to the matrices $\mathbf{R}_k$. For specific channel realizations, solutions for $\mathbf{R}_k$ are found by the waterfilling algorithm. Such a precoding, of course, only works if a priori channel state information (CSI) is available at the transmitter. If this is not the case, the best strategy is to transmit with equal power over all transmit antennas and over all frequencies. The throughput of a system following this strategy is bounded by the mutual information (MI) $I(P_{Tx})$:

$$I(P_{Tx}) = \frac{B}{K} \sum_{k=1}^{K} \log_2 \det\left(\mathbf{I} + \frac{P_{Tx}}{\sigma_n^2 N_T} \mathbf{H}_k \mathbf{H}_k^H\right) \quad (2)$$

Since the difference between the channel capacity $C(P_{Tx})$ and the mutual information $I(P_{Tx})$ is only caused by the absence of CSI at the transmitter side, we define the difference between them as the *absolute CSI loss* $L_{CSI}(P_{Tx})$ and the *relative CSI loss* $L_{CSI\%}(P_{Tx})$:

$$L_{CSI}(P_{Tx}) = C(P_{Tx}) - I(P_{Tx});$$
$$L_{CSI\%}(P_{Tx}) = 100 \cdot \frac{C(P_{Tx}) - I(P_{Tx})}{C(P_{Tx})} \quad (3)$$

The definitions of capacity and mutual information do not take specific constraints of current cellular systems into account. For example, current systems cannot utilize the whole frequency band and/or the whole transmit power for transmitting information bits. Parts of the spectrum and/or parts of the transmit power have to be allocated to pilot, control, and synchronization information. This is reflected in the measure $I_a(P_{Tx})$, which we call the *achievable mutual information*:

$$I_a(P_{Tx}) = \max_{\mathbf{W} \in \mathcal{W}} \frac{\beta B}{K} \sum_{k=1}^{K} \log_2 \det\left(\mathbf{I} + \frac{\alpha P_{Tx}}{\sigma_n^2 N_T} \mathbf{H}_k \mathbf{W} \mathbf{W}^H \mathbf{H}_k^H\right) \quad (4)$$

The parameter $\alpha \le 1$ accounts for transmit power losses, and the parameter $\beta \le 1$ accounts for spectrum and transmission time losses. In our setup, we find $\alpha = 1$ and $\beta \approx 0.65$ for WiMAX and $\alpha \approx 0.4$ and $\beta \approx 0.77$ for HSDPA. The matrix $\mathbf{W}$ accounts for non-optimal precoding/space-time coding. In the case of HSDPA, the precoding matrix $\mathbf{W}$ is strongly quantized and chosen adaptively out of a predefined codebook $\mathcal{W}$. Depending on the transmission mode, the inclusion of the precoding $\mathbf{W}$ in the calculation of the achievable mutual information may or may not (e.g., if the matrices $\mathbf{W}$ are chosen unitary) impact $I_a(P_{Tx})$. Note that in HSDPA only one precoding matrix $\mathbf{W}$ is chosen for the entire bandwidth; future standards, such as Long Term Evolution (LTE), allow for frequency-dependent precoding, thereby increasing the achievable mutual information. In the case of WiMAX, the transmitter employs Alamouti space-time coding, which can be rewritten as a constant precoding.[4] In contrast to the channel capacity and MI, the achievable MI is therefore a capacity constrained by the specific cellular standard.

We define the difference between the MI $I(P_{Tx})$ and the achievable MI $I_a(P_{Tx})$ as the *design loss*, as it quantifies the loss imposed by the system design:

$$L_d(P_{Tx}) = I(P_{Tx}) - I_a(P_{Tx});$$
$$L_{d\%}(P_{Tx}) = 100 \cdot \frac{I(P_{Tx}) - I_a(P_{Tx})}{C(P_{Tx})} \quad (5)$$

The design loss accounts for the inherent system design losses caused by, for example, losses due to the transmission of pilot and synchronization symbols, quantized precoding, or suboptimal space-time coding. By alternating design parameters, which can partially also be achieved adaptively on the transmission scenario, the design loss can be reduced.

Finally, there is the throughput $D_m(P_{Tx})$ that can be measured in bits per second in a given link, allowing the so-called *implementation loss* to be defined,

$$L_i(P_{Tx}) = I_a(P_{Tx}) - D_m(P_{Tx});$$
$$L_{i\%}(P_{Tx}) = 100 \cdot \frac{I_a(P_{Tx}) - D_m(P_{Tx})}{C(P_{Tx})} \quad (6)$$

as the difference between the achievable mutual information $I_a(P_{Tx})$ and the measured data throughput $D_m(P_{Tx})$. It accounts for losses caused by non-optimum receivers and channel codes. A more detailed discussion of the individual parts of the implementation loss in WiMAX or HSDPA is provided further ahead.

In order to complete the picture we also need to define the *relative throughput*,

$$D_{m\%}(P_{Tx}) = 100 \cdot \frac{D_m(P_{Tx})}{C(P_{Tx})}. \quad (7)$$

Summing up the relative losses and relative throughput results in 100 percent. We display the relative losses together with the relative throughput.

In Fig. 2, all terms defined above are shown in their absolute and relative measures for a SISO WiMAX transmission. Comparing the CSI loss to the other two losses, it is almost negligible, especially at higher transmit powers and thus higher SNR [11]. Note that in the plot of relative losses, the CSI loss appears as a monotone decreasing function over transmit power with high values (50 percent) at the lower end. However, at very low transmit power (approximately –10 dB SNR) the capacity is extremely small, and losing 50 percent does not make a noticeable performance difference. The initial loss of 50 percent varies depending on the offered selectivity in frequency and spatial domain, as shown later. The more diversity offered, the larger the CSI loss.

In the literature, we find the measures defined above to be a function of the receive

[4] *This is to be seen in contrast to the option of including space-frequency Alamouti coding.*

SNR. Note, however, that in modern 3G systems such a definition is obsolete. There are three reasons, the first being that we measure in terms of transmit power P_{Tx}, and comparing two systems, we should compare them at equal transmit power and not at equal receive SNR. Once channel adaptive precoding is introduced, the receive power is systematically increased, also increasing the receive SNR without changing P_{Tx}, providing a second reason. Third, depending on the scenario and the construction, one receive antenna will generally receive a different amount of power than another. Such differences can be as large as 6 dB in the mean. The reason for such differences is the scattering environment not being symmetric in nature, as has been assumed in most earlier MIMO channel models. The Winner Phase 2+ model is the first to allow for defining such asymmetries. We thus have to plot over transmit power P_{Tx} rather than SNR. Nevertheless, in order to facilitate comparisons to previous simulation results, we provide two scales in our plots, one being P_{Tx} and the other a computed equivalent average SISO receiver SNR. In the case of a SISO receiver this is simply the average value obtained over the individual SISO channels. For a MIMO receiver the situation is more complicated, as one antenna may consistently receive less energy than another; we average over all individual SISO channels, or four channels for a 2 × 2 MIMO system.[5]

Measurement Results

WiMAX Results

Now that we have explained how we obtained the data, we present them decomposed into their various relative losses. Figure 3 shows our measurement results for the selected scenarios.

The reader may recognize the upper right subfigure of Fig. 3 being identical to the bottom subfigure of Fig. 2.

• The *relative CSI losses* are monotonically decreasing functions. For moderate to high transmit power values P_{Tx}, the impact of missing CSI at the transmitter can be neglected, questioning why channel information beyond the receiver SNR is useful for the transmitter.

• The *relative design losses* show a gradually increasing behavior with P_{Tx}. This behavior is more pronounced in MIMO scenarios, in which Alamouti space-time coding plays a crucial role. In the case of 2 × 2 MIMO transmission, the Alamouti space-time code is suboptimal; hence, the achievable mutual information is limited by the system design. Especially at large transmit power levels, the loss due to Alamouti space-time coding becomes dominant, wasting approximately 50 percent of the available channel capacity.

• For lower to moderate P_{Tx}, the *relative implementation losses* are severe. Only at high P_{Tx} do they fall below the design losses. The reason for this is found in the measured throughput. As at low P_{Tx} there is no AMC scheme supporting the transmission, the throughput is zero for SNR values lower than approximately –3 dB (equivalent to 15 dBm transmit power).

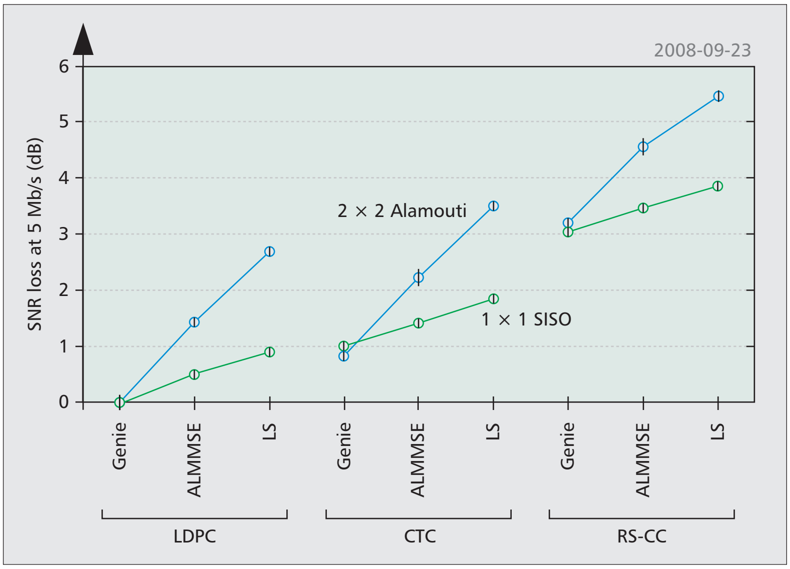

Figure 4. *SNR losses of different channel estimators and channel coding schemes with respect to genie-driven channel estimation and LDPC coding, measured in the alpine scenario.*

Until there, the implementation loss increases with the capacity. Once the throughput begins to rise, the relative implementation loss drops to 10–20 percent at high SNR and then remains roughly constant.

• The *relative throughput* appears to be a monotonically increasing function. However, at high P_{Tx}, the AMC schemes run out and no higher rate can be transmitted. Consequently, the relative throughput starts decreasing. As our maximum transmit power was 36 dBm, we hardly reached this area. Only in the alpine MIMO scenario with a strong LOS component can we observe such a decline at high P_{Tx}. Surprisingly, the relative throughput is lower in MIMO (30 percent) than in SISO (40 percent), even though the absolute values are much higher in MIMO, showing that this version of WiMAX is not capable of taking advantage of what MIMO offers.

In summary, we conclude that WiMAX behaves quite similar in alpine and urban environments as well as in SISO and MIMO. Comparing SISO with MIMO performance, we can only conclude that the various components show larger variations in SISO and less in MIMO.[6]

Let us now take a closer look at the implementation loss. In our displayed measurements we employed the CTC with an advanced channel estimation scheme (ALMMSE). Figure 4 depicts the SNR loss of various channel estimation and channel coding schemes at a throughput of 5 Mb/s when compared to a genie channel estimator and LDPC channel coding. The genie channel estimator knows not only the 200 pilot symbols, but all 9400 transmitted data symbols and uses them to achieve a very high channel estimation quality. Similarly, a regular LDPC code was designed to provide a high-quality channel code for a given SNR. Figure 4 reveals that between genie knowledge of the channel in combination with high-quality channel coding on one hand and poor LS channel estimation in combination

[5] *We have provided such "equivalent" SISO SNR to facilitate comparison by researchers working with receive SNR and simplified channel models.*

[6] *When compared to absolute loss values (not displayed here, see [12]), the results are also quite the same for all WiMAX scenarios. The largest loss is the design loss, monotonically increasing over P_{Tx}, followed by the implementation loss, which is dominant at low P_{Tx} but only small at large P_{Tx}. In the regions of interest, the CSI loss is the smallest of the losses, only showing its existence in MIMO conditions at low to moderate P_{Tx}.*

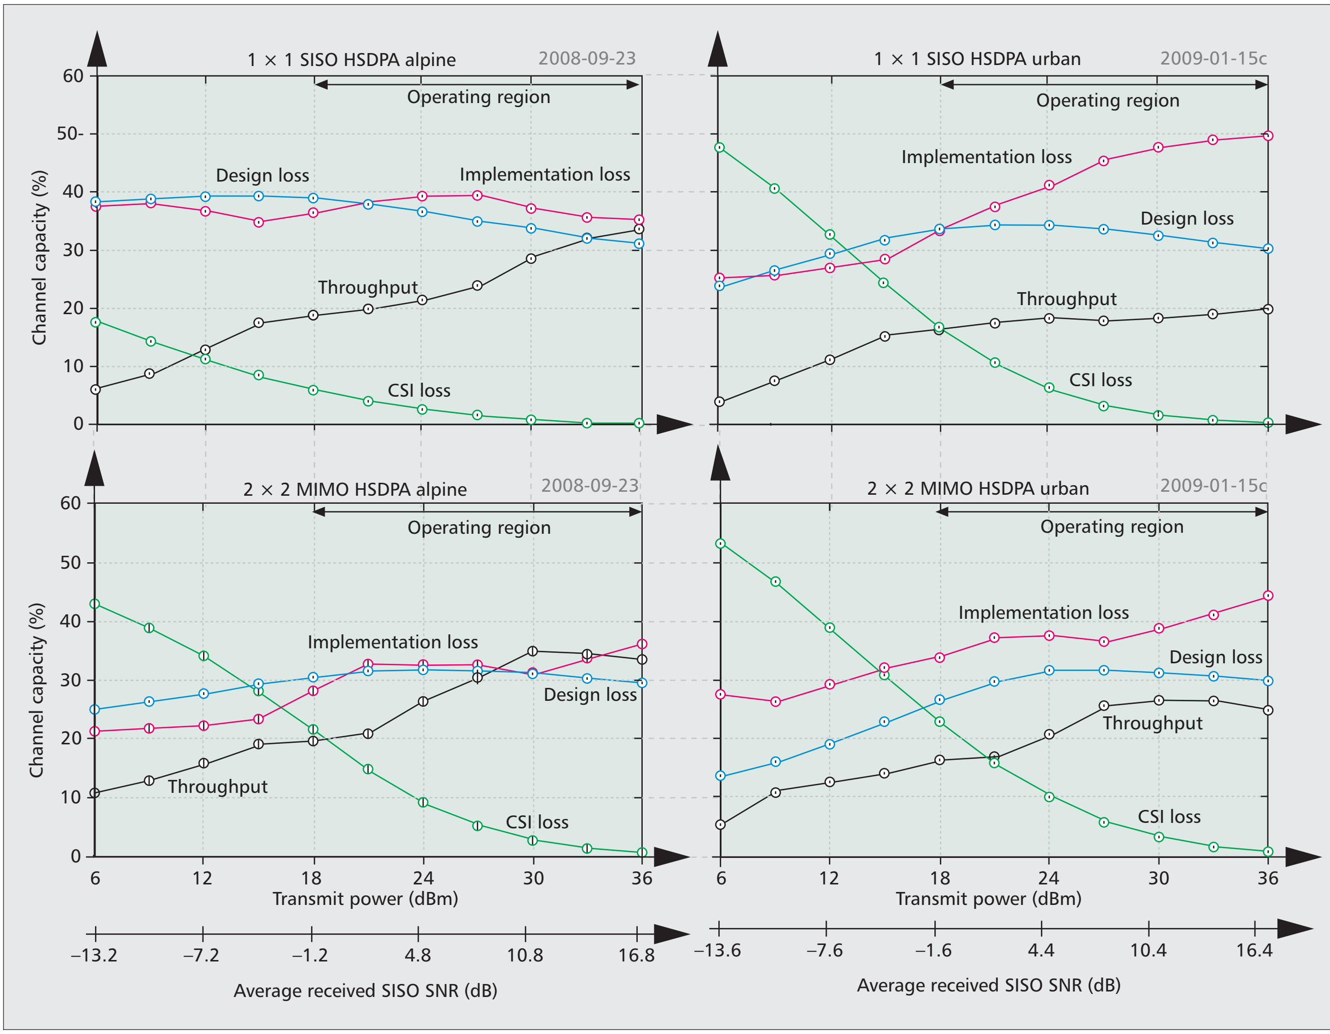

Figure 5. *Relative losses in HSDPA 1 × 1 and 2 × 2 transmissions: left, alpine; right, urban; upper, SISO; lower, MIMO.*

with the mandatory RS-CC channel coding of WiMAX on the other hand can be more than 6 dB distance. In the case of 2 × 2 transmissions, when Alamouti space-time coding is employed at the transmitter, we observe a much higher dependence on the type of channel estimator chosen. This can be explained by two facts:

• In the case of Alamouti space-time coded transmission, the available transmit power and thus also the training signal power are equally distributed on the two transmit antennas. Therefore, only half the training signal power is available per channel coefficient to be estimated. As a consequence, the channel estimation performance is poorer than in the one transmit antenna case.

• The comparison in Fig. 4 is carried out at a constant throughput value of 5 Mb/s. The Alamouti space-time coded transmission achieves this throughput at a much lower receive SNR than the SISO transmission. This lower SNR in turn causes poorer channel estimation quality and thus substantial SNR loss in comparison to the genie channel estimator.

In order to compensate for the above effects, one would have to apply a much better channel estimator for MIMO than in the case of SISO transmissions.

HSDPA Results in Standard-Compliant Setting

All measurement results presented in this section are for a Category 16 HSDPA user equipment. In the MIMO case, D-TxAA with adaptive precoding is applied [13].

The various losses in HSDPA (Fig. 5) show a much less lively picture when compared to WiMAX.

• As before, the relative CSI loss is decreasing monotonically in exactly the same dimensions, as the CSI loss is independent of the transmission standard, proving that both measurements have experienced the same equipment as well as the same wireless conditions. Small differences when compared to the WiMAX results are due to a slightly different occupied bandwidth.

• The other losses are more or less constant functions or slightly increasing with P_{Tx}. The relative design loss with values between 30 and 40 percent is a particularly flat curve when compared to WiMAX. Only at low P_{Tx} in the urban MIMO scenario the values can become as small as 15 percent .

• The *relative implementation loss* is either of same value than the design loss or larger. The distinct behavior in the urban scenario is of

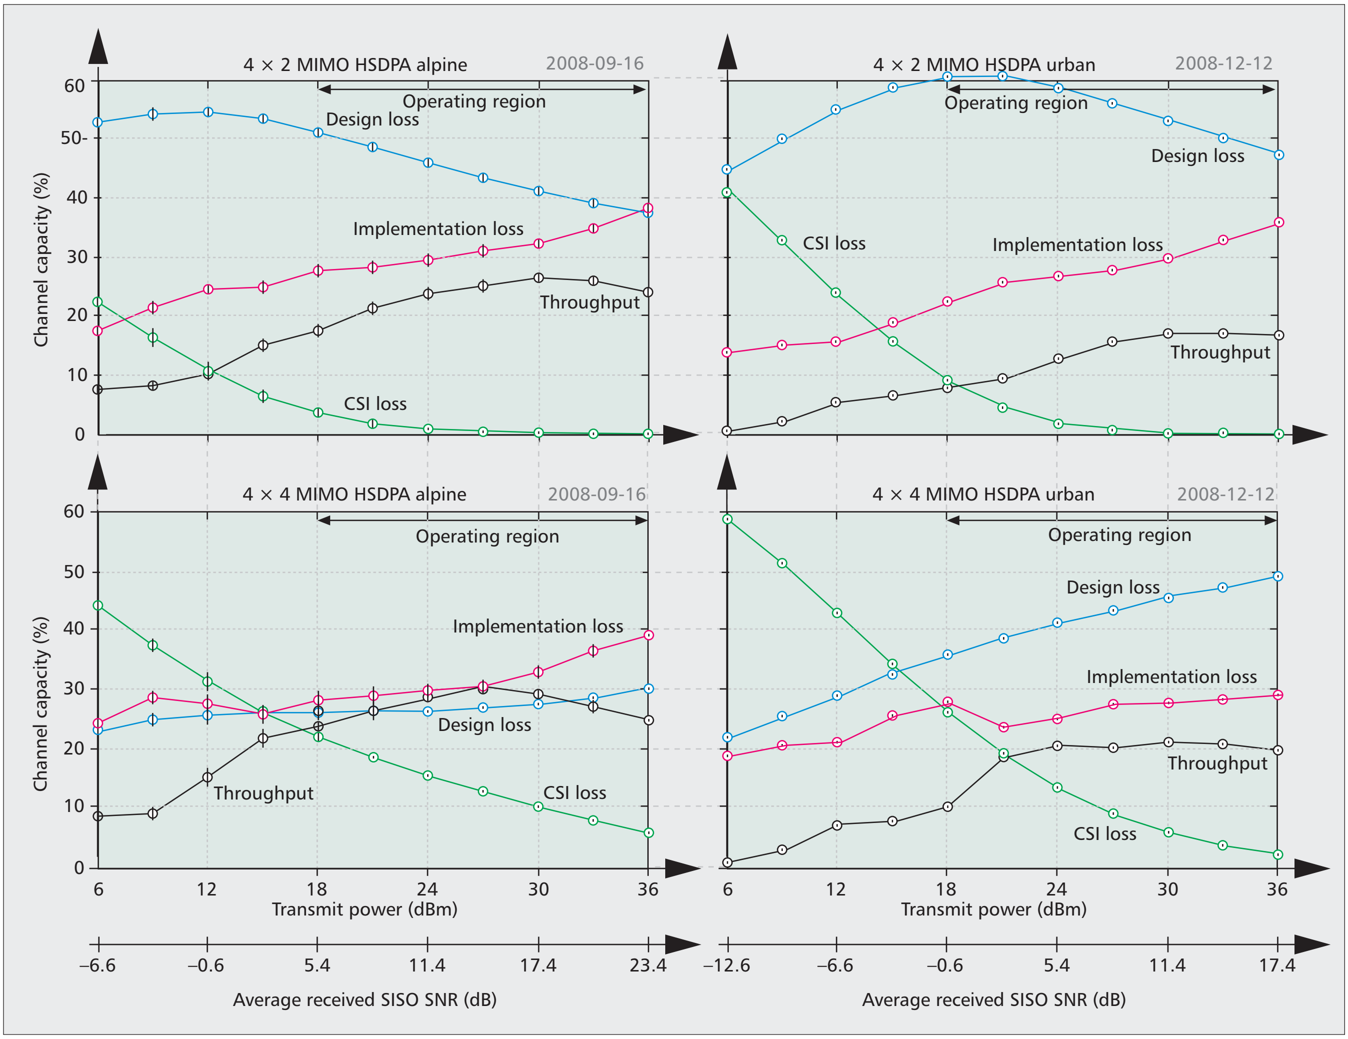

Figure 6. *Relative losses in HSDPA 4 Tx antenna transmissions: left, alpine; right, urban; upper, 4 × 2 MIMO; lower, 4 × 4 MIMO.*

interest. Due to the larger RMS delay spread in the urban scenario, the inter code interference is increased and becomes a dominant part at high P_{Tx}. Therefore, we recognize a higher implementation loss and a smaller design loss for high P_{Tx}.

• Similar to WiMAX, the *relative throughput* is an increasing function in SNR. For low values of P_{Tx} the spreading functions of HSDPA improve the situation considerably, allowing a transmission even for very low receiver SNR at low bit rates. Furthermore, HSDPA has many more AMC schemes than WiMAX, especially at low SNR. At high P_{Tx} we obtain in the alpine environment with its strong LOS the close to 40 percent values in relative throughput just as for WiMAX. However, in the urban scenario with high RMS delay spread the relative throughput is much lower as the transmitter is producing strong inter code interference in such channels, visible in SISO as well as in MIMO transmissions.

Although the results in absolute values [13] show a considerable performance increase of the different MIMO schemes when compared to the SISO transmission, all measured throughput curves are losing between 3 and 9 dB SNR compared to the achievable mutual information. The following effects contribute (next to maybe others) to this loss:

• The rate-matched turbo code utilized in HSDPA is good but not optimal. By carrying out a set of comparative additive white Gaussian noise (AWGN) simulations, we found that at higher code rates, the rate-matched turbo code loses up to 2 dB when decoded by a max-log-MAP decoder (approximately 2 dB).

• The LMMSE equalizer representing a low complexity and cost-effective solution is also not optimal. Better receivers such as the LMMSE-MAP have the potential to improve the performance by approximately 1 dB.

• In the urban scenario, a larger throughput loss was measured than in the alpine scenario because of the larger delay spread and, consequently, the larger inter-code interference. For example, in the alpine scenario the SISO system loses approximately 6 dB to the achievable MI, whereas the loss in the urban scenario is approximately 9 dB (3 dB *additional loss in the urban scenario compared to the alpine scenario*).

• In addition to the above mentioned losses, channel estimation errors and over-/underestimation of the post-equalization signal-to-interference-plus-noise ratio (SINR) degrade the measured throughput. Exactly quantifying the loss caused by these effects is difficult because

A system designer should equally take all losses into consideration, but note that the design loss imposed by the system design can never be reduced in the future by a clever receiver implementation.

neither perfect CSI nor perfect post-equalization SINR is available in measurements.

It should be noted that HSDPA supports hybrid automatic repeat request (HARQ) retransmissions, while the chosen WiMAX standard did not support this yet. We indeed implemented retransmissions in HSDPA, but the measurement plots shown here are not utilizing such retransmissions to keep the comparison to WiMAX fair. Note that including up to three retransmissions hardly changed the figures. We conclude that while HARQ offers some advantages in service quality (e.g., lower latency), it has no impact on the throughput performance. This may be due to the fact that in our measurements the block error rate (BLER) is kept at approximately 1 percent throughout the whole measurement range.

HSDPA Results in an Advanced Setting

By employing a straightforward extension of the two transmit antenna precoding, we extended the HSDPA standard toward four antennas at the transmit side and utilized either two or four receive antennas [13]. The results obtained when employing 4 × 2 and 4 × 4 MIMO systems are very different from the ones shown above:

- The relative design loss is now clearly the major loss, dominating the implementation loss in every scenario, while above they were roughly of equal size. This indicates that precoding for four antennas should be implemented with much more sophistication than we did in our experiment.
- The relative throughput has reached 30 percent of the capacity at best. We can thus conclude that although advanced settings on HSDPA improve absolute values of throughput by up to 50 percent, relatively speaking they utilize less of the now much higher capacity.

Conclusions

The measured throughput D_m actually observed accounts for only 15–45 percent of the available channel capacity. In terms of the individual losses, our findings are summarized as:

- **The CSI loss** L_{CSI}: This loss is given by absence of full channel state information at the transmitter. The CSI loss is typically small, and for SISO and cross polarized 2 × 2 MIMO systems it can be neglected compared to the other losses. Only when the transmission system employs more than two antennas and is operating at low SNR is it worth combatting the CSI loss by employing detailed CSI feedback. A potential of 50 percent improvement in capacity may sound appealing; note, however, that the absolute values are rather small in the low SNR regime.
- **The design loss** L_d: Here the differences between HSDPA and WiMAX are not very pronounced. HSDPA is generally better designed as the design loss is roughly constant over the full SNR range, while in WiMAX the design loss is small for low SNRs and increases with SNR. In the case of 2 × 2 MIMO WiMAX this is caused by suboptimal Alamouti space-time coding. Precoding is not expected to have an impact on the design loss, as optimal precoding matrices are unitary, and so are the proposed precoding schemes in the HSDPA and LTE standards. The design loss is the vehicle of the future to achieve improvements at the design stage. Once the standard is released, the design loss is final.
- **The implementation loss** L_i: In environments with high RMS delay spread (urban environment with 1.1 μs), the implementation loss of HSDPA is extremely large due to self-interference, while in small RMS delay spread areas (alpine environment with 260 ns) the behavior is different. In WiMAX the different delay spread has no impact, as the cyclic prefix was selected sufficiently large. Both WiMAX and HSDPA suffer losses of several decibels due to channel estimation and non-optimal coding. The much higher number of AMC schemes in HSDPA does not show much of a benefit compared to WiMAX. A good strategy for standardization bodies could be to accept a rather large initial implementation loss, anticipating that manufacturers will figure out how to improve implementation quality by employing more digital signal processing complexity. However, some implementations are close to optimal, no longer offering much to be gained. Once manufacturers are offering close to optimal implementations, they cannot differentiate their products between each other.

What counts at the end of the day when designing a system is the throughput actually achieved. Therefore, a system designer should take *all* losses into consideration equally, but note that the design loss imposed by the system design can never be reduced in the future by a clever receiver implementation. Therefore, it may pay off to start with a smaller design loss and a larger implementation loss.

References

[1] C. Shannon, "A Mathematical Theory of Communication," *Bell Sys. Tech. J.*, vol., July, Oct. 1948, http://cm.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf27, pp. 379–423, 623–56.

[2] G. J. Foschini and M. J. Gans, "On Limits of Wireless Communications in a Fading Environment When Using Multiple Antennas," *Wireless Pers. Commun.*, vol. 6, no. 3, Mar. 1998, http://www.springerlink.com/content/h1n7866218781520/fulltext.pdf, pp. 311–35.

[3] I. E. Telatar, "Capacity of Multi-Antenna Gaussian Channels," *Euro. Trans. Telecommun.*, 1999, tech. memo, Bell Labs, Lucent Technologies, vol. 10, no. 6, Oct. 1998, http://mars.bell-labs.com/papers/proof/proof.pdf, pp. 585–95.

[4] P. Vandewalle, J. Kovacevic, and M. Vetterli, "Reproducible Research in Signal Processing — What, Why, and How," *IEEE Sig. Proc. Mag.*, vol. 26, no. 3, 2009, pp. 37–47.

[5] G. Durisi *et al.*, "Noncoherent Capacity of Underspread Fading Channels," *IEEE Trans. Info. Theory*, vol. 56, no. 1, Jan. 2010, http://www.nari.ee.ethz.ch/commth/pubs/p/dsbsIT08, pp. 367–95.

[6] J. Andrews *et al.*, "Rethinking Information Theory for Mobile Ad Hoc Networks," *IEEE Commun. Mag.*, vol. 46, no. 12, 2008, pp. 94–101.

[7] A. Tulino, A. Lozano, and S. Verdu, "MIMO Capacity with Channel State Information at the Transmitter," *Proc. ISSSTA*, Sydney, Australia, Aug. 2004, pp. 22–26.

[8] S. Caban *et al.*, "Testbedding MIMO HSDPA and WiMAX," *Proc. VTC '09-Fall*, Anchorage, AK, Sept. 2009, http://publik.tuwien.ac.at/files/PubDat_176574.pdf

[9] A. Ghosh *et al.*, "Broadband Wireless Access with WiMax/802.16: Current Performance Benchmarks and Future Potential," *IEEE Commun. Mag.*, vol. 43, no. 2, 2005, http://ieeexplore.ieee.org/iel5/35/30297/01391513.pdf?tp=&arnumber=1391513, pp. 129–36.

[10] H. Holma *et al.*, "High-Speed Packet Access Evolution in 3GPP release 7," *IEEE Commun. Mag.*, vol. 45, no. 12, pp. 29–35, Dec. 2007, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=4395362
[11] M. Rupp *et al.*, "On Mutual Information and Capacity in Frequency Selective Wireless Channels," *Proc. IEEE ICC '10*, Cape Town, South Africa, May 2010, http://publik.tuwien.ac.at/files/PubDat_184660.pdf
[12] M. Rupp, C. Mehlführer, and S. Caban, "On Achieving the Shannon Bound in Cellular Systems," *Proc. 20th Int'l. Conf. Radioelektronika 2010*, Brno, Czech Republik, Apr. 2010, http://publik.tuwien.ac.at/files/PubDat_185403.pdf
[13] C. Mehlführer, S. Caban, and M. Rupp, "Measurement-Based Performance Evaluation of MIMO HSDPA," *IEEE Trans. Vehic. Tech.*, vol. 59, no. 9, Nov. 2010, http://publik.tuwien.ac.at/files/PubDat_187112.pdf, pp. 4354–67.

Biographies

CHRISTIAN MEHLFÜHRER [M'09] received his Dipl.-Ing. degree in electrical engineering from the Vienna University of Technology. Besides his diploma studies he worked part time at Siemens AG, where he performed integration tests of GSM carrier units. After finishing his diploma thesis on implementation and real-time testing of space-time block codes at the Institute of Communications and Radio Frequency Engineering, Vienna University of Technology, for which he received the Vodafone Förderpreis 2006 (together with Sebastian Caban), he started his doctoral thesis at the same institute. In 2009 he finished his Ph.D. about measurement-based performance evaluation of WiMAX and HSDPA with summa cum laude. His research interests include experimental investigation of MIMO systems, UMTS HSDPA, WLAN (802.11), WiMAX (802.16), and the upcoming 3GPP LTE system.

SEBASTIAN CABAN [M'09] went to a technical high school and received his Bachelor's degree in electrical and communication engineering from the Vienna University of Technology. He received his Master's degree in 2005 and his Ph.D. degree in 2009. In between, he spent a year with the University of Illinois at Urbana-Champaign. He received B.B.A. and M.B.A. degrees, with a major in economics, from the University of Vienna in 2010. He is currently with the Institute of Communications and Radio Frequency Engineering, Vienna University of Technology. His research interests include developing measurement methodologies and building testbeds to quantify the actual performance of wireless cellular communication systems (e.g., UMTS and LTE) in realistic scenarios. He received 12 national scholarships and prizes, and the Förderpreis 2006 from the German Vodafone Stiftung für Forschung for his research.

MARKUS RUPP [SM'06] received his Dipl.-Ing. degree from the University of Saarbrücken, Germany, in 1988 and his Dr.Ing. degree from Technische Universität Darmstadt, Germany, in 1993, where he worked with E. Hänsler on designing new algorithms for acoustical and electrical echo compensation. From November 1993 to July 1995 he held, with S. Mitra, a postdoctoral position with the University of California, Santa Barbara, where he worked with A. H. Sayed on a robustness description of adaptive filters with impact on neural networks and active noise control. From October 1995 to August 2001 he was a member of technical staff with the Wireless Technology Research Department, Bell Laboratories, Crawford Hill, New Jersey, where he worked on various topics related to adaptive equalization and rapid implementation for IS-136, 802.11, and UMTS. Since October 2001 he has been a full professor of digital signal processing in mobile communications with the Vienna University of Technology, where he founded the Christian Doppler Laboratory for Design Methodology of Signal Processing Algorithms at the Institute of Communications and Radio Frequency Engineering in 2002. He served as dean from 2005 to 2007. He is currently an Associate Editor for the EURASIP *Journal of Advances in Signal Processing* and EURASIP *Journal on Embedded Systems*. He is an author or coauthor of more than 350 papers and is the holder of 15 patents on adaptive filtering, wireless communications, and rapid prototyping, as well as automatic design methods. He was an Associate Editor for *IEEE Transactions on Signal Processing* from 2002 to 2005. He has been an Administrative Committee Member of EURASIP since 2004 and served as the President of EURASIP from 2009 to 2010.